

Computing and the National Science Foundation, 1950 - 2016

*Building a Foundation for
Modern Computing*

**Peter A. Freeman
W. Richards Adrion
William Aspray**



ASSOCIATION FOR COMPUTING MACHINERY

**Computing and the
National Science Foundation,
1950–2016**

ACM Books

Editor in Chief

Sanjiva Prasad, *Indian Institute of Technology Delhi*

ACM Books is a new series of high-quality books for the computer science community, published by ACM in collaboration with Morgan & Claypool Publishers. ACM Books publications are widely distributed in both print and digital formats through booksellers and to libraries (and library consortia) and individual ACM members via the ACM Digital Library platform.

Computing and the National Science Foundation, 1950–2016: Building a Foundation for Modern Computing

Peter A. Freeman, *Georgia Institute of Technology*
W. Richards Adrion, *University of Massachusetts Amherst*
William Aspray, *University of Colorado Boulder*
2019

Providing Sound Foundations for Cryptography: On the work of Shafi Goldwasser and Silvio Micali

Editor: Oded Goldreich, *Weizmann Institute of Science*
2019

Concurrency: The Works of Leslie Lamport

Editor: Dahlia Malkhi, *VMware Research* and *Calibra*
2019

The Essentials of Modern Software Engineering: Free the Practices from the Method Prisons!

Ivar Jacobson, *Ivar Jacobson International*
Harold “Bud” Lawson, *Lawson Konsult AB (deceased)*
Pan-Wei Ng, *DBS Singapore*
Paul E. McMahon, *PEM Systems*
Michael Goedicke, *Universität Duisburg–Essen*
2019

Data Cleaning

Ihab F. Ilyas, *University of Waterloo*
Xu Chu, *Georgia Institute of Technology*
2019

Conversational UX Design: A Practitioner’s Guide to the Natural Conversation Framework

Robert J. Moore, *IBM Research–Almaden*
Raphael Arar, *IBM Research–Almaden*
2019

Heterogeneous Computing: Hardware and Software Perspectives

Mohamed Zahran, *New York University*

2019

Hardness of Approximation Between P and NP

Aviad Rubinfeld, *Stanford University*

2019

The Handbook of Multimodal-Multisensor Interfaces, Volume 3: Language Processing, Software, Commercialization, and Emerging Directions

Editors: Sharon Oviatt, *Monash University*

Björn Schuller, *Imperial College London and University of Augsburg*

Philip R. Cohen, *Monash University*

Daniel Sonntag, *German Research Center for Artificial Intelligence (DFKI)*

Gerasimos Potamianos, *University of Thessaly*

Antonio Krüger, *Saarland University and German Research Center for Artificial Intelligence (DFKI)*

2019

Making Databases Work: The Pragmatic Wisdom of Michael Stonebraker

Editor: Michael L. Brodie, *Massachusetts Institute of Technology*

2018

The Handbook of Multimodal-Multisensor Interfaces, Volume 2: Signal Processing, Architectures, and Detection of Emotion and Cognition

Editors: Sharon Oviatt, *Monash University*

Björn Schuller, *University of Augsburg and Imperial College London*

Philip R. Cohen, *Monash University*

Daniel Sonntag, *German Research Center for Artificial Intelligence (DFKI)*

Gerasimos Potamianos, *University of Thessaly*

Antonio Krüger, *Saarland University and German Research Center for Artificial Intelligence (DFKI)*

2018

Declarative Logic Programming: Theory, Systems, and Applications

Editors: Michael Kifer, *Stony Brook University*

Yanhong Annie Liu, *Stony Brook University*

2018

The Sparse Fourier Transform: Theory and Practice

Haitham Hassanieh, *University of Illinois at Urbana-Champaign*

2018

The Continuing Arms Race: Code-Reuse Attacks and Defenses

Editors: Per Larsen, *Immunant, Inc.*

Ahmad-Reza Sadeghi, *Technische Universität Darmstadt*

2018

Frontiers of Multimedia Research

Editor: Shih-Fu Chang, *Columbia University*
2018

Shared-Memory Parallelism Can Be Simple, Fast, and Scalable

Julian Shun, *University of California, Berkeley*
2017

Computational Prediction of Protein Complexes from Protein Interaction Networks

Sriganesh Srihari, *The University of Queensland Institute for Molecular Bioscience*
Chern Han Yong, *Duke-National University of Singapore Medical School*
Limsoon Wong, *National University of Singapore*
2017

The Handbook of Multimodal-Multisensor Interfaces, Volume 1: Foundations, User Modeling, and Common Modality Combinations

Editors: Sharon Oviatt, *Incaa Designs*
Björn Schuller, *University of Passau and Imperial College London*
Philip R. Cohen, *Voicebox Technologies*
Daniel Sonntag, *German Research Center for Artificial Intelligence (DFKI)*
Gerasimos Potamianos, *University of Thessaly*
Antonio Krüger, *Saarland University and German Research Center for Artificial Intelligence (DFKI)*
2017

Communities of Computing: Computer Science and Society in the ACM

Editor: Thomas J. Misa, *University of Minnesota*
2017

Text Data Management and Analysis: A Practical Introduction to Information Retrieval and Text Mining

ChengXiang Zhai, *University of Illinois at Urbana-Champaign*
Sean Massung, *University of Illinois at Urbana-Champaign*
2016

An Architecture for Fast and General Data Processing on Large Clusters

Matei Zaharia, *Stanford University*
2016

Reactive Internet Programming: State Chart XML in Action

Franck Barbier, *University of Pau, France*
2016

Verified Functional Programming in Agda

Aaron Stump, *The University of Iowa*
2016

The VR Book: Human-Centered Design for Virtual Reality

Jason Jerald, *NextGen Interactions*

2016

Ada's Legacy: Cultures of Computing from the Victorian to the Digital Age

Robin Hammerman, *Stevens Institute of Technology*

Andrew L. Russell, *Stevens Institute of Technology*

2016

Edmund Berkeley and the Social Responsibility of Computer Professionals

Bernadette Longo, *New Jersey Institute of Technology*

2015

Candidate Multilinear Maps

Sanjam Garg, *University of California, Berkeley*

2015

Smarter Than Their Machines: Oral Histories of Pioneers in Interactive Computing

John Cullinane, *Northeastern University; Mossavar-Rahmani Center for Business and Government, John F. Kennedy School of Government, Harvard University*

2015

A Framework for Scientific Discovery through Video Games

Seth Cooper, *University of Washington*

2014

Trust Extension as a Mechanism for Secure Code Execution on Commodity Computers

Bryan Jeffrey Parno, *Microsoft Research*

2014

Embracing Interference in Wireless Systems

Shyamnath Gollakota, *University of Washington*

2014

Computing and the National Science Foundation, 1950–2016

Building a Foundation for Modern Computing

Peter A. Freeman

Georgia Institute of Technology

W. Richards Adrion

University of Massachusetts Amherst

William Aspray

University of Colorado Boulder

ACM Books #31



Copyright © 2019 by Association for Computing Machinery

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means—electronic, mechanical, photocopy, recording, or any other except for brief quotations in printed reviews—without the prior permission of the publisher.

Designations used by companies to distinguish their products are often claimed as trademarks or registered trademarks. In all instances in which the Association for Computing Machinery is aware of a claim, the product names appear in initial capital or all capital letters. Readers, however, should contact the appropriate companies for more complete information regarding trademarks and registration.

*Computing and the National Science Foundation, 1950–2016:
Building a Foundation for Modern Computing*

Peter A. Freeman, W. Richards Adrion, William Aspray

books.acm.org

ISBN: 978-1-4503-7274-9 eBook

ISBN: 978-1-4503-7275-6 EPUB

ISBN: 978-1-4503-7276-3 paperback

ISBN: 978-1-4503-7277-0 hardcover

Series ISSN: 2374-6769 print 2374-6777 electronic

DOIs:

10.1145/3336323 Book	10.1145/3336323.3336332 Chapter 8
10.1145/3336323.3336324 Front Matter	10.1145/3336323.3336333 Chapter 9
10.1145/3336323.3336326 Chapter 1	10.1145/3336323.3336334 Chapter 10
10.1145/3336323.3336326 Chapter 2	10.1145/3336323.3336335 Chapter 11
10.1145/3336323.3336327 Chapter 3	10.1145/3336323.3336336 Chapter 12
10.1145/3336323.3336328 Chapter 4	10.1145/3336323.3336337 Chapter 13
10.1145/3336323.3336329 Chapter 5	10.1145/3336323.3336338 Appendixes
10.1145/3336323.3336330 Chapter 6	10.1145/3336323.3336339 Index/Bios
10.1145/3336323.3336331 Chapter 7	

A publication in the ACM Books series, #31

Editor in Chief: Sanjiva Prasad, *Indian Institute of Technology Delhi*

Founding Editor: M. Tamer Özsu, *University of Waterloo*

Area Editor: Thomas Misa, *University of Minnesota*

This book was typeset in Arnhem Pro 10/14 and Flama using ZzT_EX.

First Edition

10 9 8 7 6 5 4 3 2 1

We dedicate this book to the thousands of NSF employees, past and present, and the hundreds of thousands of investigators, graduate students, educators, and reviewers who built and sustained what is often called the “gold standard” of peer-reviewed fundamental scientific research.

Contents

Preface xvii

PART I CHRONOLOGICAL HISTORY 1

Chapter 1 1950–1974: Science Information, Computing Facilities, Education, and Basic Research 3

W. Richards Adrion

- 1.1 Science Information—1950s to 1980s 4
- 1.2 Filling the Demand for Computing Infrastructure 9
- 1.3 Computers in Education 14
- 1.4 Finding a Home for Computer Science Research 17
- 1.5 Summary and Conclusions 20
- Notes 21

Chapter 2 1974–1986: CER, CSNET, NSFNET, and the Founding of CISE 27

W. Richards Adrion

- 2.1 My Background and Perspective on the 1974–86 Period 28
- 2.2 Making the Case for NSF’s Computing Research Programs 29
- 2.3 The Importance of Computing Research and Infrastructure to the Nation 33
- 2.4 Computing and Information Research in NSF, 1974–1978 35
- 2.5 Funding the Innovators in Computer Science 39
- 2.6 Facilities 41
- 2.7 Cryptography and Interactions with the National Security Agency 43
- 2.8 The Computer Science Section, 1979–1984 46

- 2.9 Addressing the Need for Academic Experimental Computer Science 48
- 2.10 CSNET 55
- 2.11 The Office of Advanced Scientific Computing and NSFNET 60
- 2.12 The Beginning of CISE 66
- 2.13 Summary and Conclusions 74
 - Notes 75

Chapter 3 1986–1998: The New Directorate in a Period of Computer Science Expansion 87

Peter A. Freeman

- 3.1 Initial Structure and Leadership of CISE: 1986–1987 88
- 3.2 Organizational Initiation 89
- 3.3 Initial Actions 91
- 3.4 Changes in CISE Leadership 92
- 3.5 Transition to Routine Operation 97
- 3.6 Analysis 108
 - Notes 110

Chapter 4 1999–2006: Broadening Computer Science with New Initiatives 119

Peter A. Freeman

- 4.1 Ruzena Bajcsy (December 1998 to August 2001) 121
- 4.2 Peter Freeman (May 2002 to January 2007) 124
- 4.3 Closing Observations 141
 - Notes 142

Chapter 5 2007–2016: The Growing Centrality of CISE to NSF 147

William Aspray

- 5.1 Deborah Crawford (Acting AD, Early 2007) 147
- 5.2 Jeannette Wing (AD, July 2007 to June 2010) 148
- 5.3 Peter Arzberger (Acting AD, 2010 and 2011) 153
- 5.4 Farnam Jahanian (AD, March 2011 to July 2014) 155
- 5.5 Suzi Iacono (Acting AD, late 2014) 161
- 5.6 James Kurose (AD, January 2015–) 162
- 5.7 Conclusions 166
 - Notes 167

PART II SELECTED SUBJECT STUDIES 173**Chapter 6 Pre-CISE Computing Facilities and Education Programs 175***William Aspray*

- 6.1 Facilities Program 175
- 6.2 The 1970s and the 1980s 180
- 6.3 Support for Computer Education, 1950–1986 181
- 6.4 Conclusions 192
- Notes 193

Chapter 7 Pre-CISE Computing Research 203*William Aspray*

- 7.1 Saburo Muroga and Computer Circuit Research⁹ 208
- 7.2 Walter Karplus, David Kuck, and Computer Architecture Research 213
- 7.3 Mary Shaw, Barbara Liskov, and Software Research 220
- 7.4 Wayne Cowell, David Young, and Numerical Analysis Research 227
- 7.5 Lofti Zadeh, Martin Hellman, and Theoretical Computer Engineering 234
- 7.6 Woodrow Bledsoe, Bertram Raphael, and Artificial Intelligence Research 241
- 7.7 Donald Greenberg, Charles Csuri, and Computer Graphics Research 249
- 7.8 Conclusions 254
- Notes 254

Chapter 8 Information Technology Research 263*W. Richards Adrion*

Notes 272

Chapter 9 Networking Research and Deployment 275*W. Richards Adrion, Peter A. Freeman*

- 9.1 Early NSF-Supported Research on Networks (1950–1980) 277
- 9.2 NSF Leads Public Networking (1980–1995) 278
- 9.3 Transitioning to the Commercial Internet 282
- 9.4 Network Research after NSFNET (1996–2001) 286
- 9.5 Networking Research (2002–2004) 288
- 9.6 Initiation of the GENI Project (2004–2006) 289
- 9.7 The Origin of GENI (2004) 290

- 9.8 Conceptual Design for GENI (2004–2006) 291
- 9.9 Implementation of GENI (2007–2016) 293
- 9.10 Conclusion 294
- Notes 295

Chapter 10 NSF Support of High-Performance Computation 303

Peter A. Freeman

- 10.1 1950–1954 304
- 10.2 1955–1983 304
- 10.3 1984–1991 306
- 10.4 1992–2000 308
- 10.5 2000–2004 308
- 10.6 2005–Present 309
- Notes 310

Chapter 11 CISE’s Role in Broadening Participation in Computing 313

William Aspray

- Conclusions 320
- Notes 321

Chapter 12 What Does an AD/CISE Do? 323

Peter A. Freeman

- Notes 333

PART III SUMMARY AND CONCLUSIONS 335

Chapter 13 Summary and Conclusions 337

Peter A. Freeman, W. Richards Adrion

- 13.1 Summary 337
- 13.2 Concluding Remarks 346

APPENDIXES 349

Appendix A Computing Organizations at NSF 351

W. Richards Adrion

- Notes 359

Appendix B CISE Oral Histories List 361

Appendix C Short Biographies 363

Appendix D CISE History Archive (CHA) 371

Appendix E Abbreviations and Acronyms 373

Bibliography of Secondary Literature on NSF and the History
of U.S. Computing 381

Author Biographies 389

Index 391

Preface

The Computer and Information Science and Engineering (CISE) Directorate and its predecessors at the National Science Foundation (NSF) have played a seminal but untold role in the growth of computing¹ from the 1950s to today. Since the mid-1990s, CISE has provided a large majority of *all* funding for basic research in computer science and closely related disciplines in the United States, as well as substantial support for other fields that study computing or push the state-of-the-art of advanced computation. The results have formed the foundations on which modern computing is built.

Two of the authors of this book, Peter Freeman and Rick Adrion, were aware of much of this history and knew also that, to date, no comprehensive record of the influential role played by CISE and its predecessors existed. As a result, in late 2016, we undertook to remedy this situation by producing a documented history of NSF's role in modern computing. Recognizing that we had no formal training as historians, we enlisted William Aspray, an historian who had published extensively on computing-related subjects including at NSF,² and with whom we had worked in other contexts, to join us on the project; his experience has been essential. This book, and a related publicly available collection of research materials³ deposited at the Charles Babbage Institute (CBI) of the University of Minnesota, are the principal results of our efforts.

Our project had four objectives. The first was to bring together as much *information* as possible that pertains to the history of computing⁴ at NSF. We have collected approximately 4,000 paper and electronic records, which were donated to the CBI.⁵ We spent considerable time talking with longtime members of the CISE staff to locate materials and develop context for later project activities. We also collected materials and consulted various archival collections.⁶

We have augmented this written material with approximately 50 oral histories,⁷ which have been transcribed and lightly edited. (Most of these will be available

through the CBI as well.) They include interviews with several NSF directors and eight of the nine living Assistant Directors (ADs) of CISE.⁸ Additional oral histories were conducted with staff within CISE (program officers, division directors, or chief scientists) as well as with other members of the Washington computer science community; for example, former members of the White House Office of Science and Technology Policy (OSTP). A list appears in the back of this book.

A second objective of this project was to rigorously *document* major events in the history of NSF support for computing research and education. Throughout the text, we have provided citations to numerous primary sources, including NSF internal memoranda and internal plans that are no longer sensitive, published documents, and other government publications. In those cases where materials we cite in this book would be difficult for readers to obtain, we have placed them with the CBI.

A third objective was to write a set of *narratives* describing the history in a readable and accessible way. This has been greatly facilitated by the fact that both Adrion and Freeman served as employees or rotators⁹ at NSF on several occasions for a combined total of 18 years, and were not only active researchers and educators (professors) but also engaged members of the professional community for almost 50 years each. Additionally, Aspray had led a team in the early 1990s that produced a large body of unpublished research on pre-CISE activities based on internal NSF documents. When he became part of this project, we gained particular advantage for satisfying this objective as he accessed those writings and drew from his 40-plus years of experience as an historian of computing.

Our fourth objective was to *analyze* what we have learned. Conclusions are indeed drawn in Chapter 13 as well throughout the rest of the text. However, generally speaking, we have not evaluated CISE programs or the individual projects that CISE supported; where we have offered judgmental opinions, these are solely the opinions of the chapter's author(s). Further analysis must await future authors. We mention other major government funders including DARPA, NASA, DoE, and the military research agencies; science policy in both the legislative and executive branches of the federal government and in the National Academies; and computing professional organizations including ACM, the IEEE Computer Society, SIAM, AFIPS, and the Computing Research Association. While we occasionally discuss the relations of these organizations to NSF, we have not identified and analyzed the many connections among the various players in this milieu and NSF. Nor have we tried to evaluate their relative contributions and merits.

Our primary focus has been on CISE (created in 1986) and its predecessor organizations, such as the Office of Computing Activities (OCA, created in 1967) and the Office of Science Information Services (OSIS). However, computing activities within NSF often extended beyond the boundaries of OCA and OSIS.¹⁰ We mention

these, but typically do not follow them in the same detail that we give to CISE and its predecessors.

Readers will find that there is some variation in the nature of the three main parts of this book, and even variation in style among its individual chapters. We wrote some chapters as participant accounts, but wrote others more objectively as historians who did not directly participate in the events described. In parts of Chapters 3, 4, and 9, for example, Peter Freeman writes from the perspective of a direct participant; in Chapter 12, he reflects on his time as the AD/CISE. In Chapters 1 and 2, Rick Adrion draws upon his early role as a program director and on his later key management experience at NSF to tell the story of critical events before and after CISE was created. In Chapter 13, Freeman and Adrion reflect on the history of NSF and computing to identify some themes that may help in future understanding. William Aspray, who has never been employed by NSF, has worked as both an historian and as the executive director of the Computing Research Association (CRA—one of the major non-profit players in Washington on computing research policy). The chapters he wrote on CISE's role in the development of modern computing are informed by this perspective.

While there has been coordination among the authors to ensure thorough coverage of the history of computing at NSF in the period from 1950 to 2016, this book is best read as a collection of linked essays rather than as a tightly written monograph. The three authors each have their individual voices, and no effort has been made to harmonize them completely. While we have all read and critiqued each other's chapters, we did not have a goal of forging a unified position throughout.

Our book is composed of distinct parts that present the results of our work on our project over the past two years. Part I provides a narrative of the history of NSF's involvement in the world of digital computing, especially as it relates to the funding activities of CISE and its predecessors. (Table P.1 gives a timeline of some of the key events in this narrative, to assist in comprehending some of the milestones passed.) Part II goes into more depth on a selected set of important topics. Part III provides our conclusions, and the appendixes present NSF organizational charts over time, a list of the interviews we conducted, a non-exhaustive set of short biographies, a description of the archive we prepared, and a list of abbreviations and acronyms,

Before providing a guide to using this book, we provide very short characterizations of each of the 13 chapters.

Chapter 1 covers computing activities related to science information, facilities, education, and basic research in the period from 1950 to 1974. The most active early support entailed providing science information and support for research in information retrieval, databases, and computational linguistics. Computing facilities and education were supported more heavily than computing research, but did

Table P.1 Select events in the history of NSF and computing, 1950–2008

Year	Event
1950	NSF enabling act signed in November; operations begin in 1951
1951	Office of Science Information (OSI) created
1953	Assistance given to buy computer for research
1954	First training/education grant.
1955	von Neumann panel recommends research on design of computers; National Science Board approves facilities program
1957	First grants for computing research
1958	Office of Science Information Services (OSIS) created; NSF expands computing facilities, research and education investments
1963	Early Training grants led to the first CS curricula and departments
1966	Rosser Report
1967	Pierce Report; Office of Computing Activities (OCA) created
1974	Division of Computer Research (DCR) created, then recreated in 1984
1978	Theorynet and Debate on Public Cryptography
1980	CER (experimental research) and CSNET (networking) programs begin
1984	Supercomputer Centers created and NSFNET begins
1986	CISE created
1995	NSFNET converted to Internet
1999	ITR program started
2003	Major reorganization of CISE
2004	GENI Program started
2005	Broadening Participation Program started
2005	Office of Cyberinfrastructure (OCI) created in O/D
2006	CCC created
2008	Expeditions in Computing, Cyber-Enabled Discovery programs begin
2013	OCI moved from O/D to CISE as Division of Advanced Cyberinfrastructure (ACI)
2016	ACI made into an office (OAC) within CISE to provide better connection with rest of NSF

enable the creation of some of the earliest computer science academic departments. Creation of the Office of Computing Activities (OCA) in 1967 was a landmark development because it strengthened support for computing research and provided organizational status; that resulted in stronger ties to other NSF programs and the NSF imprimatur to fledgling academic computer science departments.

Chapter 2 covers the years from 1974 until the founding of CISE in 1986. In addition to organizational changes and further strengthening of computing programs, there was support for efforts to professionalize and define computer science. A series of reports (e.g., Feldman, Snowbird, Hopcroft-Kennedy, Lax, Bardon-Curtis) shaped NSF's computing efforts. Among the results were cryptologic research, the Coordinated Experimental Research (CER) program, and the CSNET and NSFET networking initiatives; these are afforded expanded discussion. This growing importance of computer science and of computing, coupled with internal efforts by several NSF staff, led to the founding of CISE.

Chapter 3 covers the years from the founding of CISE in 1986 through 1998. There was a succession of short-term ADs: Gordon Bell, William Wulf, Nico Habermann, Paul Young, and Juris Hartmanis (all served approximately two years each). In spite of some internal pushback, the new Directorate quickly established its structure and importance within both NSF and the federal government. During the 1990s the first easily usable browser (Mosaic), conversion of NSFNET into the Internet, and the emergence of Google were all enabled in some way by CISE support.¹¹ By 1999 CISE started to receive greater funding from Congress, increased respect within NSF, and sustained leadership from its scientific community.

Chapter 4 covers 1999–2006 when Ruzena Bajcsy and Peter Freeman served as CISE ADs. Major initiatives increased support for cyberinfrastructure, greatly expanded the field with the Information Technology Research (ITR) program, reorganized CISE, started new funding programs in networking research (GENI—the Global Environment for Network Innovations), cybersecurity research (a centers program), and the Broadening Participation in Computing (BPC) program. Direct actions by CISE made significant management changes in the supercomputer centers and strengthened the cyberinfrastructure and basic research programs. Initial plans were laid for later initiatives including the Expeditions in Computing program and Cyber-enabled Discovery and Innovation (CDI).

Chapter 5 covers 2007–2016. Three individuals served as AD/CISE: Jeannette Wing, Farnam Jahanian, and James Kurose. Budgets were tweaked to ensure that basic computer science research was protected and that CISE received fewer but better proposals. A major one-time appropriation was received and successfully managed as part of President Obama's stimulus package. Several major programs,

such as GENI, Expeditions in Computing, and Cyber-enabled Discovery and Innovation were furthered during this time. The Computing Community Consortium was continued and there was increased partnering with other directorates and industry.

Chapter 6 provides a detailed analysis of the NSF programs in computer facilities and computer education prior to the founding of CISE in 1986.

Chapter 7 provides case studies of early NSF support for research in circuits, computer architecture, software, numerical analysis, computer engineering theory, artificial intelligence, and computer graphics.

Chapter 8 covers the Information Technology Research Program from its beginning in FY (fiscal year) 2000 through to 2005, when it became part of base CISE research funding.

Chapter 9 provides a case study of NSF's support of research on concepts and mechanisms of networking, and deployment of operational networks.

Chapter 10 covers High Performance Computing, an activity NSF has supported even as the power of such machines has grown exponentially.

Chapter 11 covers CISE's programs to broaden participation in computing to women, underrepresented minorities, and the disabled.

Chapter 12 provides a personal view of what a CISE AD does.

Chapter 13 recaps the narratives in Chapters 1 to 5 and provides a set of high-level conclusions about the history of computing and NSF funding.

Readers seeking an overview of NSF activities in computing research and education, as well as related activities, are encouraged first to read Chapters 1–5, and then follow up by reading any deeper studies that are of particular interest. The organizational charts in the appendixes may also be useful in understanding one aspect of the changing relationship between NSF and computing.

Readers with limited time and/or scope of interest may want to read only the chapter(s) in Part II that speak to their interests. A quick scan of the chapters' beginnings may help to determine whether one of them addresses the reader's interest.

Table P.2 may be of use in connecting Part I chapters with Part II chapters. It illustrates the major connections between a given chronological chapter and one or more subject study chapters.

Work on this project was supported in part by NSF Grant #1743282, EAGER: Exploring the History and Impact of the Computing and Information Science and Engineering (CISE) Directorate of the National Science Foundation, a grant made to the Massachusetts Green High Performance Computing Center (MGHPCC). We worked independently of the NSF. Any views expressed in this book are solely due

Table P.2 Relation between Part I chapters and Part II chapters

Narrative Chapter	Related Study Chapter(s)	Example Usage of Study Chapter(s)
1	6, 7, 9, 10	What are examples of early support of education for computing? What parts of NSF provided support to provide computer access?
2	7, 9, 10	When did NSF first create a specific organization for computing research? When did the supercomputer centers start?
3	8, 9, 10, 12	Why was the Information Technology Research (ITR) program created? What role did NSF have in creating the Internet?
4	8–12	What happened to NSF involvement in networking after the Internet? Is the AD/CISE involved with anything besides computing?
5	10–12	What is Cyber-Enabled Discovery & Innovation (CDI)? What is NSF doing about cybersecurity?

to us or named third-party sources, not the NSF nor the MGHPCC. Any errors of fact are our responsibility.

Work on this project would not have been possible without NSF support and the help of many people. Erwin Gianchandani, currently Deputy AD/CISE, guided us on the usage of NSF materials and other issues. NSF Historian Leo Slater answered questions and Assistant NSF Historian Emily Gibson provided access to some NSF records. Janet Abbate (Virginia Tech), Thomas Haigh (University of Wisconsin-Milwaukee), and Jeffrey Yost (Charles Babbage Institute) served as our historical advisory committee. We have worked closely with Amanda Wick, the Acting CBI Archivist, on the deposit of project materials at CBI. Several former and current CISE staff have donated material to the project. Over 50 individuals have agreed to sit for oral history interviews. A succession of four people provided diligent support to the project: Jana Vetter, Julia Fan, Jessica Ewen, and Kayla Heslin. In particular, we want to thank Julia and Jessica for their work on the oral histories and Kayla for the work at the end of the project as we compiled this book and readied materials to be sent to the CBI. Finally, we wish to thank our families, who have sometimes missed us and been neglected as we worked on this project. Many thanks to all!

Notes

1. Before the early 1960s there was no computer science, but by the 1970s the term was widely known and departments of computer science rapidly became a dominant academic unit and scientific discipline. We will use “computer science” primarily to refer to the research discipline and “computing” to refer to the broader activity of using computers and studying that usage.
2. William Aspray, Bernard O. Williams, and Andrew Goldstein, “Computing as Servant and Science: Impact of the National Science Foundation” (unpublished, 510 pages, 1992).
3. The CISE History Archive (CHA) is described in Part III of this book.
4. We use the ambiguous term “computing” to denote computer science and closely related disciplines, but not all uses of computing by other fields; however, especially in the early days, the distinction was not yet clear.
5. One motivation for this collection effort was the physical move of NSF headquarters in September 2017 from Arlington to Alexandria, Virginia, and the knowledge that valuable documents might be discarded. Another motivation was that early NSF employees are starting to pass on—and their memories and their documentation with them. Two interviewees passed away during the project and several potential interviewees were incapacitated.
6. These included documents from Gordon Bell, Mel Ciment, Mike Foster, John King, Irene Lombardo, Jack Minker, Rick Adrion, and Peter Freeman. Archival collections consulted included those of Ed Feigenbaum and John McCarthy at the Stanford archives.
7. The oral history record is strong but not complete. Many of the principal people involved with CISE and its predecessors have been interviewed, but a few are deceased, a few we could not reach or they did not agree to be interviewed, and due to oversight or lack of time, no doubt a few were missed. While there were perhaps 10 or 20 oral histories concerning the NSF computing story existing at the time we began this project (mostly at the Charles Babbage Institute Archives, the IEEE History Center, and the Computer History Museum), the new interviews we have added represent a major increase in coverage of this topic.
8. The AD is the head of the directorate; “Assistant” indicates they also have NSF-wide responsibilities, reporting to the NSF Director.
9. A “rotator” at NSF is a person on leave from their home institution to work at NSF under the Intergovernmental Personnel Act (IPA) or as temporary employees.
10. This is true of early work on science information and information science. High-performance computing and cyberinfrastructure have sometimes been housed within CISE, but at other times either in the Office of the Director or in their own freestanding office reporting to the Director. At times, computing activities have existed in other NSF directorates: especially Engineering, Mathematics and Physical Sciences, Biology, and Education and Human Resources.
11. See the list of acronyms and abbreviations that appears in Appendix E of this book.



PART

**CHRONOLOGICAL
HISTORY**



1950–1974: Science Information, Computing Facilities, Education, and Basic Research

W. Richards Adrion

As the National Research Council report *Funding a Revolution* states, “rather than a single, overarching framework of support, federal funding for research in computing has been managed by a set of agencies and offices that carry the legacies of the historical periods in which they were created.”¹ This chapter traces the parallel development of NSF programs in science information, computing facilities, computer-supported education, computational science, numerical computation, and the beginning of computing and information research programs. The NSF role in the federal support of computer science, computer engineering, and information science advanced within separate units and programs until they began to consolidate in the 1980s.

Prior to the Second World War, academic research funding for most disciplines came from universities’ internal resources, industry, foundations, and philanthropic sources.² The war years saw a large investment by the federal government. In 1941, President Roosevelt established an Office of Scientific Research and Development (OSRD),³ an arm of the Office of Emergency Management, with Vannevar Bush as director. OSRD remained in existence through 1945. During the 15 years following the Second World War, research in computing and communications was supported by mission agencies connected to the military, atomic power, and space.⁴

During and following the war, a number of efforts were underway to establish a “science foundation,”⁵ mainly led by Senator Harley Kilgore (D-WV), who chaired the Senate Subcommittee on War Mobilization of the Military Affairs Committee (the “Kilgore Committee”). As the debate over the appropriate agency or structure for supporting scientific research continued, President Roosevelt asked OSRD Director Vannevar Bush to have a say. Bush delivered his report in 1945, entitled “Science—The Endless Frontier,”⁶ to Roosevelt’s successor, President Harry Truman. Truman vetoed the National Science Foundation Act of 1947⁷ primarily because it did not give the president authority to name a single, politically appointed director of the agency.⁸ After three more years of debate, Congress passed and President Truman signed Public Law 81-507,⁹ creating the National Science Foundation; operations began in 1950.

While the Foundation had been interested in science information as early as 1951, following the Sputnik launch on October 4, 1957, the NSF role in science information increased and it was given a new emphasis on addressing the need for computing in both research and education. The NSF did not become a significant player in computing research, however, until the 1970s. Several threads of NSF support for science information, computing infrastructure, computers in education, and early computer science and information science research funding led to the NSF divisions, offices, and programs that later comprised the Computing and Information Science and Engineering (CISE) Directorate starting in 1986.

1.1 Science Information—1950s to 1980s

NSF’s Office of Scientific Information (OSI) was established in 1951 with Robert Tumbleson as head. OSI initially had four programs: Publication Support and Scientific Documentation, Foreign Science Information, U.S. Government Research Information, and Exhibits.¹⁰ Between 1952 and 1955, OSI supported the publication of scientific books and journals, Soviet-focused projects (translation, including machine translation, and symposia), studies of information processes and methods, abstracts and indexes of government, professional society and international science publications, and linguistics research related to machine translation.¹¹ The NSF Advisory Panel on Scientific Information—made up of scientists, publishers, a university president, and the assistant librarian of the Library of Congress—held its first meeting in 1953. As OSI expanded, Alberto Thompson succeeded Tumbleson as its head in 1955. Among the OSI program directors was Helen Brownson, “an outspoken advocate and significant figure in many pivotal events which formed what is now known as information science,”¹² who was responsible for guiding

many of the research efforts funded by OSI and its successor, the Office of Science Information Services.

While OSI's primary mission was managing and coordinating science information across federal agencies, NSF also began to support applied and basic research activities. In May 1956, NSF sponsored a meeting¹³ of representatives of the Department of Defense, National Bureau of Standards, and the Patent Office, as well as experts in linguistics, logic, information theory, operations research, computer design, and library science, to discuss fundamental research on the organization of information. On April 15–17 of the following year, Western Reserve University (WRU) hosted a Symposium on “Systems for Information Retrieval.”¹⁴

Following a period when Thomas Jones was acting head, Burton Adkinson became head of OSI.¹⁵ As Adkinson noted, “In 1957, two unrelated events made a big impact on NSF/OSI. The first was the untimely death of Alberto Thompson, who had barely started to develop a vigorous scientific information support program. Second, the launching of Sputnik surprised most Americans.”¹⁶

In 1958, the President's Science Advisory Committee (PSAC) created the “Baker Panel.”¹⁷ Packed with luminaries and influential figures,¹⁸ this panel issued a report on “Scientific Judgments on Foreign Communications Intelligence” that called for improving the availability of U.S. scientific and technical information.¹⁹ PSAC endorsed the recommendations of the Baker Report; and the President's Science Assistant, James Killian, Jr., urged presidential approval. A White House press release in December 1958 directed “the National Science Foundation [to] take leadership in bringing about effective coordination of various scientific information activities within the Federal Government.”²⁰

The post-Sputnik National Defense Education Act (NDEA) became law on September 2, 1958. It contained major provisions²¹ for loans to higher education students; fellowships for advanced study of mathematics and science; guidance counseling and testing to identify able students; improvement of K–12 science, mathematics, and foreign language programs; vocational programs; and research on effective uses of television and other media for educational purposes. In addition, the NDEA authorized the National Science Foundation to establish a Science Information Service: first to address indexing, abstracting, translating, and to provide other services leading to a more effective dissemination of scientific information; and next to undertake programs to develop new or improved methods for making scientific information available.²²

On December 11, 1958, NSF established the Office of Science Information Service (OSIS) with Adkinson as head. By the end of the decade, OSIS had made 146 grants totaling about \$3.8 million under four major programs: Documentation

Research (through which most of the research and development was funded), Foreign Science Information, Publications and Information Services, and Unpublished Research Information. Among these grants²³ were projects on linguistic transformation for information retrieval at the University of Pennsylvania and mechanical translation projects at Harvard Computation Laboratory, Georgetown University, the University of California, Massachusetts Institute of Technology, and the Cambridge Language Research Unit in England. OSIS also funded the National Bureau of Standards to establish a Research Information Center and Advisory Service on Information Processing in 1959.²⁴

In the late 1950s, it was unclear how to classify the various fields that encompass the basic sciences behind computing, computers, information, communications, and the fields that depend on them. Louis Fein, a Stanford Research Institute (SRI) consultant, was asked by Frederick Terman and Albert Bowker of Stanford University to design a computing curriculum. Fein began studying university programs “in the fields of computers, data processing, operations research, and other relatively new and apparently closely related fields.”²⁵ His goals were to identify not only computing-related organizations, curricula, research programs, and facilities, but also computing-related fields of study, and the role of the universities in these fields. As Fein noted in 1959,²⁶ “universities, as institutions, are having a hard time . . . learning how to effectively incorporate these new fields into the academic structure.” In recommending the creation of a Graduate School of Computer Sciences at Stanford, Fein defined two research-oriented departments.²⁷ “Information and Communication” encompassed instruction and research activities in information theory, switching theory, coding theory, automata theory, artificial intelligence, learning, language translation, and theory of simulation. “Systems” comprised instruction and research activities in management science, econometrics, systems theory, information classification, indexing and retrieval, model theory, self-organizing systems, and adaptive mechanisms. Today, the former might fall under a computer science (or engineering) department, while the latter might be divided among departments of information systems, information technology, and management information science. Fein saw a divide between the science of computing, communications, and information and the application and use of computing, communications, and information.

As we describe in Chapter 2, efforts to formally establish computer science as a discipline accelerated in the late 1960s and early 1970s. By the early 1960s, the fields and practitioners of information technology and information science were becoming better defined. Information technology—the more applied side—was staffed by information specialists, while information science—the research

side—was staffed by information scientists. As we relate later, Altman and Brown²⁸ described the creation in the 1980s of the CISE Directorate as a move away from the library scientists and specialists supported under OSIS, to support for computer and information scientists.

Dorothy Crosland organized a series of conferences²⁹ at the Georgia Institute of Technology, for the first time making a distinction between information specialist and scientist. A *specialist* was someone who applied technology to the storage, indexing, and archiving of information, while a *scientist* was concerned with the nature of information and its representation. These conferences had a significant impact on the establishment of new information research programs at Georgia Tech, Lehigh University, and Drexel University.³⁰

The OSIS programs continued to expand. In 1967, OSIS made grants to Georgia Tech (Vladimir Slamecka) and Ohio State University (Marshall Yovits) to expand programs in information science. It also made grants to professional scientific societies to improve their literature services. The Georgia Tech center had two principal activities: mathematical models for information in the scientific disciplines and control of information for problem solving and decision making in an academic environment.³¹ By 1968, NSF awards to various professional societies to develop computerized information retrieval systems had grown to \$17.7 million,³² up from \$9 million in 1958. While the percentage of OSIS funds going to research projects was approximately 5.5% in 1958, eventually 50% of OSIS funding was spent on disciplinary information research centers.

In 1969, OSIS was moved organizationally from reporting directly to the NSF Director to reporting to the Assistant Director for National and International Programs, where OSIS staff were less able to make a case for funding directly to the Office of Management and Budget (OMB). With declining interest in supporting OSIS within its new directorate, science information activities declined as its appropriations waned. OSIS also had to assume responsibility for the Committee on Scientific and Technical Information (COSATI), which was transferred from the President's Office of Science and Technology. This greatly increased the burden on OSIS staff³³ and its resources. These changes also resulted in a termination of operating grants for information services and unrestricted grants to university research centers for information science by 1972.³⁴ OMB further reduced the OSIS appropriations to \$5 million in 1974 and asked NSF to phase out support to the university-centered information systems programs at Pittsburgh and Ohio State and to the New England Board of Higher Education science information network. These and similar organizations at the University of Georgia, UCLA, and Lehigh University continued at their own expense.³⁵

During the period from 1971 to 1973, OSIS also experienced a rapid change in staffing.³⁶ Adkinson retired and moved to the American Geographical Society in 1971. Melvin Day, who replaced him as head of OSIS, left NSF in 1973 to accept a position as Director of the National Library of Medicine. Lee C. Burchinal was named as Day's replacement.³⁷ NSF meanwhile established priorities among the five OSIS programs: Research Support, National Information, User Support, Economics of Information, and Foreign Science (with the major emphasis remaining on the Research Support program).³⁸

NSF undertook a major reorganization in 1975, creating four new directorates: Mathematical, Physical, and Engineering Sciences; Astronomical, Earth, and Ocean Sciences; Biological and Social Sciences; and Scientific, Technological, and International Affairs, which joined Science Education, Research Applications, and Administration.³⁹ OSIS was renamed the Division of Science Information (DSI) in 1976⁴⁰ within the reorganized Directorate for Scientific, Technological, and International Affairs. At this time, the Office of Computing Activities, which briefly had joined OSIS in the Directorate for National and International programs, became the Division of Computer Research (DCR) in Mathematical, Physical, and Engineering Sciences.

DSI became the Division of Information Science and Technology (DIST) in 1978 and responsibility for supporting the dissemination of scientific information was distributed among the research divisions within NSF, making it appear that NSF was shifting away from efforts to support the *users* of scientific information and would concentrate instead on funding the development of new information science technology and its applications.⁴¹ Altman and Brown⁴² called the 1978 reorganization “a major cleavage between past and future,” noting a shift from focus on publication, distribution, and dissemination of documents, and improving access to and indexing of documents, to a prioritization of “information science research.”

Following the creation of DIST, former DSI head, Lee Burchinal, transferred to another NSF office and Harvey Averich served as acting head of DIST with a staff of 12 and a budget of approximately \$4.5 million. Program directors Edward Weiss, Harold Bamford, and Richard Lee all moved from DSI to DIST.⁴³ Altman and Brown noted that the DIST managers “shied away from defining ‘information,’ and consequently its science” largely because the term meant “different things in different disciplines.”

Howard Resnikoff, a mathematician who had been brought in as the founding DIST director in 1980, noted that the new program in information science “incorporates certain research responsibilities of previous Foundation programs which were primarily concerned with science information dissemination [but the]

focus of effort [is] so different, that prior award and funding patterns are not comparable. . . .”⁴⁴ Resnikoff attempted in his few years (1979–1981) at NSF to create a significant role for DIST, assembling a distinguished advisory group that included Gordon Bell, Seymour Cray, Ed David, John Gibbons, Ralph Gomory, George Heilmeier, Donald Knuth, and Joshua Lederberg. His goals were for DIST to support research on the structure of information, infometrics, behavioral aspects of information transfer, measures of fundamental quantities, and standards for assessing the predictions of theory and comparing the results of experiments. Resnikoff left NSF in 1981 to join Harvard University and later co-founded Thinking Machines Corporation. He also founded FutureWave, an intellectual property company.

Resnikoff left DIST when it moved to the Directorate for Biological, Behavioral, and Social Sciences (BBS). Edward Weiss became acting division director of DIST and its three programs: Information Science, Information Technology, and Information Impact. Information Science was concerned with the properties of information and the dynamics of information transfer, including biological and human information processes. Information Technology dealt with improving theory underlying the design of systems and problems with user-system interaction emphasizing human factors. Information Impact was interested in the economic and social consequences of information and information technologies.⁴⁵ Weiss argued that BBS as a research directorate was likely to provide a more favorable climate for the division.⁴⁶

Following the creation of the Computer and Information Science and Engineering (CISE) Directorate in 1986, Harold Bamford and Charles Brownstein discussed the emergence of information science research as a more fundamental question being revealed by the unfolding structure of knowledge. They argued that the “evolution of units supporting information science research”⁴⁷ in CISE was a “recognition of the unity and coherence of the intellectual streams, which converge in computer and information science and engineering and in the great importance which [NSF] attaches to the confluence.”⁴⁸ Several unmet needs focused NSF’s attention as CISE evolved.

1.2 Filling the Demand for Computing Infrastructure

In the years following the Second World War, a commercial computer industry came into being, including leading efforts at IBM and Remington (later Sperry) Rand and other companies such as Bendix, Burroughs, General Electric, Honeywell, Raytheon, and RCA. Federally funded projects constituted roughly three-quarters of the total computing infrastructure. Government facilities, government-funded

research centers, and private federal contractors were typically pushing the technical cutting edge.⁴⁹

During these years, computing research was supported primarily by mission agencies of the federal government, especially defense and energy agencies (initially the Atomic Energy Commission), and later NASA. The Foundation was beginning, however, to recognize that the computer was an important tool for scientific research. The 1955 Annual Report noted that:

. . . a revolution has occurred in scientific work in that much of it now calls for exceedingly expensive structures and equipment . . . which already have outrun the financial capacity of private resources, and this will increasingly be the case. Only the Federal Government . . . will be able to meet the deficiency after all possible private resources have been utilized.⁵⁰

Scientists and engineers outside the military and atomic laboratories were having difficulty accessing computers due to heavy security constraints. The high cost of maintaining a modern computation laboratory and the challenge and pitfalls of charging usage fees, “a practice which affects the character of its scientific program,”⁵¹ limited access to academic computing centers.

The NSF entered into an agreement with the Applied Mathematics Laboratories of the National Bureau of Standards (NBS) for “advice on the methods of numerical analysis and the choice of machines for specific computation involved in requests . . . ”⁵² That year (1955), NSF made computational grants (with advice from NBS) to the Ohio State University; the University of Texas; the University of California, Berkeley; and the University of Illinois.⁵³

In February 1955 the NSF appointed an ad hoc Advisory Panel on University Computing Facilities, led by John von Neumann.⁵⁴ The panel recommended “that the Foundation establish a limited program to provide computing equipment and partial support for appropriate staff in order to carry on research and training in high-speed computation.” The report also noted that research in the advanced design of computing machines should be recognized as being of basic importance: “it is desirable that the speed of computing machines be increased by a factor of at least 50 and that their capacity be substantially increased.”⁵⁵ At its October 1955 meeting the panel recommended that “\$5 million be expended for the development of a fast, large computing machine of advanced design.”⁵⁶

Leading this panel was not the only instance where von Neumann played a role in developing NSF’s computing facilities program. He earlier had proposed the stored program concept in his “First Draft of a Report on the EDVAC,”⁵⁷ and he built such a machine at the Institute for Advanced Studies (IAS) in Princeton. Com-

puter simulations were frequently used for both meteorology and nuclear weapons and von Neumann had realized that these two fields were closely connected scientifically. Both were centrally concerned with highly nonlinear fluid dynamics.⁵⁸ Von Neumann was the principal investigator on an NSF grant to organize the Conference on High-Speed Computing in Meteorology and Oceanography⁵⁹ held May 13–15, 1954, at the University of California, Los Angeles. Following this meeting, NSF funded the aforementioned advisory panel convened by von Neumann, then at the Atomic Energy Commission (AEC). In May 1956, von Neumann outlined the needs for facilities, which were critical to the advancement of science yet beyond the financial means of universities and the National Science Board; it subsequently approved a computer facilities program.⁶⁰ Von Neumann died early the following year.

The career of John Pasta connected von Neumann, his IAS machine, the AEC, the Los Alamos National Laboratory (LANL), and NSF. Pasta had a long and unusual career, beginning as a New York City police officer, then an Army Signal Corps officer, a physics PhD student, and eventually a staff member at Los Alamos. In 1953, Pasta, Stanislaw Ulam, and Enrico Fermi used the LANL MANIAC computer, based on von Neumann's design for the IAS computer, to identify the Fermi-Pasta-Ulam (FPU) problem,⁶¹ a fundamental advance in soliton theory. In 1956, von Neumann invited Pasta to head what became the AEC Division of Mathematics and Computer Research. In 1961, Pasta left the AEC to join the University of Illinois as chair of the computer science department and later became director of the NSF Office of Computing Activities, director of the NSF Division of Computer Research (DCR), and director of the NSF Division of Mathematical and Computing Sciences (DMCS).

NSF continued to make grants for university computing centers and research in numerical analysis through the 1950s, for example at Cal Tech, MIT, Oregon State, Washington, and Wisconsin in 1956. Research grants went to Cal Tech, Berkeley, Cornell, MIT, Oregon State, Penn, Princeton, Purdue, Stanford, Washington, and Wisconsin the following year.

In July 1960, an institutional grants program was created to assist institutions to strengthen their general research and training functions. NSF made 6 grants in 1961 totaling \$1,685,000 for the acquisition or rental of high-speed computers and 20 grants totaling \$796,000 for computing centers and procurement of small computers. Because NSF funding was limited, the Foundation limited computer center support to an amount equal to 5% of a proposing institution's research grant income, capped at \$50,000 (later reduced to \$37,500). Using this formula, NSF made institutional grants for computing infrastructure totaling \$1,496,604 to

248 institutions; more than half the awards amounted to \$2,000 or less, while just 10 institutions received the maximum grant of \$37,500.⁶²

In June 1962, NSF Director Alan Waterman requested that the National Academy of Sciences' National Research Council undertake a study of "the status and likely growth of computer uses. . . ." J. Barkley Rosser prepared the National Academy of Sciences report, "Digital Computer Needs in Universities and Colleges." The Rosser Report⁶³ was completed in 1966 and made a strong case for universities having access to high-performance computers, but it said little about education. In 1963, the Foundation was able to provide only limited support for computing facilities due to the magnitude of the need. Institutions were required to provide as much as two-thirds of the purchase price from a non-federal source. Even though funding increased to \$4,980,000 in fiscal year 1963,⁶⁴ only 13 grants were made.

Arthur Grad administered the computer facilities grants at NSF beginning in 1959 and he recalled that the Rosser Report:

. . . all started with Phil Morse at MIT. They needed a bigger computer. They estimated they would need about ten million dollars. And I told them, well, there wasn't much I could do about it since my entire budget was only five (million). And I suggested to him that probably the best thing he could do was to have a National Academy study done pointing out the need for more money for computers. So, the Academy duly appointed the committee to make those studies. . . . But it all started from Phil Morse's need for a big computer.⁶⁵

At the time Morse was seeking additional funding, MIT had received a 7094 computer from IBM on which MIT faculty began development of the CTSS operating system.⁶⁶ The CTSS operating system, a forerunner of Project MAC, Multics, and eventually Unix, was based on an idea of John McCarthy, then at MIT. In an influential memo titled "A Time-Sharing Operator Program for Our Projected IBM 709," he proposed interactive time-shared debugging. Herb Teager and McCarthy gave a presentation entitled "Time-Shared Program Testing"⁶⁷ at the national ACM meeting in September 1959.⁶⁸ Much of the CTSS research was funded by NSF grants to the MIT Computation Center. This is clearly an example where fundamental advances occurred through NSF funding of infrastructure. McCarthy started working at BBN with JCR Licklider and others at around that time, and it is said that McCarthy influenced Licklider's thinking about time-sharing. Licklider later went to ARPA, where he funded Project MAC at MIT, based on CTSS, and many other important initiatives.

NSF established the Office of Computing Activities (OCA) in July 1967 to provide federal leadership in the use of computers for research and education. Later,

the directive was added as a statutory requirement to the NSF charter. Faced with ever-increasing demand for computing facilities from all sectors of academe, OCA established regional centers. In fiscal years 1968 and 1969, the Foundation explored various computer-based cooperative arrangements. Typically, each regional activity was centered on a major university, which provided computer services and technical assistance to help a cluster of nearby institutions introduce computing. Altogether, 15 regional centers were established, including 12 major universities, 116 participating colleges, 11 junior colleges, and 27 secondary schools located in 21 states. By the early 1970s, 30 regional computing networks were connecting approximately 300 institutions at all levels of education and including minority institutions.

As the number of college and university computing centers grew, NSF also began to recognize the need for programmers and technicians to staff these centers. In its 1957 Annual Report, it noted:

The rapid development of computing machines and their usefulness in a wide variety of research investigations have created a demand for persons trained in the use and operation of computers. Although such training may be considered a proper responsibility of colleges and universities, there is a severe shortage of teachers competent to give instruction. The Foundation has provided support for a program of training for experienced mathematicians on the faculties of colleges and universities to prepare them to develop courses of instruction in the use and operation of modern computing machines.

In 1954, Wayne University had held a Conference on Training Personnel for the Computing Machine Field⁶⁹ with a focus on educating mathematicians and on scientific rather than business applications of computing. Participants in the 1954 NSF-funded meeting identified a large but unspecified demand for people highly skilled in computation; however, the attendees were unsure whether the primary use of computers was for scientific calculations or business calculations. Educating the needed workforce led to the conclusion that there were “not enough mathematicians.”⁷⁰ Leon W. Cohen, the program director for Mathematical Sciences, made the first public announcement of NSF’s support for computing infrastructure at this meeting.⁷¹

By 1957, NSF was providing support for training experienced mathematicians on the faculties of colleges and universities to prepare them to develop courses of instruction in the use and operation of modern computing machines.⁷² This activity formed the basis for creating academic computer science programs. Training programs continued with the Office of Computing Activities created in 1967.

1.3 Computers in Education

While the Rosser Report said very little about the use of computing for education, the issue did not go away. The President’s Science Advisory Committee (PSAC) commissioned another study of computers in higher education in 1967, chaired by physicist John Pierce of Bell Labs. Following extensive hearings, the committee concluded that “an undergraduate college education without adequate computing was as deficient as an undergraduate education would be without an adequate library . . . [and that] there was value in using computers for precollege education.”⁷³ The Pierce Report’s focus on education supported NSF’s expanded involvement.

Andrew Molnar, a leader in the computing education field, asserted that:

The most significant event [related to computers in education] occurred when President Lyndon Johnson . . . directed the National Science Foundation to work with the U.S. Office of Education to establish an experimental program to develop the potential of computers in education. In response to the directive, NSF created the Office of Computing Activities (OCA) in July of 1967 to provide Federal leadership in the use of computers for research and education.⁷⁴

When OCA was created, Molnar moved over to the NSF from the Department of Education, first on detail and later as a program director, to work on the computers in education programs.

NSF has a long history of involvement in early efforts to use computers for education. It funded three pioneers⁷⁵ in educational technology projects: The Children’s Television Workshop,⁷⁶ the computer-based learning system PLATO, and the curriculum sharing network CONDUIT.

PLATO, the first large-scale, computer-based education system, was developed at the University of Illinois at Urbana-Champaign under the guidance of Donald Bitzer beginning in 1959. With NSF support, Bitzer showed that computers could serve thousands of students, at many different geographic locations, with hundreds of courses, at a reasonable cost. Most of the financial support for PLATO initially came from NSF. Control Data Corp. (CDC) was eventually licensed by the University of Illinois to produce and market the PLATO system.

One unique feature of the PLATO system was a plasma display that provided high quality, low-cost graphics. The PLATO authoring language helped educators create thousands of instructional programs. Bitzer eventually moved PLATO to a Control Data 6000-class machine that served several thousand student stations and provided hundreds of lessons simultaneously. When distributed by Control Data

Corporation, PLATO primarily was used for in-service training in industry, but it continued in use in many universities and secondary schools through the 1980s.

James Johnson at Iowa, Gerald Weeg at Iowa, Thomas Kurtz at Dartmouth, and Jim Parker at North Carolina Educational Computing Service, together with representatives from Texas and Oregon State, formed CONDUIT, a consortium of five regional networks involving approximately 100 colleges and universities for sharing computer-based curricula in seven fields of science.⁷⁷ In 1971, when CONDUIT was conceived, the major barrier to instructional computing was a lack of quality learning materials and computer software. CONDUIT faced significant challenges in validating shared curricula,⁷⁸ but the concept of regional networks would return as a critical part of the NSFNET project.

In addition to computer-aided instruction (CAI) systems such as PLATO and CONDUIT, NSF had an uneven but long history with some of the leaders in the cognitive and learning sciences. As Molnar stated,⁷⁹ “no other name is more closely connected to computer-assisted instruction (CAI) than that of Patrick Suppes.” As Director of the Stanford Institute for Mathematical Studies in the Social Sciences, Suppes began a program of research and development in computer assisted instruction in 1963. He and Richard C. Atkinson, who later would become NSF Director, developed sophisticated mathematical models of student learning to help design instructional materials and strategies.⁸⁰ Suppes noted that John McCarthy of Stanford’s computer science department (having moved from MIT) played an important role in the design and operation of the institute’s computer facilities. Suppes wanted to demonstrate that computers could have an immediate impact on education, even using existing equipment. He and Atkinson began initially with 12 six-year-old children who came to their lab daily and spent 30 minutes at the computer. From 1966 to 1968, Suppes used an IBM 1500 and an audiotape device for CAI. Students responded to questions displayed on a CRT via light pen and keyboard. Suppes later developed a wide variety of CAI courses. The National Science Foundation, the U.S. Office of Education, and the Carnegie Corporation of New York supported Suppes’s research projects.

In 1963 at Dartmouth, John Kemeny and Thomas Kurtz transformed the role of computers in education from primarily a research activity to an academic one. They did not like the idea that students had to stand in long lines with punch cards for batch processing. So they adopted the recently demonstrated concept of time-sharing, which enabled many students to interact directly with the computer. The university developed its own time-shared system and expanded it into a regional computing center for colleges and schools. Kemeny, a mathematician who later became Dartmouth’s president, had applied for an NSF grant to bring a GE-225

computer to campus and to build the first fully functional general-purpose time-sharing system.⁸¹ He received the funding despite reviewers' serious doubts about his plan to employ undergraduates as his research team. Together, Kemeny, Kurtz, and their undergraduate students built a time-sharing system at Dartmouth. At the same time, they developed a new programming language, BASIC (Beginner's All-purpose Symbolic Instruction Code). It turned out to be ideal for introducing beginners to programming and nevertheless was powerful enough to be used for most applications. BASIC worked on any computer. It spread rapidly and was used for the creation of computer-based instructional materials for a wide variety of subjects at all levels of education.

In the early seventies, Seymour Papert at MIT set out to develop a new and different approach to computers in education. He developed a programming language, Logo, to encourage rigorous thinking about mathematics. He wanted it to be accessible to children and be easy to use to express procedures for simple, non-numerical tasks familiar to children. He used it for mathematics education by teaching it in a wide variety of interesting "micro world" environments such as music and physics. Papert insisted that one should not teach mathematics but instead should teach children to be mathematicians. Logo soon became the language of the elementary school computer literacy movement. After OCA was created, the Logo group wanted to do more testing in schools in collaboration with Wally Feurzeig at Bolt, Baranek and Newman (BBN).⁸² The joint project did receive NSF funding, but only following extensive arguments and considerable reservations. NSF was concerned with giving research funding to a private company such as BBN. At the time, NSF preferred a non-profit, research-oriented institute or university such as MIT. "BBN was a suspect as being a money-grabbing kind of place rather than pure as a drift of snow like universities. So, he [the head of OCA, Dr. Milton Rose] said: 'Why should I fund you? You are not a university.'"⁸³ However, Feurzeig's group at BBN was the only group then doing this type of research, and so the NSF obliged. Because of differing viewpoints between the Logo Group's goal to revolutionize mathematics teaching and NSF's focus on educational applications, NSF cut the project's funding in 1977. "These cuts succeeded in allowing the NSF to better control Logo's development as an educational tool rather than a revolution."⁸⁴

In October 1972, OCA's Computer Innovations in Education Section⁸⁵ was transferred to the Education Directorate where funds for research and education started to tighten. To bolster support for their programs, the group decided to support two demonstration projects: PLATO IV⁸⁶ and the Time-shared Interactive Computer Controlled Information Television system (TICCIT),⁸⁷ directed by John Volk of the MITRE Corporation. While PLATO was a large centralized system, TICCIT used a

minicomputer and two-way television in a more distributed system.⁸⁸ The National Science Board, at first skeptical, was impressed with the demonstrations and the result slowed budget reductions *temporarily*.

President Ronald Reagan's fiscal year 1982 budget for NSF included major reductions for education and social science funding. As a result, all funding in the Education Directorate, except for graduate fellowships, was slashed.⁸⁹ Molnar was left to close out all of the existing grants. However, he was able to find ways to fund computers in education researchers. He and Dorothy Deringer, an information scientist from Case Western Reserve serving as an NSF program officer, recruited vendors to donate equipment to NSF and this equipment was made available to researchers. The Education Directorate was eventually restored, and Molnar remained there. Attempts failed to move the computers in education programs into the Computer and Information Science and Engineering Directorate when it was created in 1986. Molnar continued to interact with CISE staff and was involved in the MOSIS VLSI fabrication facility and worked with DARPA and CISE staff members John Lehmann and Bernard Chern to provide access to that system.⁹⁰

1.4 Finding a Home for Computer Science Research

By the late 1950s, the Mathematical Sciences Section was making computer research grants, for example to Delaware, Harvard, Kansas, Michigan State, Michigan, Princeton, Syracuse, and Yale as well as for computing facilities at Northwestern.⁹¹ Grants were later awarded to Oregon State University, Columbia, Delaware, and Rice.⁹²

Under the leadership of Donald Laird, program director for Computer Sciences, and Milton Rose, program director and, later, section head for Mathematics, the NSF program in the mathematical sciences began in the early 1960s to include grants for theoretical symbolic logic, computer sciences, artificial intelligence, and pattern recognition.⁹³ In 1965, 10% of the NSF fellowships in mathematics went to computer scientists; by 1974, the percentage grew to 20%.⁹⁴ The computing facilities and research activities and program managers were transferred from the Mathematical Sciences Section to the Office of Computing Activities when it was created in 1967, with Rose as its head.⁹⁵

The NSF leadership's view lingered that computer science was primarily a form of scientific infrastructure, rather than a discipline in its own right, but OCA fulfilled the hopes of ACM activists by bringing computer science out from under the shadow of mathematics, where its status as a research field had always been in question.⁹⁶

The shift also kept computer science out of the Engineering Division, which had been lobbying since 1965 for control over computing activities. The placement of the Office of Computing Activities under the NSF Director, and its emphasis on education rather than engineering, was a disappointment to NSF's engineers.

The Office of Computing Activities' initial budget was \$22 million, a 73% increase from the \$12.7 million allocated for computer education and research in mathematics and other NSF offices in the previous year. OCA had three sections: the Institutional Computing Services Section (for funding universities to purchase computers as a tool for scientists), led by Kent Curtis; the Special Projects/Computer Innovations in Education Section, led by Arthur Melmed; and the Computer Science Education, Research, and Training Section, led by Fredrick Weingarten. The initial OCA Advisory Committee included a number of leading figures in the developing discipline.⁹⁷

The primary initial role of OCA was to support computing facilities, computers in education, and training of computing professionals. In 1968, Donald Aufenkamp assumed management of the facilities programs and Curtis moved over to head the new Computer Science and Engineering Section with Tom Keenan, John Lehmann, and later Val Tareski as program managers. The concurrent growth in academic computer science programs and researchers led OCA's computing research portfolio to grow. A discipline of computer science was emerging but was not yet sufficiently well-defined to provide an obvious blueprint for the new Computer Science and Engineering (CS&E) Section. With leadership from Rose and input from the advisory committee, Curtis and his program team began to define a set of programs. As Keenan noted:

Well, computer science had achieved the title computer science without much science in it, early. And I think we—here I have to say that Kent Curtis was a prime person . . . I loved the man very much; he was a great guy—we decided that to be a science you had to have theory, and not just theory itself as a separate program, but everything had to have a theoretical basis. And so, whenever we had a proposal, we encouraged, as much as we could, some kind of a theoretical background for this proposal—not just software, and not just write a program, but there should be some basis for it.⁹⁸

The CS&E staff worked together to define a set of programs:

. . . we decided that there was a minimum of three—smallest integer greater than two—things that went to make up computing. The first was theory; the second was hardware; the third was software. So, John Lehmann became the hardware

person. I became the software person. And in the beginning, I think Val Tareski was the theory person . . . each of these programs had probably something less than a million dollars to spend. I think the section had perhaps a \$2 million budget in 1969 or 1970.⁹⁹

The CS&E portfolio of grants, taken together with support from engineering and information science programs, represented a growing investment in the emerging computing research field. Some of the early OCA research grants were awarded to Niklaus Wirth (Stanford), Michael Harrison (Berkeley), Sam Conte (Purdue), Patrick Fischer (Cornell), Juris Hartmanis (Cornell), and Martin Davis (NYU Courant). Computer science-related facilities awards went to Edward Feigenbaum (Stanford), John Pasta (UIUC), Conte, and Richard Conway (Cornell). The Engineering Section in Mathematics, Physical Sciences, and Engineering (MPE) funded Walter Karplus (UCLA), Melvin Breuer (USC), Edward Coffman (Princeton), and Steve Ungar (Columbia), while the Division of Information Sciences funded Vladimir Slamecka (Georgia Tech) and Naomi Sager (NYU).

When Pasta joined NSF in January 1970, he became extremely important in navigating NSF “politics.” Pasta was respected by the senior NSF staff and other division directors in MPE due to his intellect and background in mathematics, physics, engineering, and computer science. This was essential to the growth of computer science funding in competition with other disciplines for budget. His death in 1981 eventually led to the Computer Science Section (CSS) being split off from the Math Section in 1984,¹⁰⁰ as a separate division in Mathematics and Physical Sciences (MPS; by this time, Engineering had become a separate directorate). There was a feeling among many that no one, other than Pasta, had the breadth of background to oversee both mathematics and computer science. His Signal Corps background and his long connection with the Atomic Energy Commission (AEC) and classified projects made it possible for him to play a key role in the conversation between NSF and NSA over cryptography research.

When NSF terminated the computer facilities program, Pasta reorganized OCA into three sections: Computer Science and Engineering, Computer Applications in Research, and Computer Innovations in Education.¹⁰¹ These three new sections reflected the changing nature of computer science and of OCA’s role within NSF.¹⁰² The Computer Science and Engineering Section continued to sponsor research in fundamental computer science, the Computer Innovations in Education Section helped bring the power of the computer to bear on the problems of education, and the Computer Applications in Research Section fostered the development of advanced computer techniques to increase science research capability.

In 1972, Pasta recruited Peter G. Lykos, an Illinois Institute of Technology computational chemist, to NSF with the explicit charge to lead a new initiative to address computer impacts on society.¹⁰³ Lykos was assigned to Aufenkamp's section until he could get the program started. During his tenure at NSF, Lykos experienced turbulent times. OCA had ended its computing facilities program and the computers in education programs were transferred to the Education Directorate. Lykos recalled¹⁰⁴ frustration working with the OCA staff and for the loss of the facilities and later the computers in education programs. He left NSF around the time the Office of Computing Activities was reorganized and transferred to the Research Directorate in November 1973.¹⁰⁵

In 1974, OCA was restructured as the Division of Computer Research (DCR) with Pasta as Division Director. The division¹⁰⁶ supported research in all areas of computing with a major emphasis on fundamental aspects of computer science and engineering (in Curtis's section), on research directed toward the development of techniques that increase the responsiveness of the computer to the requirements of scientific disciplines (in Aufenkamp's section), and on privacy and computer system security, human-machine interface, and societal impacts of computing (in a newly formed section led by Fredrick Weingarten based on Lykos's initiative).

1.5 Summary and Conclusions

The first 24 years of NSF were marked by changing roles and outcomes for its computing and information programs. The Cold War had a strong influence on the science information and computing facilities programs. Interest in foreign intelligence increased the science information budgets. Defense and atomic energy agencies created a rapid growth in the number, capabilities, and providers of computers and computing facilities. Scientists who had limited or no access to Department of Defense (DoD) and AEC laboratories increased their demand on NSF to provide campus facilities. When NSF was given responsibility for applications in science information and computing facilities, the need to provide the underlying technology resulted in NSF investments that advanced fundamental and applied research. Program and office managers in mathematics, engineering, and the OCA began to make grants to the early pioneers in computing research that with DoD support helped establish early computer science programs. OSIS initiated a number of academic information science and systems programs.

By the mid-1970s, OSIS had been greatly weakened and was moved to a non-supportive directorate. OCA lost its facilities and education programs and had yet

to gain the respect of the NSF management. In the 1980s everything would change dramatically.

Much credit for protecting the NSF computing and information programs and building grant portfolios that advanced the underlying technologies is due to a few individuals. Burt Adkinson, the long-term head of OSI and OSIS (1957–1970), was a champion for science information and information science across the government and the discipline. Helen Brownson (1951–1966) was responsible for guiding many of the research efforts funded by OSI and OSIS. Milton Rose (1963–1969), Mathematics Division Director and first head of OCA, recruited to government service a veritable who's who of computing and was a significant force in the rapid development of computing and computer science in academia. Milt was replaced by John Pasta (1969–1981), who led OCA, DCR, and DMCS through many changes and who with Kent Curtis (1967–1987) established the programs that led to the current strong position of NSF in computer science research.

Notes

1. National Research Council. 1999. *Funding a Revolution: Government Support for Computing Research*. Washington, DC: National Academies Press. DOI: [10.17226/6323](https://doi.org/10.17226/6323).
2. D. Kaiser. 2011. The search for clean cash. *Nature*, 472: 7341. DOI: [10.1038/472030a](https://doi.org/10.1038/472030a).
3. J. F. Sargent Jr. and D. A. Shea. 2017. *Office of Science and Technology Policy (OSTP): History and Overview*. Congressional Research Service. <https://fas.org/sgp/crs/misc/R43935.pdf>.
4. National Research Council. 1999.
5. D. L. Kleinman. 1995. *Politics on the Endless Frontier: Postwar Research Policy in the United States*. Duke University Press.
6. V. Bush. 1945. *A Report to the President by the Director of the Office of Scientific Research and Development*. United States Government Printing Office: Washington, DC. <https://www.nsf.gov/od/Ipa/nsf50/vbush1945.htm>.
7. W. A. Blanpied. 1998. Inventing US science policy. *Physics Today*, 51(2): 34–40. DOI: [10.1063/1.882140](https://doi.org/10.1063/1.882140).
8. G. T. Mazuzan. July 1994. *The National Science Foundation: A Brief History*, p. 1. <https://www.nsf.gov/about/history/nsf50/nsf8816.jsp>; last accessed 29 January 2018.
9. National Science Foundation Act of 1950, as amended through PL 112-166, enacted August 10, 2012. legcounsel.house.gov/Comps/81-507.pdf; last accessed 29 January 2018.
10. E. Altman and S. Brown. 1991. From retrieval to robotics: The National Science Foundation's support of information research. *Government Information Quarterly*, 8(4): 359–373.
11. B. W. Adkinson. 1976. National Science Foundation—science information. In A. Kent, H. Lancour, and J. Daily, eds. *Encyclopedia of Library and Information Science: Volume 19*. CRC Press.

12. T. Jayroe. 2012. Homage to Helen Brownson: Information science pioneer. *Proceedings of the American Society for Information Science and Technology*, 49(1): 4-1. DOI: [10.1002/meet.14504901264](https://doi.org/10.1002/meet.14504901264).
13. *National Science Foundation Annual Report*, June 30, 1956. https://www.nsf.gov/pubs/1956/annualreports/ar_1956.pdf.
14. J. H. Shera and J. W. Perry, eds. 1957. *Information Systems in Documentation Advances*. New York: Interscience.
15. Shera and Perry, 1957, *op. cit.*
16. Adkinson, 1976, *op. cit.*, p. 156.
17. Declassified memorandum from the NSA Archives; available in Eisenhower Library. https://www.eisenhower.archives.gov/research/online_documents/declassified_fy_2017.html.
18. Members included Luis Alvarez (University of California, Berkeley), Hendrik Bode (Bell Labs), Richard Garwin (IBM), David Huffman (MIT), John Milnor (Princeton), John Pierce (Bell Labs), Nathaniel Rochester (IBM), Oliver Selfridge (Lincoln Labs), John Tukey (Princeton & Bell Labs), and William Baker (Bell Labs).
19. J. P. Emard. 1976. An information science chronology in perspective. *Bulletin of the American Society for Information Science*, 2(8): 51–56.
20. Adkinson, 1976, *op. cit.*, p. 156.
21. G. Lappan. October 1997. Lessons from the Sputnik era in mathematics education. In National Academy of Sciences symposium titled *Reflecting on Sputnik: Linking the Past, Present, and Future of Educational Reform*, 7: 2003.
22. T. E. Pinelli, M. Henderson, A. P. Bishop, and P. Doty. 1992. *Chronology of Selected Literature, Reports, Policy Instruments, and Significant Events Affecting Federal Scientific and Technical Information (STI) in the United States*. Washington, DC: National Aeronautics and Space Administration. NASA TM-101662: National Defense Education Act (NDEA) (P.L. 85-864) United States Statutes at Large, 72: 1580–1605. http://wwwedu.oulu.fi/tohtorikoulutus/jarjestettava_opetus/Troehler/NDEA_1958.pdf; last accessed 30 January 2018.
23. *NSF Annual Report*, 1959. https://www.nsf.gov/pubs/1959/annualreports/ar_1959.pdf.
24. National Bureau of Standards. *Research Highlights of the National Bureau of Standards Annual Report, Fiscal Year 1959*, Miscellaneous Publication 229. Washington, DC: U.S. Government Printing Office. <https://nvlpubs.nist.gov/nistpubs/Legacy/MP/nbsmiscellaneouspub229.pdf>; last accessed 19 August 2018.
25. L. Fein. 1959. The role of the university in computers, data processing, and related fields. *Communications of the ACM*, (2): 7–14. DOI: [10.1145/368424.368427](https://doi.org/10.1145/368424.368427).
26. Fein, 1959, *op. cit.*, p. 8.
27. Fein, 1959, *op. cit.*, p. 13.
28. Altman and Brown, 1991, *op. cit.*, pp. 359–373.
29. J. Moriarty. 1963. Proceedings of the conference on training science information specialists. *Georgia Institute of Technology, College & Research Libraries*, 24(5): 432–433.
30. R.S. Taylor. 1976. On education. *Bulletin of ASIS*, 2(8): 34.

31. *NSF Annual Report*, 1970. https://www.nsf.gov/pubs/1970/annualreports/ar_1970.pdf.
32. E. Altman. 1993. National Science Foundation's support of information research. *Encyclopedia of Library and Information Science*, 52: 273–291.
33. Adkinson, 1976, *op. cit.*, p. 166.
34. Adkinson, 1976, *op. cit.*, p. 166.
35. *NSF Annual Report*, 1974.
36. Adkinson, 1976,, *op. cit.*, pp. 154–176.
37. *NSF Annual Report*, 1973.
38. *NSF Annual Report*, 1974.
39. *Physics Today*, 1975, 28(9): 80. DOI: [10.1063/1.3069197](https://doi.org/10.1063/1.3069197).
40. Altman and Brown, 1991, *op. cit.*, pp. 359–373.
41. Association of Research Libraries (February 17, 1978) Newsletter no. 90, in *A.R.L. Newsletter*, vols. 95–113.
42. Altman and Brown, 1991, *op. cit.*, pp. 359–373.
43. Association of Research Libraries (February 17, 1978) Newsletter no. 90, p. 8.
44. H. L. Resnikoff. 1980. Information science and technology at the National Science Foundation. *Information Processing and Management*, 16: 243–250. DOI: [10.1016/0306-4573\(80\)90013-8](https://doi.org/10.1016/0306-4573(80)90013-8).
45. Altman and Brown, 1991, pp. 364–365.
46. Association of Research Libraries (April 27, 1981) Newsletter no. 106, in *A.R.L. Newsletter*, vols. 95–113.
47. Office of Scientific Information (OSI), Office of Science Information Services (OSIS), Division of Information Science (DIS), Division of Information Science and Technology (DIST), Division of Information, Robotics, & Intelligent Systems (IRIS), Division of Information & Intelligent Systems (IIS), and the Computer and Information Science and Engineering (CISE) Directorate.
48. H. E. Bamford and C. N. Brownstein. 1986. National Science Foundation support for computer and information science and engineering. *Information Processing & Management*, 22(6): 449–452. DOI: [10.1016/0306-4573\(86\)90095-6](https://doi.org/10.1016/0306-4573(86)90095-6).
49. National Research Council, 1999.
50. *NSF Annual Report*, 1955. https://www.nsf.gov/pubs/1955/annualreports/ar_1955.pdf.
51. *NSF Annual Report*, 1955.
52. *NSF Annual Report*, 1955.
53. *NSF Annual Report*, 1955.
54. Members included Herbert L. Anderson (Chicago), Jule G. Charney (IAS), Gerald M. Clemence (Naval Observatory), Joseph O. Hirschfelder (Wisconsin), Ralph E. Meagher (Illinois), Walter H. Munk (Scripps), J. Barkley Rosser (Cornell), Martin Schwazschild (Princeton), Edward Teller (Livermore Labs), S. M. Ulam (Los Alamos), and John von Neumann (AEC).

55. *NSF Annual Report*, 1955.
56. National Science Foundation. July 1956. *Preliminary Summary: Foreseeable Basic Research Facility Needs for MPE Disciplines*. MPE STUDY S-29.
57. J. von Neumann. 1945. First draft of a report on the EDVAC, 30 June 1945. Contract no. W-670-ORD-492. Moore School of Electrical Engineering: University of Pennsylvania.
58. E. Winsberg. 2010. *Science in the Age of Computer Simulation*. University of Chicago Press.
59. W. Aspray and B. O. Williams. 1994. Arming American scientists: NSF and the provision of scientific computing facilities for universities, 1950–1973. *IEEE Annals of the History of Computing*, 16(4): 60–74. DOI: [10.1109/85.329758](https://doi.org/10.1109/85.329758).
60. *NSF Annual Report*, 1956.
61. T. Dauxois, M. Peyrard, and S. Ruffo. 2005. The Fermi-Pasta-Ulam “Numerical experiment”: History and pedagogical perspectives. *European Journal of Physics*, 26(5): S3. DOI: [10.1088/0143-0807/26/5/S01](https://doi.org/10.1088/0143-0807/26/5/S01).
62. *NSF Annual Report*, 1962. https://www.nsf.gov/pubs/1962/annualreports/ar_1962.pdf.
63. J. Barkley Rosser. 1966. *Digital Computer Needs in Universities and Colleges*. National Academy of Sciences/National Research Council. <http://digitalcollections.library.cmu.edu/awweb/awarchive?type=file&item=33671>.
64. *NSF Annual Report*, 1963. https://www.nsf.gov/pubs/1963/annualreports/ar_1963.pdf.
65. Oral history, Arthur Grad, interviewed by William Aspray, 1990. Charles Babbage Institute. <http://purl.umn.edu/107338>; last accessed 21 August 2018.
66. T. Van Vleck. 2003. *The IBM 7094 and CTSS*. <https://multicians.org/thwv/7094.html>; last accessed 30 December 2018.
67. H. Teager and J. McCarthy. September 1959. Time-shared program testing. In preprints of papers presented at the 14th national meeting of the Association for Computing Machinery, ACM. DOI: [10.1145/612201.612216](https://doi.org/10.1145/612201.612216).
68. J. McCarthy. 1983. *Reminiscences on the History of Time Sharing*. Oxford University Computing Lab. <http://www-formal.stanford.edu/jmc/history/timesharing.dvi>; last accessed 30 January 2018.
69. A. W. Jacobson, ed. 1955. *Proceedings of the First Conference on Training Personnel for the Computing Machine Field*, held at Wayne University, Detroit, Michigan, June 22–23, 1954. Wayne University Press.
70. W. Aspray. 2004. The supply of information technology workers, higher education, and computing research: A history of policy and practice in the United States. In R. Coopey (ed.), *Information Technology Policy: An International History*. Oxford University Press: 54–98.
71. W. Aspray, B. O. Williams, and A. H. Goldstein. 1992. *Computing as Servant and Science: The Impact of the National Science Foundation*. Unpublished manuscript; sections published elsewhere and included in this book.
72. *NSF Annual Report*, 1957. https://www.nsf.gov/pubs/1957/annualreports/ar_1957.pdf.
73. J. Pierce, et al. 1967. *Computers in Higher Education*. The President’s Science Advisory Committee. Washington, DC: U.S. Government Printing Office. <https://eric.ed.gov/?id=ED029518>.

74. Oral history, Andrew Molnar, interviewed by William Aspray, September 25, 1991. Charles Babbage Institute. <http://hdl.handle.net/11299/107509>; last accessed 24 September 2019.
75. Office of Technology Assessment. November 1982. Informational technology and its impact on American education. <https://www.princeton.edu/~ota/disk3/1982/8216/8216.pdf>; last accessed 19 July 2018.
76. The Children's Television Workshop, with funding from NSF, the U.S. Office of Education, the Ford Foundation, and the Carnegie Corporation, created three highly successful educational series: *Sesame Street*, *The Electric Company*, and *3-2-1 Contact*.
77. Molnar interview, 1991, *op. cit.*
78. Molnar interview, 1991, *op. cit.*
79. A.R. Molnar. 1990. Computers in education: A historical perspective of the unfinished task. *THE Journal* 18(4).
80. P. Suppes. May 19, 1971. *Computer-Assisted Instruction at Stanford*. Technical Report No. 174, Institute for Mathematical Studies in the Social Sciences, Stanford University.
81. Dartmouth College. 2014. "BASIC Begins at Dartmouth." <http://www.dartmouth.edu/basicfifty/basic.html>.
82. A. Chakraborty, R. Graebner, and T. Stocky. December 10, 1990. *Logo: A Project History*. <http://web.mit.edu/6.933/www/LogoFinalPaper.pdf>; last accessed 14 February 2019. Also W. Feurzeig. June 1967. New instructional potentials of information technology. *IEEE Transactions on Human Factors in Electronics*, 8(2): 84–88. DOI: [10.1109/THFE.1967.233624](https://doi.org/10.1109/THFE.1967.233624).
83. A. S. Agalianos. 1997. *A Cultural Studies Analysis of Logo in Education*, Ph.D. thesis at the University of London, p. 30. <http://discovery.ucl.ac.uk/id/eprint/10018887>.
84. Agalianos, 1997, *op. cit.*
85. *NSF Annual Report*, 1973.
86. E. McWilliams. 1974. *PLATO IV System Progress Report on Field Testing (Large-Scale Field Test of PLATO Computer-Assisted Instruction System in 1974)*. National Science Foundation. <https://www.atariarchives.org/bcc1/showpage.php?page=31>; last accessed 14 March 2019.
87. E. McWilliams. 1974. *TICCIT System Progress Report on Field Testing (MITRE Corporation Field Test of Time-shared Interactive Computer-Controlled Information Television System of Computer Aided Instruction in 1974)*. National Science Foundation. <https://www.atariarchives.org/bcc1/showpage.php?page=33>; last accessed 14 March 2019.
88. Molnar interview, 1991, *op. cit.*
89. National Science Foundation. March 10, 1981. *A Timeline of NSF History*. <https://www.nsf.gov/about/history/timeline80s.jsp>; last accessed 21 August 2018.
90. Molnar, 1990, *op. cit.*, p. 13.
91. *NSF Annual Report*, 1958. https://www.nsf.gov/pubs/1958/annualreports/ar_1958.pdf.
92. *NSF Annual Report*, 1959.
93. *NSF Annual Report*, 1963.
94. Aspray, 2004, *op. cit.*

95. T. Misa, ed. 2016. *Communities of Computing: Computer Science and Society in the ACM*. Morgan & Claypool (in particular Chapter 2). Also: J. Abbate. From handmaiden to “proper intellectual discipline”: Creating a scientific identity for computer science in 1960s America.
96. Misa, ed., 2016, *op. cit.*, p. 44.
97. Advisory Committee members included Daniel Alpert (UIUC), Gordon W. Blackwell (Furman), Sidney Fernbach (Lawrence Livermore), John G. Kemeny (Dartmouth), Alan J. Perlis (CMU), Jack T. Schwartz (NYU), Jules I. Schwartz (Systems Development Corp.), Harrison Shull (Indiana), and Patrick C. Suppes (Stanford).
98. Oral history, Thomas Keenan, interviewed by William Aspray, September 28, 1990. Charles Babbage Institute. <http://hdl.handle.net/11299/107401>; last accessed 24 September 2019.
99. Keenan interview, *op. cit.*, 1990.
100. J. Minker. Letter to Donald N. Langenberg. June 1, 1981. Attached to the Summary Minutes, Advisory Subcommittee for Computer Science. National Science Foundation, 1800 G. Street NW, Washington, DC 20550 (May 28–29, 1981).
101. R. Coopey. 2004. *Information Technology Policy: An International History*. Oxford University Press.
102. *NSF Annual Report*, 1971. https://www.nsf.gov/pubs/1971/annualreports/ar_1971.pdf.
103. P. G. Lykos, M. G. Morgan, and F. W. Weingarten. 1973. Computer impact on society: Perspective on a new NSF initiative. *Proceedings of the ACM Annual Conference*. New York: ACM, 374–380. DOI: [10.1145/800192.805738](https://doi.org/10.1145/800192.805738).
104. Oral history, Peter Lykos, interviewed by William Aspray, November 13, 1990. Charles Babbage Institute. <http://hdl.handle.net/11299/107447>; last accessed 21 August 2018.
105. *NSF Annual Report*, 1973 and 1974.
106. *NSF Annual Report*, 1974.



1974–1986: CER, CSNET, NSFNET, and the Founding of CISE

W. Richards Adrion

As we discussed in Chapter 1, computing and information programs and activities existed from the beginning of the National Science Foundation. After several major NSF reorganizations, the computer science and engineering research programs in the Office of Computing Activities were transferred to the Research Directorate in 1974 and the Office was renamed the Division of Computer Research (DCR) in 1975. After the Research Directorate was divided into several discipline-based directorates, the DCR programs were moved into the Computer Science Section of the Mathematical and Physical Sciences, and Engineering (MPE) Directorate in 1976.¹ Programs for scientific computing resumed in the early 1980s as support for high-performance computing, and then in the 2000s for “cyberinfrastructure.” Educational applications of computing moved to the Education Directorate in 1972 and, following a brief hiatus during the Reagan administration, remained there. The programs in the Office of Science Information Services (OSIS) moved to the Directorate for National and International Programs in 1969, where they suffered substantial reductions in funds and significant changes in staffing. The NSF science information/information science programs evolved to focus on essential technologies for addressing fundamental questions of information science.

By the 1980s, NSF programs supporting computer science, computer engineering, and information science research had moved from the administrative side of NSF to, or were created within, various divisions and sections in the research directorates. Computing research was housed in Mathematics and Physical Sciences (MPS). The Division of Information Science and Technology was moved to

the Biological and Behavioral Sciences Directorate in 1978. After an Engineering and Applied Science Directorate was created in 1978 (becoming the Engineering Directorate in 1980), NSF developed explicit programs for computer engineering and housed them in a new Electrical, Computer, and Systems Engineering Division. A new office of Advanced Scientific Computing was created in 1984 to meet the demand for supercomputer centers and associated networking access. The formation of the Computer and Information Science and Engineering (CISE) Directorate in 1986 brought these programs together in a single directorate.

2.1 My Background and Perspective on the 1974–86 Period

Much of this chapter is based on my experience and memory of events, augmented with documents and references. I first joined the National Science Foundation in late summer 1976 and for two years was the program director for the Theoretical Computer Science (TCS) program.² I will describe the creation and operation of the Computer Science Section (CSS) within NSF, issues that arose around cryptography research funded from the TCS program, and the roles of the CISE Equipment program and Theorynet in influencing the Coordinated Experiment Research (CER) initiative.

I returned to NSF in January 1980 as the program director for Special Projects in the Computer Science Section (CSS). My responsibilities included the new Coordinated Experimental Research initiatives: CER (facilities), CSNET, a New Faculty Investigators program, and a Postdoctoral program. In FY 1981, CER and CSNET and the New Faculty Investigators program became separate programs, while the Postdoctoral program was terminated. I managed CER and oversaw C. William Kern, the CSNET project manager. I assumed the role of CSNET project manager in 1982 when Kern left for Ohio State. I was also responsible for other programs in Special Projects including research on databases, privacy and security, and social impacts as well as conferences, symposia, and special studies.

In 1984, I joined the Office of Advanced Scientific Computing as program director for Networking while maintaining responsibility for Special Projects programs, CER, and CSNET. I was on an Independent Research and Development (IR&D) assignment at the University of California, Berkeley, for the 1984–1985 academic year, handing over the CER program to Harry Hedges and the Special Projects program to Larry Oliver. I continued to manage the OASC Networking and the CSNET programs until Dennis Jennings took over the OASC Networking program in January 1985 and when CSNET had become more or less independent under management by the University Center for Atmospheric Research (UCAR) and BBN. I will describe

the early efforts for “Sciencenet” that led to NSFNET and the successful spinoff of CSNET and its eventual merger with BITNET to form the Corporation for Research and Educational Networking (CREN).

While I was at Berkeley, I was hired as the Deputy Division Director for the new Division of Computer Research (DCR), which had split off from Mathematics but remained in MPS. When I returned to Washington in the fall of 1985, I had mostly administrative duties in DCR, including upgrading the computing infrastructure within DCR and working with Connie McLindon, the NSF CIO, on NSF-wide technology.

During fall 1985, I also began working with Chuck Brownstein, Division Director of Information Science and Technology, to assist Director Erich Bloch with plans to develop a full-blown computing directorate. In March 1986, he announced that Gordon Bell would be joining NSF to lead the effort. A week or so earlier, Bell had requested that Brownstein take on the role of Executive Officer of the new directorate, Jerry Daen be added as the Planning and Administrative Officer, and I join half-time on loan from DCR. Eventually, I became the Senior Scientist for the Computer and Information Science and Engineering (CISE) Directorate. I will describe the negotiations and planning that went into the first nine months of the CISE Directorate.

I returned to NSF in January 2000 as Division Director for Experimental and Integrative Activities (EIA). Chapter 8 includes a description of the President’s Information Technology Advisory Committee (PITAC) report that led to the government-wide initiative on Information Technology for the 21st Century (IT²), the designation of NSF as lead agency, and the planning and experiences that led to the NSF implementation of IT², the Information Technology Research (ITR) program.

In addition, I served on a number of advisory committees and was involved in three more reorganization efforts: chairing the NSF/CISE Committee on CISE Organization in 1995–1997 for Paul Young, chairing the divisional NSF/CISE/EIA Reorganization Working Group in 1997–1998 for Juris Hartmanis, and—as a part-time CISE senior advisor—chairing a committee that advised Peter Freeman on his 2003 reorganization.

2.2 Making the Case for NSF's Computing Research Programs

NSF provided funding for computing, communications and information infrastructure, applications, and fundamental research from its beginning. The physical scientists who ran the NSF were not quite sure there was a “discipline” of computer science, but they clearly appreciated the growing importance of computing,

communications and information infrastructure, and applications. Scientific and engineering disciplines typically turned to related professional societies or the National Academies to describe the field, its accomplishments, and its future promise. An influential report was needed to define computer science, its value to the nation, and the need for investment and support.

The professional societies—ACM, IEEE-CS, AFIPS, SIAM, and AAAI—established the conferences and journals in this new field. None of them adequately represented academic computer science research in Washington, DC. This gap led to the creation of the Computer Science Board in 1972, later renamed the Computing Research Association (CRA), which created a Washington presence in 1988. Ever since, CRA has played an important role in advocating for the computing research community.

From 1978 to 1986, the National Academy Board on Telecommunications and Computer Applications primarily published reviews of information technology issues and challenges experienced by federal agencies such as the Social Security Administration, the Internal Revenue Service, NASA, and the Departments of Defense and Commerce. One exception was a 1982 report from an ad hoc committee on the roles of industry and the university in computer research and development.³ The National Research Council created the Computer Science and Telecommunications Board (CSTB) in 1986 to replace the Board on Telecommunications and Computer Applications.

Earlier in the 1960s, a number of individuals attempted to define computer science as a discipline. In addition to Louis Fein's⁴ efforts described in Chapter 1, Saul Gorn of the University of Pennsylvania wrote in 1963 that “a new basic discipline is emerging which might be called ‘The Computer and Information Sciences’ [that] makes application of concepts from the traditional fields of mathematics, philosophy, linguistics, psychology, engineering, management science, library science, etc.”⁵ George Forsythe,⁶ the founder of Stanford's computer science department and ACM President, commented on Gorn's analysis, suggesting that computer scientists are concerned with the pragmatics of the applications of mathematics. In 1967 Allen Newell, Alan Perlis, and Herbert Simon⁷ defined computer science as the study of phenomena related to computers. Donald Knuth's definition⁸ of computer science as the study of algorithms appeared in 1968. Curriculum 68⁹ defined computer science as the study of information structures. Edsger Dijkstra defined computer science as the study and management of complexity.¹⁰ Historian Janet Abbate observed that computer scientists, in arguing for scientific status of their field, drew on “three distinct meanings of science (sometimes in combination)”¹¹: (1) science as the study of natural phenomena (information in this instance),¹²

(2) science as the derivation of abstract ideas from concrete phenomena,¹³ and (3) the experimental method as the defining characteristic of science.¹⁴

These assertions about computer science as a science did not persuade NSF management that computing was or was beginning to be a mature scientific discipline. Abbate notes: “. . . organizational control wielded by the established disciplines, as well as NSF's emphasis on basic research, put the emerging field of computer science at a disadvantage. In this context, the notion of computing as a 'science' and the appropriateness of NSF funding for computing researchers were both contested.”¹⁵

After NSF moved the Office of Computing Activities into the Research Directorate, renaming it the Division of Computing Research (DCR) in 1974, the weak support for computer science as a discipline resulted in DCR programs being placed in a section (CSS) within a Mathematical and Computer Sciences Division in the Mathematical and Physical Sciences, and Engineering Directorate in 1976. When DCR was created, Gordon Bell, then with Carnegie Mellon University and Digital Equipment Corporation, was “concerned about funding for computer science within the National Science Foundation and that we [the computer science community] lack representation on the National Science Board.”¹⁶ Saunders MacLane, a Chicago algebraist on the National Science Board (NSB), was a good supporter of computer science but not a true representative of the discipline. NSF provided a 12.2% increase for Computer Science research for FY 1976, while MPE overall was increased 6.3%. The \$13.22 Computer Science research budget, however, was only 6.6% of the total MPE budget.

To offset the perception that computing research was well-served by industry, Bell argued that funding for basic research in computing should be directed to universities and not industry. Bell added that while mission agencies, such as ARPA, played a significant role, NSF had the role of supporting basic computer science research. Bell also suggested that NSF funding of basic computer science research introduce a “question of scale” and that NSF consider investments of an ARPA-like magnitude in several non-ARPA-funded, leading computer science programs.¹⁷

Facing skepticism from NSF leadership about the emerging field of computer science and its core research questions, John Pasta and Kent Curtis mobilized influential scientists. In 1974, they funded the Computer Science and Engineering Research Study (COSERS) under the direction of Bruce W. Arden of Princeton University. “For the first time in its quarter century of activity . . . this discipline will be given a comprehensive examination by researchers in the field. . . . The report will define what computer science and engineering is, describe major research problems now under investigation, and point out future educational and research

opportunities.”¹⁸ Apart from a brief progress report¹⁹ in 1976, the massive 1000+ page report, *What Can Be Automated?: Computer Science and Engineering Research Study*,²⁰ unfortunately did not appear until 1980. By that time, other influential reports had appeared and diminished its impact.

By the late 1970s, Curtis and Pasta were working with leading members of the computer science community to address a serious concern about the health of academic computer science. Academic salaries were falling behind industry salaries, there was a significant lack of computing equipment except at the ARPA-funded departments, undergraduate enrollments were rising, and many scientists, engineers, and policymakers still viewed computer science as consisting only of programming, computing applications, and hardware development. Faculty, new PhDs, and promising graduate students were leaving academia for industry in large numbers.

The NSF sponsored a workshop in Washington, DC, on November 2, 1978, led by Jerome Feldman (Rochester) and including Gordon Bell (DEC), Bernard Galler (Michigan), Patricia Goldberg (IBM), John Hamlin (Missouri), Eliot Pinson (Bell Labs), Ivan Sutherland (CalTech), and William Robert “Bert” Sutherland (Xerox PARC). The “Feldman Report,”²¹ also published in the *Communications of the ACM*, called for universities to recognize the special resource needs of experimental computer science, use appropriate criteria in evaluating experimental computer science programs and faculty, and encourage cooperative programs. While it called for industry to exchange and share people and technology with universities and provide funds and equipment, it looked to government to modernize tax and patent policies, develop funding of adequate scale and time-horizon for experimental computer science, and identify a lead agency responsible for computing. The report proposed large, 5-year capital resources to produce 25 well-equipped university laboratories for a total cost of about \$15 million yearly. This recommendation led to the Coordinated Experimental Research (CER) program described below.

The Feldman Report was enthusiastically expanded upon by the ACM Executive Committee,²² the Computer Science Board (sponsor of a 1980 meeting of computer science department chairs²³ now known as the biennial Snowbird Meeting), the 1979 NSF Computer Science Advisory Committee,²⁴ and a series of ACM letters and articles by Peter Denning that began with his well-known “eating our seed corn” letter.²⁵ For example, a panel at the 1980 Snowbird Meeting²⁶ addressed the nature of computer science, advances in computer technology, and how computer scientists might address societal implications. These panels and reports changed the perception of computer science within the NSF and across the federal government.

2.3 The Importance of Computing Research and Infrastructure to the Nation

For a time after the Computer and Information Science and Engineering (CISE) Directorate was created in 1986, the computer science advisory committees remained associated with the CISE division that Kent Curtis directed. At the request of the NSF Advisory Committee for Computer Research, a subcommittee appointed by Curtis submitted a preliminary report²⁷ at the committee's meeting held on December 5–6, 1986. This report was revised, published in 1989, and became known as the “Hopcroft-Kennedy Report.”²⁸ Kent Curtis passed away December 17, 1987, and the Hopcroft-Kennedy Report was completed after Peter Freeman had replaced Curtis as DD/CCR.

Frank Press, then president of the National Academy of Sciences and chairman of the National Research Council, created in 1986 what came to be called the Computer Science and Telecommunications Board (CSTB), with Joseph Traub as chair and Marjory Blumenthal as executive director. Under Traub, Blumenthal, and their successors, CSTB published many influential (and occasionally controversial)²⁹ reports. Their first efforts did not try to identify the achievements and opportunities of computing research as did the Hopcroft-Kennedy Report. Traub noted, “CSTB decided that beginning with a report on the nature of the field would be self-serving. We wanted first to build a record of reports dealing with critical national issues.”³⁰ The Board published *Toward a National Research Network*³¹ in July 1988 and *The National Challenge in Computer Science and Technology*³² in September 1988. Among the many CSTB reports are ones the Academy characterizes as “explaining how information technology evolves, the role of R&D, and the role of different contributors, public and private, to that process.” These include³³ *Innovation in Information Technology*, *Making IT Better*, *Funding a Revolution*, *Evolving the HPCCI to Support the Nation's Information Infrastructure*, and *Computing the Future*. While clearly influential, one of the criticisms of the Academy and the National Research Council, as voiced by Ed Feigenbaum (perhaps a bit sharply), is that they are “extremely slow and conservative organizations, unwilling to say things that make anyone bristle. So, a lot of what CSTB might try to do is either squashed or squashed in advance by this elaborate structure.”³⁴ Until the CRA was chosen to create the Computing Community Consortium in 2006, the options for fast response “blue ribbon” reports remained limited.³⁵

Beginning in the 1970s when NSF reduced its support for computing facilities, concern grew in the scientific community that future scientific advances would be impeded by the lack of advanced computers. Moreover, a number of countries

were developing “supercomputers”³⁶ and programs to increase access³⁷ for their scientists. An interagency study group, led by Peter Lax of NYU, made the case³⁸ (known as the “Lax Report”) for a program that would increase access to supercomputers via high bandwidth networks; increase research on computation, software, and algorithms; train personnel; and increase R&D on new supercomputer systems. Ken Wilson, then at Cornell and a recent Nobel laureate, was one of the leading proponents of a program in supercomputers and a national network to connect them. In an attachment to the Lax Report, Wilson stated that “the lack of large scale scientific computing resources for basic university research has become a major problem for U.S. science.”³⁹ In this, he advocated for a national network linking all scientists and support for a collaborative program in large-scale scientific computing hardware, software, and algorithms led by the science, computer science, and electrical engineering communities and industry. As part of his advocacy, Wilson coined the term *grand challenges*.

A four-part federal program was proposed⁴⁰ and, in mid-1983, an internal NSF working group, led by Marcel Bardon and Kent Curtis, laid out specific actions they recommended that NSF take (the “Bardon-Curtis Report”).⁴¹ These actions included providing “supercomputer services for academic research and science education . . . ” and supporting “networks linking universities and laboratories with each other. . . . ” Following a report⁴² from a panel on “Computer Architecture,” led by Jack Schwartz of the NYU Courant Institute on behalf of the National Academies’ Committee on Science, Engineering, and Public Policy (COSEPUP), NSF Director Erich Bloch asked the engineering and physical sciences directorates and the newly formed Office of Advanced Scientific Computing to comment on Schwartz’s suggestion that NSF “might strive for a position of higher importance and impact” in high-performance computing.⁴³ The response recommended that reaching the Schwartz panel’s recommendation would “require a well coordinated federal effort among at least the following agencies: DOD (including DARPA, ONR, AFOSR, and ARO), DOE, NASA, NBS, and NSF . . . [and] it is appropriate that NSF provide more leadership because of its independence from mission criteria in selecting research projects for support and because of the excellent technical judgment it can bring to bear.”⁴⁴

The emphasis on networking in the Lax and Bardon-Curtis reports led to a series of reports on networking including the Sciencenet⁴⁵ proposal and the initial ideas for NSFNET.⁴⁶ These and CSTB reports provided background for the Federal High-Performance Computing program and NSF’s programs in advanced scientific computing and networking.

2.4 Computing and Information Research in NSF, 1974–1978

By the late 1970s, the programs that would be joined to form the Computer and Information Science and Engineering (CISE) Directorate were in place but divided among several NSF research directorates. In 1974, NSF transferred the Office of Computing Activities to the Research Directorate and renamed it the Division of Computer Research (DCR). In 1976 DCR merged its sections and programs into the Computer Science Section (CSS) of the Division of Mathematical and Computer Sciences (DMCS) within the new Mathematics, Physical Sciences and Engineering (MPE) Directorate.

DCR had two sections from 1974 to 1975: computer science and engineering, and computer applications in research. The former ran programs in theory, programming languages and systems, and systems design. The latter ran programs in techniques and systems, software quality research, and networking for science. The FY75 NSF Annual Report includes these comments:

The discipline of computer science is barely 10 years old, only vaguely defined, and mushrooming. . . . In a field as new and as rich as computer science it is not surprising that new areas appear, create a flurry of activity, and then level off or stagnate; automata theory, mechanical translation, and theory of formal languages are a few such . . . researchers in computer science are anxious to follow new leads into uncharted regions. This kind of process of extension to new areas and pruning of less productive ones partly accounts for the lack of definition of the field.⁴⁷

The report goes on to suggest that the Arden COSERS initiative, described above, was a necessary disciplinary self-examination. By the time of my arrival at NSF,⁴⁸ toward the end of Transition Quarter 1976,⁴⁹ the Assistant Director for the Mathematical and Physical Sciences and Engineering (MPE) Directorate, Ed Creutz, had decided to merge computer research with mathematics.

John Pasta became Division Director for the Division of Mathematical and Computer Sciences (MCS). The three sections in DCR (Computer Science and Engineering, Computer Applications, and Computer Impact on Society) became one section within MCS. Kent Curtis, who had been section head for Computer Science and Engineering (CS&E), became section head for the Computer Science Section (CSS). William H. Pell led the Mathematics Section. Don Aufenkamp, who had been section head for applications, took over the NSF US-USSR program. The program directors in the DCR Computer Science and Engineering Section—Bruce Barnes (Theory), Thomas Keenan (Programming Languages and Systems), and

John Lehman (Computer Architecture)—moved with their programs into CSS. Sally Sedelow from Techniques and Systems became the Intelligent Systems program director in CSS. Fredrick Weingarten became Special Projects program director. Walter Sedelow came over from the applications section, where he had overseen computer networking-related grants, to join Weingarten in Special Projects. While Kent had recruited me for the Software Engineering program, he decided to have Bruce Barnes head that program because of his experience and interest. I was assigned instead to the Theoretical Computer Science program. The Sedelows left in 1977, and Sally was replaced by Eamon Barrett (from ESL Inc.)⁵⁰ and Walter by Larry Oliver (from NSF Education).

Engineering, which also was a division in MPE in 1976, had an Electrical Sciences and Analysis Section, which funded research on digital systems and communications, and information theory. Later, after a possibility that a separate National Engineering Foundation might be created, NSF merged applied research and engineering to create a new Engineering Directorate with an Electrical, Computer, and Systems Engineering (ECSE) Division. Steve Kahne, Thelma Estrin, and others served as ECSE division directors. The Division of Science Information in the Scientific, Technological, and International Affairs (STIA) Directorate supported fundamental research on information sciences and applied research on information access and user requirements. This division would later be renamed the Division of Information Science and Technology and moved to the Biological and Behavioral Sciences (BBS) Directorate.

The new Computer Science Section had six programs—Theoretical Computer Science, Software Systems Science, Software Engineering, Intelligent Systems, Computer Systems Design, and Special Projects—each described in the NSF Guide to Programs as shown in Figure 2.1 (before Software Engineering was added). The programs had no deadlines, target dates, or solicitations; and all proposals were essentially “unsolicited” without restrictions on page length, format, font size, etc. Prospective principal investigators were encouraged to submit proposals in the fall if they wanted summer funding for the following year.

William Aspray writes in Chapter 7, “Foundation staff did not generally set a research agenda for funding. They relied instead on the scientific community to set the agenda, both through the proposals individual scientists submitted and the reviews the scientific community gave to these proposals.” I would argue that, while we placed no constraints on what could be submitted and solicited no proposals, the program directors, Kent Curtis, and John Pasta were very proactive in encouraging people to submit and in publicizing the programs. The proposals the section funded and the people we encouraged, in effect, defined an agenda.

Theoretical Computer Science—The theory of computation, numerical analysis and computational mathematics, theory of formal languages, analysis of algorithms, and other topics concerned with the theoretical foundations of computer science.

Software Systems Science—Fundamental questions of communicating with and controlling computer systems. Topical areas include advanced procedural and nonprocedural languages, the semantics of programming languages, information structures, file management and data base systems, control and allocation of computing resources, and other topics concerned with the structure and representation of numeric or non-numeric software.

Software Engineering—The methods, tools, and techniques for specifying, designing, and implementing quality software. The program scope includes development of prototypes or experimental implementations where these are integral parts of the research program, and verification, testing, portability, reliability, and human interfacing to numeric or non-numeric software systems.

Intelligent Systems—Computer-based systems, which have some of the characteristics of intelligence. Relevant areas include pattern recognition, pattern generation, knowledge representation, problem solving, natural language understanding, theorem proving, and others, which relate to the automatic analysis and handling of complex tasks.

Computer Systems Design—The principles of computer systems design, including computer system architecture, performance, graphics, man-machine interactions, logic design, and others, which relate to the structure of computer systems or the process of systems design. The program may include experimental implementation where that is an integral part of the research.

Special Projects—Research projects, studies, workshops, and other activities, which encourage the development of new fields of computer science research that are responsive to the problems and opportunities arising from the widespread use of computers in society.

Figure 2.1 FY 1977 NSF guide to programs: Computer science research. (Source: National Science Foundation (1976) Guide to programs, FY 1977.)
<https://hdl.handle.net/2027/mdp.39015043526683>; last accessed 6 June 2018.

Before FastLane⁵¹ made web-based submissions possible, proposals were mailed to NSF with approximately 25 copies arriving in the one office that processed all arriving proposals. After the CSS administrative officer picked up the proposals from central processing and distributed copies to the program officers, they would do a quick check on the appropriateness and redistribute if needed. Since the volume of applications was modest,⁵² program directors took time reading each proposal in detail and consulting colleagues for suggested reviewers. One also could walk down the street to the George Washington University Library (or use the much smaller NSF library) to read related or cited papers to help in understanding the proposals and selecting reviewers.

Typically, a program director needed three to four reviews to support a recommendation. Proposals were sent to six to eight reviewers, given the low response

rate. These reviews were carried out as “mail reviews,” that is, copies of the proposals along with a review form and check boxes for an “adjective” review (poor, fair, good, very good, excellent). Proposals were triaged: the clearly fundable proposals were recommended as soon as possible, the clearly non-fundable proposals were declined, and the remaining were held for discussion in weekly meetings with all six program directors, Kent Curtis, and often John Pasta. In these meetings, we discussed the status of our programs and the awards and declinations we were planning. These were often lively discussions about priorities and high-risk proposals.

The primary issue delaying recommendations was the time it took to get solid reviews. “We read the comments very carefully, used our best judgement, and did not really put much weight on the adjective ratings.”⁵³ The directorate, however, did consider the ratings and compared our recommendations against the other programs in MPE/MPS. The field was young and the “shooting inward”⁵⁴ phenomenon was at its height. Our first strategy was to plead with the researchers to evaluate proposals fairly and to understand that there were risks that could be overcome with good new approaches. Our second strategy to address both response rate and review quality was to employ John Lehmann’s skill at mining the NSF databases. We gathered data for every reviewer on the time to review, the number of reviews, and the average review, and compared their performance with other reviewers of the same proposals. So, if Mary Smith seldom gave “excellent” ratings and typically gave ratings below those of other reviewers, we could use that in the recommendation. Our next strategy was to remove the adjective ratings from the review forms entirely. This had two good outcomes: it left interpretation more to the program directors rather than depending on scoring, and the lack of the option to just check a box resulted in longer and more thoughtful reviews. In the long run and because of a desire to have uniform measures across the Foundation, however, we were asked to return to using adjective reviews.

Computing research funding rose relatively slowly over the period 1974–1980 (see Figure 2.2) with the first significant increase coming with the establishment of the Coordinated Experimental Research Programs (described below) in 1980. There were several ways in which we managed our program portfolios.

Once an adequate number of reviews arrived, we would seek out other program managers in computer science, mathematics, engineering, or information sciences to discuss “split funding” if appropriate. Typically, there was little budget gain from these transactions inasmuch as we might co-fund as many proposals in other programs and divisions as we would get co-funding from them. It did contribute, however, to a broader understanding of computing and computer science and, as I will discuss, some quite important joint-funded grants were made. One

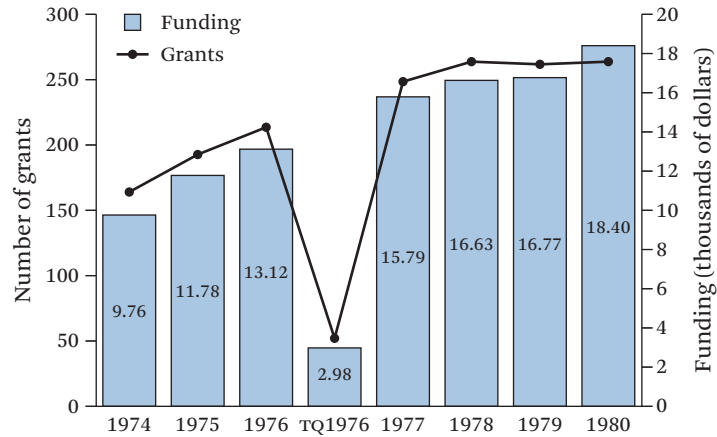


Figure 2.2 Computing research funding FY 1974–1980.

important feature of the 1800 G Street NW NSF headquarters building was a “senior staff lunchroom” on the 12th floor, where program officers would grab lunch and join program officers from other offices and directorates at the few available tables. These casual meetings led to collaborations, joint funding, and collegiality. Unfortunately, due to its size, entrance to the lunchroom was limited by grade level, thus barring junior program officers and clerical and administrative staff. Erich Bloch closed the lunchroom for just this reason.

I believe the process I have described led to thorough and thoughtful reviews and recommendations, which corrected the perception that computing research proposals were of comparatively lower quality. The number, breadth, and quality of the research the Computer Science Section supported under its constraints and with limited funds demonstrates an effective stewardship of NSF investments in a growing field.

2.5 Funding the Innovators in Computer Science

It is not easy to measure the impact of individual funding decisions on the health and growth of computer science. One indicator might be the role NSF played in the careers of Turing awardees. The A. M. Turing Award is the oldest and most prestigious award⁵⁵ in computing. It is presented annually by the Association for Computing Machinery (ACM)⁵⁶ “to an individual who has made lasting contributions of a technical nature to the computing community.”

Many of the Turing Award winners from the 1960s and 70s were in industry (Maurice Wilkes, Richard Hamming, Charles Bachman, John Backus, and Kenneth Iverson), Europe/Israel (Wilkes, Jim Wilkinson, Edsger Dijkstra, Michael Rabin), or the (D)ARPA-funded universities (Alan Perlis, Marvin Minsky, John McCarthy, Don Knuth, Allen Newell, Herb Simon, Dana Scott, and Bob Floyd). However, Don Knuth received significant NSF funding for the work that went into his *The Art of Computer Programming*⁵⁷ series and the development of T_EX.⁵⁸ Both Alan Perlis and John McCarthy were involved with NSF facilities grants in the 1960s, which provided them with an environment for their early work on programming languages and operating systems. John McCarthy was funded by multiple NSF programs during the later 1970s.

During the period from 1976 to 1978, the Computer Science Section launched the research careers of many future Turing Award winners. Of the winners from the 1980s through 2017, again some spent most or all of their careers in industry/government or outside the United States.⁵⁹ From 1976 to 1978, the Theoretical Computer Science (TCS) program made grants to Richard Karp, John Hopcroft, Robert Tarjan, Juris Hartmanis, Manuel Blum, Amir Pnueli (as a visitor at University of Pennsylvania), Andrew Yao, Leonard Adelman, Ronald Rivest, Adi Shamir, Judea Pearl (with Intelligent Systems), Martin Hellman, and Whitfield Diffie (Hellman and Diffie with Engineering Systems). Michael Stonebreaker was funded from the Special Projects program in 1980. At a later time, the TCS program funded Leslie Valliant and Shafi Goldwasser. Edward Clarke, Alan Emerson, Barbara Liskov, and Leslie Lamport all received NSF grants from the Software Systems Science program.

Many of the Turing Award winners, including those in industry, played significant roles in advising and advocating for computer science within NSF, including John Hopcroft and Fredrick Brooks, who both served on the National Science Board (NSB). Vinton Cerf and Robert Kahn were important to the development of CSNET and NSF, and Cerf recently served on the NSB.

In addition to the Turing awardees, there were other important grants made in the 1976 to 1980 period. I discuss cryptography and security below, including the work of Hellman and Diffie (Stanford); Rivest, Adelman, and Shamir (MIT); George Davida (Wisconsin Milwaukee); and Dorothy Denning (then at Purdue and SRI). Lawrence Landweber's Theorynet project and the concurrent analysis of collaboration over networks by Starr Roxanne Hiltz were important in building support for CSNET and later NSFNET. The Coordinated Experiment Research (CER) program addressed the national issue described by the Feldman and Snowbird reports, but the successful grantees would not have succeeded without the equipment grants

(typically VAXes and PDP-11s), which initiated experiment work in the grantee departments and established their credibility as potential centers of experimental research.

In Chapter 7, William Aspray describes research done by some of the most respected, NSF-supported scientists. During 1976–1978, most of these people were funded by the Computer Science Section. Mary Shaw and Barbara Liskov were among many influential women researchers funded by CSS in the late 70s—a group that included Sue Graham, Sherry Turkle, Irene Grief, Lori Clarke, Anita Jones, Mary Jane Irwin, Ruzena Bajcsy, Nancy Lynch, Diane O’Leary, Shari Pfleeger, Elaine Cohen, Sheila Griebach, Dorothy Denning, and Naomi Sager.

Additional grants from this period illustrate the impact of the Computer Science Section. The work of Arthur Burks and John Holland became the basis of classifiers used in machine learning. The Stanford AI lab (with John McCarthy, Edward Feigenbaum, and Cordell Green) moved artificial intelligence ahead. The Ingres Relational Database developed by Michael Stonebreaker, Lawrence Rowe, and Eugene Wong was arguably the first practical research relational database. Concurrently, the Division of Information Science and Technology funded early work in Information Retrieval by Gerald Salton, Naomi Sager, Michael McGill, and several others.

2.6 Facilities

In his book on applications of case study research, Robert Yin took David Gries’s abstract from the final report on the first Coordinated Experimental Research (CER) grant to Cornell and analyzed it as a case study. Yin quotes Gries’s final report and identifies the outcomes:

From 1980 to 1986, the Computer Science Department at Cornell was radically transformed from a theoretical, pencil-and-paper research operation to one with a high degree of experimental computing. The departmental computing facility grew from a VAX780 and a PDP11/60 to an integrated complex of almost 100 workstations and UNIX mainframes. All faculty and graduate students now use computing daily, and much research that was hitherto impossible for us is now being performed.

The CER grant enabled the department to attract bright young faculty who would not have joined a department with inadequate facilities. As a result, the department has been able to branch out into new areas, such as VLSI, parallel architectures and code optimization, functional programming, and artificial intelligence. The CER program did what it set out to do: It made it possible for the

department to expand its research activity, making it far more experimental and computing intensive while still maintaining strong theoretical foundations.⁶⁰

The CER program was transformative in the ways that Gries describes. Earlier support for the VAX780 and a PDP11/60 likely came from the Computer Science Research Equipment program. The program made seven grants in FY 1977, totaling \$753,200, to departments that later received CER grants: Cornell, Arizona, Illinois, Pennsylvania, Utah, Washington, and Wisconsin. It also made grants to other departments, totaling \$817,700, some of which competed unsuccessfully in the CER program. In FY78 the program made eight more grants, totaling \$790,249, to departments that later received CER grants: Brown, Stony Brook, Berkeley, Illinois, UMass, Utah, and Wisconsin; and additional grants totaling \$755,403. In many ways, the equipment program was as important to the computer science community as CER, moving departments from “pencil and paper” to a point where they had facilities for experimental research. The program continued, under various titles, from the 1980s until 2001, when the CISE Research Instrumentation program was incorporated into a revised CISE Research Infrastructure program along with the successors to the CER program. I will return to the CER infrastructure programs below.

There were attempts to develop a national research network, or regional ones, prior to the Office of Computing activities move to the Research Directorate. Don Aufenkamp, then head of the OCA Applications in Computing Section, announced at the 1972 EDUCOM meeting that NSF was going to sponsor research that might lead to a network linking universities and other institutions.⁶¹ He and Ed Weiss delivered a paper⁶² at the International Conference on Computer Communication in October 1972 discussing further details. In *Science* in October 1973, Greenberger et al. noted that NSF had funded

. . . EDUCOM to bring together interested users and administrators with those possessing shareable resources and relevant experience in a series of three 2-day working seminars. . . . The seminars . . . were designed to help identify the central organizational, political, and economic issues in building and operating networks on a national basis.⁶³

What happened to this effort is not at all clear. It may have been inspired by the success of the CONDUIT regional networks described in Chapter 1, but with a broader national vision. Historian Janet Abbate⁶⁴ speculated that it was because the Office of Computing Activities (and its successor, the Division of Computing Research) had a limited budget or because the importance to researchers was not

yet realized. When I arrived in 1976, NSF leadership was not interested in anything of the scale and management demands of an ARPANET-like national network and remained unconvinced that a network for sharing resources and collaboration had value, given the cost.

An opportunity for NSF to be involved in networking arose in 1977. Fredrick Weingarten inherited what was left of the DCR applications efforts in his Special Projects program. I knew that he was supporting research on the economics and social impacts of networks and computer-based collaboration, often jointly with Edward Weiss and others in the Division of Information Science and Technology. An opportunity surfaced at the 1977 Foundations of Computer Science (FOCS) conference in Providence, RI, where I met with several researchers during the conference reception at the Marriott Inn. We discussed whether NSF might entertain a proposal to support an email system for collaboration. The group included Lawrence Landweber, Richard Lipton, Richard Demillo, and Edward Robertson. After the meeting, Fredrick Weingarten and I decided to encourage Landweber to submit a proposal (NSF 7801689, An Electronic Mail-Box and Teleconferencing Network for Theoretical Computer Science). Landweber agreed to add Starr Roxanne Hiltz of Uppsala College to analyze the impact on collaboration and research output.

Thirty or so theoretical computer scientists in the United States and Australia participated in the Theorynet project by using Telemail running on a University of Wisconsin computer and accessing it over Telenet,⁶⁵ a commercial packet-switched network. Research collaboration rose steadily and the 1978 ACM SIGACT program committee communicated via Theorynet/Telemail. Although she had some difficulty monitoring usage and interviewing users, Hiltz⁶⁶ was able to show positive outcomes in terms of collaboration and jointly published papers. Theorynet's modest success lent credibility to the CSNET and NSFNET projects.

2.7 Cryptography and Interactions with the National Security Agency

The various controversies that arose around the National Bureau of Standards (NBS) Digital Encryption Standard (DES) were a prologue to issues related to cryptography. IBM submitted a cryptographic algorithm as a candidate for the DES in 1974 and NBS requested that the NSA evaluate it.⁶⁷ NBS also asked IBM to grant the U.S. government “a nonexclusive, royalty-free license to make, use, and sell apparatus that implemented the algorithm.” NBS published a notice in the Federal Register in August 1975 of the proposed standard and requested comments. Martin Hellman

and Whitfield Diffie of Stanford University criticized the proposed DES standard and outlined a potential attack on the algorithm.⁶⁸ In April 1977, NBS issued the DES standard.⁶⁹

In November 1976, in the midst of the DES controversy, Diffie and Hellman published “New Directions in Cryptography,”⁷⁰ which introduced several new concepts: *public key cryptosystems*, *one-way authentication* (or *functions*), *trap-door one-way functions*, and *digital signatures*. At about that time, El (Elias) Schutzman, the program director in engineering systems, approached me about co-funding grants to Hellman at Stanford and I agreed. Diffie was at the time a research assistant working with Hellman. I was also funding Ronald Rivest, who was developing a number-theoretic public-key encryption algorithm⁷¹ with Leonard Adleman and Adi Shamir (all at MIT), which became the RSA algorithm.

What we did *not* know at the time was that James Ellis of the British Communications Electronics Security Group (CESG) had published a classified paper⁷² containing the idea of public key cryptosystems and that Clifford Cocks, also with CESG, had proposed an implementation similar to RSA.⁷³ Both of these British papers predate the Americans’ work by four to five years; however, since they were classified, the NSF-funded researchers would not have known about them before CESG de-classified the work in 1997. The National Security Agency, however, was aware of Ellis’s and Cocks’s work.

In August 1977, J. A. Meyer of Bethesda, Maryland (later identified as an employee of the National Security Agency), wrote to the IEEE suggesting that some attendees at the September 1977 IEEE Symposium on Information Theory held at Cornell might be violating provisions of the International Traffic in Arms Regulations (ITAR) Act. Hellman and Rivest turned the problem over to their universities’ lawyers and opted to wait until the lawyers finished looking into the issue. Cleared to attend, they both limited their public discussion to the mathematical and technical aspects of cryptography and did not discuss possible national security implications of their work.⁷⁴

In April 1978, the NSA placed under a secrecy order a patent application from the Wisconsin Alumni Research Foundation on behalf of George Davida of the University of Wisconsin-Milwaukee.⁷⁵ NSA invoked provisions of the Invention Secrecy Act preventing Davida from discussing any aspect of this research and severely limiting his ability to pursue research in cryptology for several months. The secrecy order was later lifted. That fall, Davida joined NSF, replacing me as program director for Theoretical Computer Science.

The American Council on Education (ACE), in response to a request by the NSA, assembled the Public Cryptography Study Group⁷⁶ in March of 1980. NSA indicated

concern that information contained in some articles in professional journals and in monographs might be a risk to national security. The study group held a series of meetings through February 7, 1981, and produced a report⁷⁷ that recommended a voluntary system of review of papers in cryptology. No author or publisher would be required to participate. Davida contributed a minority report⁷⁸ that argued against restraints on non-governmental research in cryptography.

According to a report in *Science* in August 14, 1980,⁷⁹ Leonard Adleman was told by an NSF program officer that parts of his grant proposal would not be funded. Vice Admiral Bobby Inman, NSA Director, was quoted as saying that the reason the NSF chose not to fund parts of Adelman's proposal is that NSA wanted to fund the research itself. Soon afterward, NSA Director Inman wrote to *Science* indicating that NSA, as the government's primary user of cryptography, was increasingly interested in investing in primary research in cryptography as well as related fields, such as mathematics. He mentioned NSA's assistance with evaluating NSF research proposals in cryptographic areas but stated, "NSA does not now have and does not intend to seek the authority to prohibit NSF funding in this area."⁸⁰ Inman hoped that NSA would become an increasingly important sponsor of research in this area.

In November 1980, NSF Acting Director Don Langenberg clarified the respective roles of NSF and NSA in support of cryptologic research.⁸¹ Since 1977, NSF routinely referred proposals with relevance to cryptology to NSA for review. The process I used⁸² as program director for Theoretical Computer Science, following guidance from the NSF management and attorneys, was to include a designated NSA expert among the referees from whom I solicited proposal reviews.

Langenberg stated that NSF long had a policy of encouraging other agencies to support basic research and had encouraged NSA to establish an unclassified basic research program, but "if an investigator prefers to apply only to NSF, the proposal will be processed in the usual manner, without prejudice." Langenberg added that the Foundation would ensure reporting requirements that would allow it to meet its responsibilities with respect to classification.⁸³ The Adleman proposal was approved by the NSF on December 9, 1980, and the award letter included a statement of NSF policy and elaborated reporting requirements. After negotiation between NSF and MIT, a grant was made to Rivest on September 25, 1981, with an altered policy on reporting.

Jack Minker of the University of Maryland, who was co-chair of the 1980–1981 computer science advisory subcommittee, asked John Guttag of MIT to head an ad hoc committee to review current NSF policy regarding cryptographic research. At the May 1981 Advisory Subcommittee,⁸⁴ a three-and-one-half-hour discussion was held on the "Role of the NSF in Supporting Cryptological Research." Guttag

was asked to prepare a final version of the report, have it approved by the subcommittee chairs (Minker and Thomas Pyke) and the section heads (Curtis and William Rosen), and transmitted to Langenberg. The report urged:

No agency or part of the government should be allowed to bypass the normal means of controlling information by using the National Science Foundation to threaten the funding of those producing the information. Most of the recommendations made in this report have as their implicit goal promoting the clean separation of the procedures for funding and otherwise promoting basic research from the procedures for handling national security and other non-scientific considerations. We believe that the applicability of most of the recommendations contained within this report is by no means limited to the area of cryptology. . . . NSF must continue to support, as Dr. Langenberg put it, “the best research it can find in all areas of science and engineering, with the fewest possible restrictions on investigators.”⁸⁵

Sometime after I had returned to NSF in 1980, John Cherniavsky, the new program director for the Theory program, and I made several trips to the NSA headquarters at Fort Meade to help them design an open and unclassified basic research program. I believe our work with NSA was in the same time period as the delivery of the Gutttag Report. NSA subsequently established an unclassified research grants program, which made its first award in FY 1982. The NSF cooperated in this program and made joint awards with the NSA.

2.8 The Computer Science Section, 1979–1984

I left NSF in the fall of 1978 for a position in the Institute for Computer Science and Technology (ICST) at the National Bureau of Standards (now the National Institute for Standards and Technology). George Davida had replaced me as program director for Theoretical Computer Science and, after a year, he was followed by Meera Blattner. Anil Jain had come in to replace Eamon Barrett in the Intelligent Systems program and later was followed by Y. T. Chien. Before I left, Frederick Weingarten left, eventually joining the Congressional Office of Technology Assessment.

The period from 1977 to 1984 saw many changes in the NSF management and, eventually, growing support, if not budget, for computing research. At the director level, in 1976 Guy Stever became Gerald Ford’s Science Advisor. Richard Atkinson replaced Stever as NSF Director through the early NSA discussions. Both Atkinson and his deputy, Donald Langenberg, were supportive of computer science. John Slaughter’s term as Director (1980–1982) was short, but he recruited Thelma

Estrin⁸⁶ of the UCLA computer science department to head the Electrical, Computer, and Systems Engineering Division. Slaughter was also supportive of the CER and CSNET programs. Ed Knapp arrived from Los Alamos in 1982 and served as Assistant Director for MPS for only two months before being named NSF Director. He was very supportive of computer science, CSNET, and NSF's role in future networking and high-performance computing. When he returned to Los Alamos, Erich Bloch became NSF Director in September 1984. Soon after, Computer Science became a separate division again and, in just two years, part of the new Computer and Information Science and Engineering Directorate.

Within MPS, Ed Creutz retired from his role as Assistant Director for MPS in 1977 and was replaced by Jim Krumhansl from Cornell. Krumhansl was much more supportive of computer science but left in 1979. Bill Klemperer came from Harvard to serve as AD from 1979 to 1981. He was supportive of computing research but skeptical about the section's leadership. When asked to create a separate Division of Computing Research after Pasta's death, he brought Jim Infante from Brown University back in⁸⁷ to head the Mathematical and Computer Sciences Division, delaying the creation of a separate computing research division until 1984. After Klemperer left, he was replaced in MPS by Marcel Bardon on an acting basis. At some point in the fall of 1984, with support from Bardon, Infante, and Bloch, the Mathematical and Computer Sciences Division was split and the Division of Computing Research (DCR) was created.

After the NSF, with the backing of Klemperer, Atkinson, Langenberg, and the National Science Board, decided to allow a new set of programs to address the crisis in experimental computer science research, Kent Curtis contacted me to see if I would be interested in returning to NSF. I was able to retain a visiting research position at NBS while having a chance to make a difference for computer science nationally at NSF. In January 1980 I returned as the program director for Special Projects, which included the Coordinated Experimental Research program and CSNET, as well as the Computer Science Section programs in databases, security and privacy, and social impacts of computing. After almost four years with the Special Projects program, I left in August 1984 for an Independent Research and Development (IR&D) assignment to the University of California, Berkeley.

Before I left for Berkeley, I was assigned part time to the new Office of Advanced Scientific Computing (OASC) to direct the networking programs. While I was at Berkeley, the Division of Computing Research (DCR) was created. Kent Curtis became DCR Division Director and immediately opened a search for a deputy division director (DDD). I applied and was hired into that position beginning in September

1985. My duties were largely administrative as DDD/DCR and were quickly overtaken by my role as a senior scientist for CISE.

In the following sections, I will describe the Coordinated Experimental Research program and its successors, CSNET, the OASC Networking program and the beginning of NSFNET, and the creation of CISE. One goal is to describe the roles of the many people within NSF who helped the computing-related programs begin to grow, thrive, and assume the significant leadership position that CISE holds today.

2.9 Addressing the Need for Academic Experimental Computer Science

The Coordinated Experimental Research (CER) program was created in response to a perceived “crisis” in academic computer research. The NSF heard from the Computer Science advisory committees, the Feldman and Snowbird reports, and Peter Denning’s articles in ACM *Communications* that serious problems were arising in the field. This drumbeat of reports began to have a significant impact on the perception of computer science within the NSF and across the federal government.

In these reports, members of the computing research community pointed to the rapid deterioration of research facilities and the flight of faculty and graduate students to industrial laboratories. Many felt that only three institutions, Carnegie Mellon, MIT, and Stanford, were adequately capitalized to perform experimental research. Only researchers associated at these three universities and, to a lesser extent, other departments and labs with specialized DARPA/IPTO projects, had adequate experimental infrastructure. Even at the major DARPA centers, access was often limited. Remote access to these facilities could be obtained via ARPANET in some cases, but ARPANET access was also limited. As a result, computer scientists at many of the major research universities were engaged primarily in theoretical research and training, graduating fewer Ph.D. computer scientists, and failing to meet the growing demand for experimental computer science faculty.

At the May 1979 Computer Science Advisory Committee,⁸⁸ Jim Krumhansl, AD/MPS, cited the beneficial effect of recent reports. He noted that Frank Press, the President’s Science Advisor, used the Feldman Report as the basis for recent remarks. Krumhansl, however, admitted that computer science would not be a part of any special initiative. The Office of Science and Technology Policy was said to be considering the “general area as one of national concern and in this is dealing with DARPA, OMB, and any other agencies involved.”⁸⁹

The Advisory Committee in May 1979 warned of “the eroding research position of the United States in experimental computer science,” applauding “the recent

action by the National Science Board to place special emphasis on computer science in FY 1981,” and encouraging “the Foundation to continue that initiative throughout the budgeting and appropriation process.”⁹⁰ The Committee placed a high priority on five-year Centers of Excellence Grants and a Computer Network for Research. The Centers of Excellence program could provide up to \$2,000,000 for capital investment plus up to \$500,000/year for operating and maintenance costs and software development. They placed a somewhat lower priority on new investigator and career development awards, graduate fellowships, and traineeships.

In a report⁹¹ issued in December 1979, the Advisory Committee recommended that NSF invest \$15 million each year in a national competition for resource grants. These grants would “total no less than \$250,000 and no more than \$2 million,” include maintenance and software support of 10% of the capital costs and be available to individual researchers and departments. The report expected that after five years, the program would “produce at least 25 well-equipped university laboratories among the 64 computer science Ph.D. degree granting universities.”

John Pasta and Kent Curtis responded to these recommendations by “taxing” the Standard Research Projects Support (SRPS) budget in FY 1980 by \$1 million, almost the entirety of the FY 1980 budget increase, to create the Coordinated Experimental Research (CER) program. It had three main thrusts: a CER facilities program, a program to assist the research community in developing networking services in support of computer science research, and grants to attract experimentalists into a university environment. Curtis sent a “dear colleague letter” in November 1979 inviting proposals for what would become the CER facilities program. Program descriptions for a New Investigator program and a Postdoctoral program came after for FY 1981 funding. In the first year the CER program funded one facilities grant and a CSNET study grant.

Today, electronic “dear colleague letters” quickly gain broad audiences, but in 1979 Curtis mailed a letter to the computer science Ph.D.-granting departments. The timeframe was short and we received only seven proposals. Predictably, most of the proposals came from people familiar with experimental computer science within NSF and the Computer Science Section. I said in a 1990 interview, “in the first set of proposals there was one good proposal, one sort of half-good proposal, and the rest of them were . . . bad proposals [although they involved] some very good people”⁹²—bad in the sense that they seemed to be independent projects “stapled together” rather than a unified coherent proposal.

Our first challenge was to determine how to review the proposals. We decided on a multi-stage process: mail reviews, site visits, and a final panel. We decided that I with two external reviewers would personally visit the sites of all of the proposing

institutions, following receipt of mail reviews. Principal Investigators would have the mail reviews to react to, along with questions by the site visitors. The final panel included no academics but instead the heads of major industry and non-profit laboratories.

The first CER award was made to the University of Washington to construct the Eden operating system with a goal to build a system coupling the performance of powerful personal machines with the resource sharing and accessing capability of a modern time-sharing system. This major research project involved a majority of the departmental faculty and produced a facility that could support a variety of research projects. The Eden Project attracted co-funding from Intel and Digital Equipment, whose technical staff collaborated with Washington on the research.

In FY 1981, CER became an official program with \$3.6 million dollars of the Special Projects budget identified as “experimental computer science” with other line items for CSNET, young investigator, and postdoctoral awards. As we discuss below, a revised set of CSNET proposals were received and approved by the National Science Board. One postdoctoral award and four new investigator awards were made. The CER program received 24 proposals responsive to the program announcement, which were distributed to other NSF, Office of Naval Research (ONR), and DARPA programs. We hoped for significant DoD involvement in developing CER sites as DARPA, ONR, the Air Force Office of Scientific Research (AFOSR), and the Army Research Office were planning a “Computer Resources Initiative” in FY 1982 with \$30 million among the DoD science agencies.

With cooperation from the DoD agencies, we selected 11 proposals for site visits similar to those of the prior year. From the eleven sites visited, five were discussed with DARPA and ONR. Following a budget negotiation, four proposals went to the National Science Board, which approved three immediately and a fourth later. DARPA eventually funded a version of a fifth proposal. The four new NSF CER awards went to Cornell to support investigation into the programming process, Illinois for the construction of computer aids to program and system development, the University of Wisconsin–Madison for construction of a 50-node network of powerful computing devices, and Yale for facilities to support artificial intelligence and natural language processing, numerical computing, and computer architecture.

In the succeeding years when I managed the CER program, five awards were made in FY 1982 to Rice, Brown, Utah, UCLA, and Texas. Four awards were made in FY 1983 to North Carolina Chapel Hill, Pennsylvania, Maryland College Park, and Duke. SUNY Stony Brook, Rochester, Arizona, and New York University received grants in FY 1984. After I left for Berkeley, Harry Hedges joined NSF from Michigan

State to run the CER program. In FY 1985, Hedges and Bruce Barnes made awards to Princeton, UMass Amherst, Colorado Boulder, and Minnesota.

As the CER program began, we did not completely agree on its goals. Some supported the concept of “Centers of Excellence”; some supported funding large, multi-investigator “experimental” research projects; and some promoted large-scale facilities grants, which would include equipment, maintenance, supplies, and technical staff. Clearly the Eden Project fell into the large, multi-investigator “experimental” research project category. While I personally favored a focus on large, collaborative research projects, there were few awards in this category. Reviewers prioritized facilities support and grants to institutions with an existing core of potential experimental computer scientists. Almost all of the grants had a unifying theme, but the available funding limited grants to support for equipment, maintenance, and support staff, with some support for the lead principal investigators, and provided few or no funds for graduate students, postdocs, or faculty salaries.

I attempted to create a CER community based upon the DARPA model. We held a two-day CER principal investigator (PI) meeting⁹³ in February 1984 where the PIs presented their research in a series of focused sessions. Even though large, integrated projects usually were not the primary focus, the new state-of-the-art facilities resulted in many significant research projects. Jack Schwartz’s 1983 taxonomy of parallel computers⁹⁴ included several that were developed or extended under CER grants. These included the NYU Ultracomputer,⁹⁵ the Illinois CEDAR machine,⁹⁶ the Texas Reconfigurable Array Computer (TRAC),⁹⁷ the Berkeley Hypertree (also at Wisconsin),⁹⁸ the Utah Applicative Multi-Processing System,⁹⁹ the Wisconsin GAMMA database machine,¹⁰⁰ the Maryland ZMOB,¹⁰¹ Yale’s ELI-512 computer,¹⁰² the Duke Boolean Vector Machine,¹⁰³ and the Blue CHip Project¹⁰⁴ at Washington (begun as the Purdue Configurable, Highly Parallel (CHiP) family). The Eden Project¹⁰⁵ expanded on ideas from the efforts at Xerox PARC, SRI, and other industry labs, and developed an influential operating system. In a similar direction, the Crystal Project¹⁰⁶ at Wisconsin developed a shared multicomputer. The Cornell CER started a long career by Ken Birman¹⁰⁷ in distributed operating systems, which included the Isis Toolkit, the Horus system, the Ensemble system, and currently Isis2, Gradient, and the reliable TCP solutions. At Cornell, Bob Constable worked with Birman on Horus and Ensemble and developed a program development system called PRL (“pearl”)¹⁰⁸ that provides automated assistance with explaining and proving. There are many additional examples.

Several CER grants became the basis for early Science and Technology Centers relating to computing: the \$38 million Science and Technology Center for Research on Parallel Computation at Rice University with the California Institute of

Technology, Syracuse University, the University of Tennessee, Argonne National Laboratory, and Los Alamos National Laboratory (NSF 9120008); the \$21 million Center for Research in Cognitive Science at the University of Pennsylvania (NSF 8920230); and the \$35 million Science and Technology Research Center in Computer Graphics and Scientific Visualization at the University of Utah with Cornell University, Brown University, the University of North Carolina, and the California Institute of Technology (NSF 89202191). The more recent Team for Research in Ubiquitous Secure Technology (TRUST) at Berkeley, with Carnegie Mellon, Cornell, Mills College, San Jose State, Smith College, Stanford, and Vanderbilt (NSF 0424422) can trace some of its activities back to research that came out of the Cornell CER some 20 years earlier.

There was concern early in the CER program that these investments were severely limiting the funds available for regular grants. Jim Ortega at the May 29, 1981, CS Advisory Committee requested a review of the impacts of CER and CSNET on Standard Research Project Support (SRPS). While the Computer Science Section budget had increased 24% from FY 1980 to FY 1982, SRPS support had only increased 10.4%. In comparison, the Mathematical and Physical Sciences budget increased 19%. After substantial discussion, the advisory committee concluded that “the CER and CSNET are essential to the furtherance of computer science research and that it is too early to modify the direction being taken.”¹⁰⁹

In 1982, the DoD planned to expand its agencies’ support to include as many as 10 or 15 institutions. This DoD program never materialized, but DARPA upgraded facilities for its major contractors and expanded its smaller (\$250–300,000) equipment contracts. ONR was able to provide a few Special Research Opportunities contracts in computer research with some facilities support. Without the planned DoD programs, the CER program grew in an attempt to fill the need.¹¹⁰ Through 1985, NSF had committed \$49.89 million to 22 institutions for experimental computer research. In addition, DARPA had major contracts with MIT, Stanford, Carnegie Mellon, and California-Berkeley, which had supported experimental computer research. When NSF began the CER activity, it expected to support approximately 15 institutions. With more than 70 Ph.D.-granting departments of computer science and engineering, it was estimated that 25 to 30 would require research facilities of the magnitude provided by the CER program.

A report in 1986¹¹¹ noted that “universities have been funding CS growth at rates significantly higher than in any other major discipline. But national funding policy has favored the growth of basic research in CS at a rate no greater than that of other scientific, mathematical, and engineering disciplines.” The authors warned that “the late 1980s will witness the departure of our best and most mobile computer

scientists and graduate students for industrial careers. Inevitably, the universities will be unable to maintain the centers of academic excellence in CS that have been so carefully developed during the past five years.” A year later, “[t]here has been a dramatic increase in federal funding for both total and academic CS research between FY 1976 and FY 1987 . . . Funding has shifted away from basic and toward applied research, both in CS federal funding as a whole and within academic CS.”¹¹² NSF convened an Infrastructure Workshop in July 1991. The workshop report¹¹³ placed a high priority on maintaining the Institutional Infrastructure programs at \$20 million per year. It also proposed developing a matching program of \$8 million to support facilities for individual and small group grants.

Recognizing that no more than 30 computer science departments¹¹⁴ would have enough experimental computer scientists to require CER-scale funding, a new Institutional Infrastructure program was announced with both “Large-Scale” (II-LS) and “Small-Scale” (II-SS) grants. Given the shortcoming of DoD funding, described above, CISE invested in, expanded, and replenished the experimental facilities at around 30 institutions. The II-LS program essentially replaced CER. II-SS was aimed at units with fewer experimental computer scientists and a reduced need for facilities support. Figure 2.3 shows both “large” (CER and II-LS) and “small” (II-SS).

The CISE Institutional Infrastructure program continued until 1993, when it was replaced by the CISE Research Infrastructure (RI) program. The RI program had institutional, instrumentation, and “shared” facilities, such as the CISE Advanced Distributed Resources for Experiments (CADRE). Figure 2.3 shows that 60 institutions benefited from the CER, II-LS, II-SS, and RI—many receiving three or more awards. In 1989, CISE introduced a facilities program directed toward minority-serving institutions (see Figure 2.4). One of those awards to University of Texas at El Paso became the basis for the CISE BPC CAHSI (Computing Alliance of Hispanic-Serving Institutions) Alliance, and in turn the CAHSI INCLUDES project, one of the first five \$10 million NSF INCLUDES Alliances.

When I returned to NSF in 2000 as Division Director for Experimental and Integrative Activities, we moved the Minority Institutional Infrastructure to the education and workforce block of programs and began to redesign the remaining infrastructure programs of CISE. Today, CISE supports a Community Research Infrastructure (CCRI) program to encourage “discovery and learning in the core CISE disciplines . . . by funding the creation and enhancement of world-class research infrastructure. This research infrastructure will specifically support diverse communities of CISE researchers pursuing focused research agendas in computer and information science and engineering.”¹¹⁵

Institution	80	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95	96	97	98	99	
U. Washington (4)	CER	CER	II-LS	II-LS	II-LS	II-LS															
Wisconsin, Madison (4)	CER	CER	II-LS	II-LS	II-LS	II-LS	RI														
Cornell (4)	CER	CER	II-LS	II-LS	II-LS	II-LS	RI														
Yale (2)	CER	CER	II-LS	II-LS	II-LS	II-LS															
Illinois, U-C (5)	CER	CER	II-LS	II-LS	II-SS	II-SS/RI	RI														RI
Rice (3)	CER	CER	II-LS	II-LS	II-LS	II-LS	RI														
Brown (2)	CER	CER	II-LS	II-LS	II-LS	II-LS	RI														
Utah (3)	CER	CER	II-LS	II-LS	II-LS	II-LS	RI														
UCLA (2)	CER	CER	II-LS	II-LS	II-LS	II-LS	R/2														
Texas, Austin (2)	CER	CER	II-LS	II-LS	II-LS	II-LS	RI														
North Carolina, CH (2)	CER	CER	II-LS	II-LS	II-LS	II-LS	RI														
Pennsylvania (3)	CER	CER	II-LS	II-LS	II-LS	II-LS	RI														
Duke (2)	CER	CER	II-LS	II-LS	II-LS	II-LS	RI														
Maryland, CP (2)	CER	CER	II-LS	II-LS	II-LS	II-LS	RI														
SUNY, Stony Brook (3)	CER	CER	II-LS	II-LS	II-LS	II-LS	RI														
Rochester (3)	CER	CER	II-LS	II-LS	II-LS	II-LS	RI														
Arizona (3)	CER	CER	II-LS	II-LS	II-LS	II-LS	RI														
NYU	CER	CER	II-LS	II-LS	II-LS	II-LS	RI														
Princeton (4)	CER	CER/II-LS	II-LS	II-LS	II-LS	II-SS	RI														
UMass, Amherst (4)	CER	CER	II-LS	II-LS	II-LS	II-LS	RI														
Colorado, Boulder (3)	CER	CER	II-LS	II-LS	II-LS	II-LS	RI														
Minnesota (2)	CER	CER	II-LS	II-LS	II-LS	II-LS	RI														
Indiana (2)	II-LS	II-LS	II-LS	II-LS	II-LS	II-LS	RI														
UC Irvine	II-LS	II-LS	II-LS	II-LS	II-LS	II-LS	RI														
Purdue (3)	II-LS	II-LS	II-LS	II-LS	II-SS & LS	II-SS & LS	RI														
Georgia Tech (2)	II-LS	II-LS	II-LS	II-LS	II-LS	II-LS	RI														
UC Berkeley (3)	II-LS	II-LS	II-LS	II-LS	II-LS	II-LS	RI														
Chicago	II-LS	II-LS	II-LS	II-LS	II-LS	II-LS	RI														
Virginia	II-LS	II-LS	II-LS	II-LS	II-LS	II-LS	RI														
Columbia (2)	II-LS	II-LS	II-LS	II-LS	II-LS	II-LS	RI														

Institution	80	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95	96	97	98	99	
Michigan State													II-SS								
RPI													II-SS								
New Mexico State													II-SS								
Penn State													II-SS								
Johns Hopkins (2)													II-SS								RI
Northeastern													II-SS								
Florida													II-SS								
Tennessee													II-SS								
Texas A&M													II-SS								
UC Santa Cruz													II-SS								
Oregon State													II-SS								
UC Santa Barbara													II-SS								
USC													II-SS								
Brandeis													RI								
Illinois, Chicago (2)													RI								RI
Virginia Tech													RI								
Harvard													RI								
Kansas													RI								
Washington U.													RI								
Kentucky													RI								
New Mexico													RI								
Stanford													RI								
Boston													RI								
Delaware													RI								
North Carolina State													RI								
Northwestern													RI								
Oregon Graduate Inst.													RI								
MIT													RI								
Dartmouth													RI								
UC San Diego													RI								

Figure 2.3 Infrastructure awards FY 1980–1998.

Institution	89	90	91	92	93	94	95	96	97	98
North Carolina A&T		II-MI								
Puerto Rico, Mayaguez		II-MI				II-MI				
Maryland, Eastern Shore		II-MI								
Spelman		II-MI								
Texas, El Paso (2)		II-MI				II-MI				
Bowie State		II-MI								
CCNY		II-MI								
Clark Atlanta (2)		II-MI					II-MI		▶	
Florida A&M			II-MI							
Florida International			II-MI/3		MI	II-MI			▶	
Hampton		II-MI								
UDC		II-MI								
Fond du Lac			II-MI							
Puerto Rico, Rio Piedras			II-MI			II-MI				
Houston, Downtown						II-MI			▶	
Xavier						II-MI			▶	
Texas, San Antonio						II-MI			▶	
Tuskegee						II-MI			▶	

Figure 2.4 Minority institutional infrastructure awards FY 1989–1998.

2.10 CSNET

In parallel with the “crisis” in experimental computer science, a number of researchers at leading universities did not have access to the ARPANET.¹¹⁶ Curtis and Pasta’s strategy for the Coordinated Experimental Research program included a computer network for research. Lawrence Landweber invited a number of researchers and government representatives to the University of Wisconsin–Madison in May 1979. His goal was to “discuss how computer network services like those of ARPANET could become available to the entire community of computer science researchers.”¹¹⁷ The attendees included Kent Curtis, Bob Kahn, and individuals who had experience with Theorynet and other similar “mailbox” systems hosted on commercial networks.¹¹⁸ The participants agreed that ARPANET’s mail, file transfer, and remote login services had “enhanced research productivity and had generated a strong community spirit among computer science and engineering departments that hosted ARPANET sites.”¹¹⁹

A consortium of universities including Georgia Tech, Minnesota, New Mexico, Oklahoma, Purdue, UC-Berkeley, Utah, Virginia, Washington, Wisconsin, and Yale submitted a proposal in November 1979 for a “CSNET” that would create a separate and independent network to provide ARPANET-like services to all U.S. computer

science departments. Given the cost of duplicating the ARPANET infrastructure (estimated at \$100,000 per institution), the proposed network would be built on commercial X.25 networks such as CompuServe, Tymnet, and Telenet. NSF declined the proposal in March 1980. Reviewers felt that the proposers were reinventing the ARPANET and not extending it, that they lacked a strong project management plan, and that for NSF to pay for the network it would have to reduce research support.¹²⁰ The reviewers' skepticism was not unlike the reaction I had heard during CER site visits in 1980–1981 when asking proposing PIs about the CSNET plans. Many of those outside the CSNET proposal development did not see a real justification for an ARPA-like network, and some not even the need for email.

The NSF offered to fund a thorough study of CSNET. Landweber organized a meeting in Berkeley on June 15, 1980, at which DARPA announced its support for CSNET and assigned Vinton Cerf to help develop a plan to connect CSNET and the ARPANET. Landweber convened a 19-person CSNET planning committee, including Cerf and others who had extensive computer networking experience. The group worked throughout the summer of 1980 to devise an implementation strategy. The outcome was a plan to design CSNET as a network on multiple communication platforms interconnected via an Internet protocol. DARPA was moving from NCP to TCP/IP and the MMDF-based¹²¹ Phonet system had been developed by David Farber and David Crocker at the University of Delaware. Phonet was a low-cost mail relay system similar to the UUCP-based mail relay developed by Bell Laboratories to connect computer science departments that had Unix platforms. The UUCP protocol¹²² supported email and file transfer, but required explicit addressing and, unlike MMDF, was not compatible with IP networks at the time. The proposal would integrate ARPANET access, X.25 networks running TCP, and Phonet to provide multiple tiers of services and costs for departments wishing to be connected to CSNET.

Landweber and colleagues from the University of Delaware, Purdue University, RAND Corporation, and the University of Wisconsin submitted a revised CSNET proposal to NSF in October 1980, and the National Science Board (NSB) approved the five-year proposal the following January. To address concerns about how CSNET would be managed required an unusual structure in which NSF itself, under Project Director C. William Kern, would directly manage the project for two years (through 1983) by means of contracts. NSF management would focus on setting up the organization to collect and disburse funds, and after two years the project would be sufficiently advanced that users would be willing to begin paying dues and fees.¹²³ Contracts were established with the University of Delaware, Purdue University, Rand Corporation, and the University of Wisconsin for CSNET development. Bolt, Beranek, and Newman (BBN) was contracted to run the CSNET Coordina-

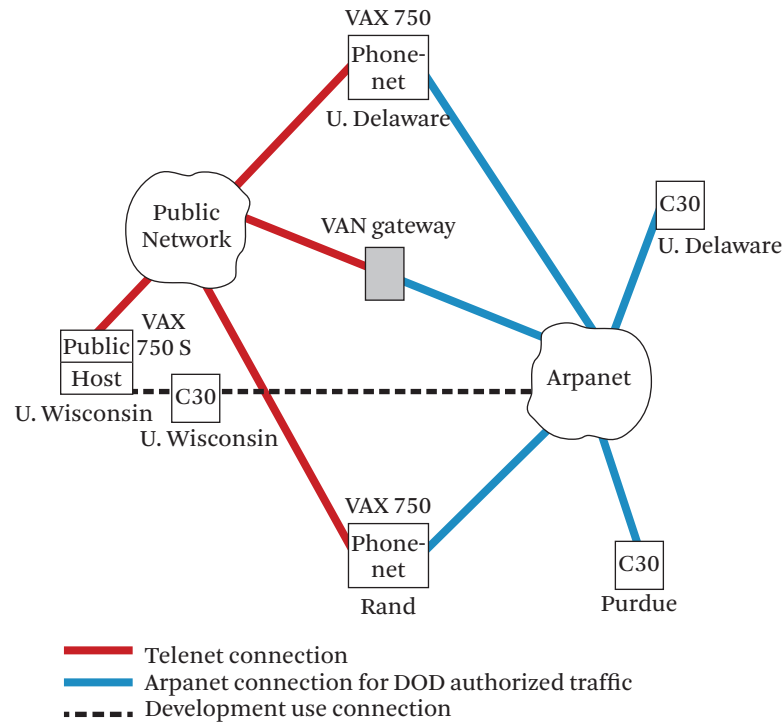


Figure 2.5 CSNET architecture 1981.

tion and Information Center (CIC) for managing the network and distributing software.

On March 6, 1981, NSF announced the establishment of CSNET, which would become a major step along the path to the Internet. On May 28, 1981, Bill Kern presented the status of the CSNET effort to the Computer Science Advisory Committee.¹²⁴ He discussed the two-year NSF management plan and the expectation that CSNET would become self-supporting in five years. He also told the Advisory Committee members that DARPA would develop the CSNET/ARPANET gateway and that software, systems, and services would target the Berkeley UNIX 4.3BSD operating system on VAX computers. He indicated that CSNET would initially comprise three subnets (Figure 2.5)—ARPANET, Telenet, and Phonet—but would be designed to support expansion to other available networks. CSNET initially provided the same services as ARPANET: mail, file transfer, remote login, and an on-line name server. CSNET's \$5 million project budget, limited staffing, and the five-year timeframe for self-sufficiency put significant pressure on the CSNET team.

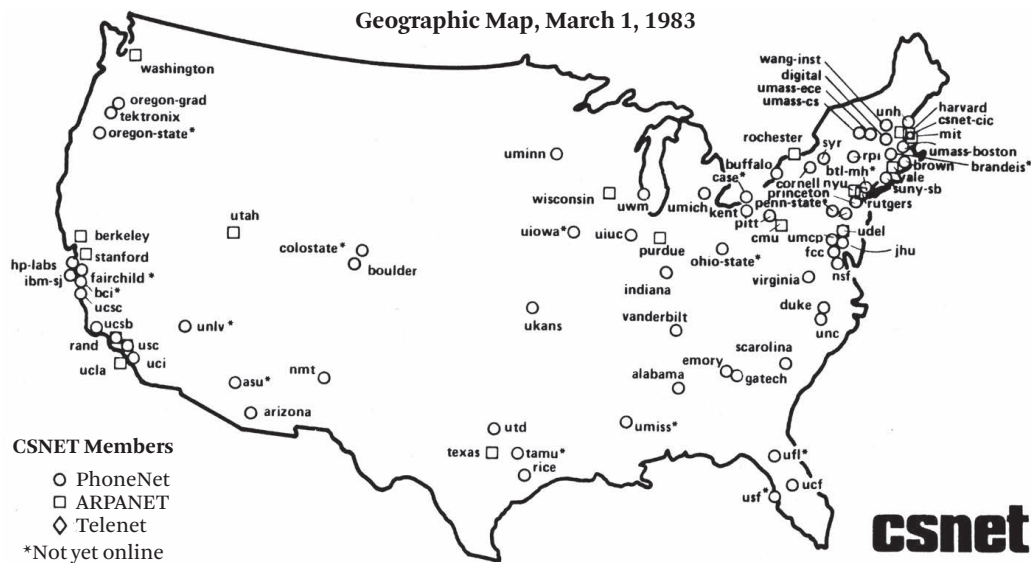


Figure 2.6 CSNET map 1983.

In just six months, CSNET was operating,¹²⁵ including a Phonenet site at NSF in the Computer Science Section, the first NSF Internet connection. In addition to NSF, Phonenet sites included Cornell, FCC-NET, HP Labs, Purdue, Princeton, UC Irvine, and Delaware, with plans to expand to New Mexico Tech, Pennsylvania, Georgia Tech, Duke/UNC, Fairchild, and Maryland-College Park.

When Kern stepped down as CSNET Project Manager in October 1982, I assumed the role of CSNET Project Director with Landweber as Chair; Peter Denning, Richard Edmiston, David Farber, Anthony Hearn, Kern, and me as members of the Management Committee. By the time of the first CSNET Newsletter,¹²⁶ 56 Phonenet sites were operational and 27 were nearing operation (see Figure 2.6). These connected through the two CSNET relays on the ARPANET at RAND Corporation and the University of Delaware. CSNET was beginning to meet with European network leaders to investigate international connections.

After its two-year management, the NSF selected the University Consortium for Atmospheric Research (UCAR)¹²⁷ on May 3, 1983, to host and manage CSNET, with Leonard Romney (UCAR executive director) as PI and a member of the CSNET Management Committee.¹²⁸ As UCAR assumed control of CSNET, the Management Committee was replaced by a larger Executive Committee¹²⁹ with Peter Denning as chair, representing the computing research community; and the operation of the CSNET relays and technical services moved entirely to BBN. BBN housed the CSNET Coordination and Information Center (CIC) to provide operational management



Figure 2.7 CSNET executive committee 1983.

of CSNET. In 1983, CIC staff included: Dr. Richard Edmiston (CIC Director); Laura Breeden (CIC User Liaison); Dan Long (CIC Technical Liaison); and Beth Johnson (CIC Staff Assistant). Leonard Romney left UCAR in May 1984 and was replaced by Stanley Ruttenberg.

By October 1983, Lawrence Landweber was leading an effort to create gateways and connections among BITNET, and Canadian and European networks, including SERCNET (United Kingdom), SUNNET (Sweden), CERNEY (Switzerland), and UNINET (Norway). CSNET connected to BITNET through a University of Wisconsin gateway. At the time, BITNET was a fast-growing network connecting university computing centers via IBM store-and-forward software and leased lines. Connecting to international and other U.S.-based networks raised issues about how to manage the costs associated with traffic transiting multiple networks. In 1983, the initial agreement called for each network to bear the costs of message traffic into other networks. Security issues also arose concerning international traffic in and out of ARPANET via BITNET and CSNET.

In June 1984, I described new NSF networking plans (see NSFNET below) to the CSNET Executive Committee and asked them how CSNET might interact with this expanded vision. CSNET established new gateways with SUNET (Sweden), the Israeli Network, and DFN (Germany). NSF paid for CSNET dues for undergraduate institutions. Dennis Jennings, then chairman of the European Academic Research Network (EARN), visited the CSNET Executive Committee in September 1984. He would soon be recruited to NSF, replacing me as the program director for Networking in the Office of Advanced Scientific Computing (OASC).

The last NSF payment for CSNET operations was in mid-1985. By 1986, CSNET connected more than 165 university, industrial, and government computer

research groups serving more than 50,000 researchers and students, including accounts for 1000 Internet hosts. Network services were operational and numerous networks outside the U.S. were connected.¹³⁰ CSNET was self-supporting and received significant industry funding. CSNET clearly demonstrated, for the first time, that users were willing to pay for network services.

CSNET actively collaborated with colleagues in other countries, supporting and often enabling the international expansion of the Internet. CSNET had mail connection via CSNET/Internet and USENET/EUNET/UUCPNet connections to foreign affiliates and their gateways. These included: CDNNET (Canadian Academic Network, via the University of British Columbia); SDN (System Development Network, with a gateway at the Korea Advanced Institute of Science and Technology); SUNET (Swedish University Network, via Chambers University of Technology); CHUNET (Swiss University Network, via ETH-Zentrum); INRIA (French University Network, through INRIA/Rocquencourt); DFN (Deutsches Forschungsnetz); JUNET (Japanese University Network, through the University of Tokyo); Finnish University Network (via Helsinki University); AC.UK (Academic Community, United Kingdom, via University College, London); ACSNET (via a UUCP-based connection at the University of Melbourne); New Zealand Academic Network (via Waikato University, Hamilton); and the Israeli Academic Network (via Hebrew University of Jerusalem).

At its meeting in Ann Arbor in June 1988, the CSNET Executive Committee discussed a potential merger of CSNET¹³¹ and BITNET. As vice chair of the Executive Committee, I was assigned to the CSNET-BITNET merger team, planning a merged network called “ONENET.”¹³² Eventually, in 1989, CSNET and BITNET were brought under the Corporation for Research and Educational Networking (CREN), a non-profit corporation initially composed of the organizations that had participated in BITNET and CSNET. NSF funded the expansion of CSNET and BITNET, as well as the development of TCP/IP services as adjuncts to NSFNET. Because of the success of NSFNET and the regionals, CREN discontinued CSNET services in 1991. CREN ended their support for BITNET in 1996, due to the growth of TCP/IP-based networks, and by 2003, CREN dissolved itself.

2.11 The Office of Advanced Scientific Computing and NSFNET

Beyond a brief overview of the high-performance computing programs, the details of which are covered in Chapter 10, this subsection examines the developments that led to the NSFNET. Chapter 9 provides details on NSF’s broader role in networking before, during, and after the NSFNET project.

The Lax Report¹³³ identified two problems: “important segments of the research and defense communities lack effective access to supercomputers and students are

neither familiar with their special capabilities nor trained in their use”; and “the capacity of today’s supercomputers is several orders of magnitude too small for problems of current urgency in science, engineering and technology.” The panel recommended a program that would increase access via high bandwidth networks; increase research on computation, software, and algorithms; train personnel; and increase R&D on new supercomputer systems.

In response to the Lax Report, NSF organized an internal working group¹³⁴ in April 1983 to help the Foundation meet the computing needs of academic science and engineering. NSF also held a workshop in May 1983 with 13 scientists from diverse disciplines to define an initiative in large-scale computing and networking. According to what became known as the Bardon-Curtis Report,¹³⁵ NSF should: (1) coordinate with other federal agencies; (2) increase support for local computing facilities; (3) encourage proposals to provide supercomputer services and access and be prepared to support 10 supercomputer systems within three years; (4) support networks linking laboratory researchers with each other and with supercomputer centers to provide access, file transfer, and scientific communication; (5) support academic research in advanced computer systems design and computational mathematics; and (6) establish an NSF advisory committee for supercomputing.

In November 1985, the House Committee on Science and Technology held hearings¹³⁶ on supercomputer and network resources for science research. During the hearings, NSF Director Edward Knapp cited the Bardon-Curtis Report and indicated that Edward Hayes, the NSF Controller, was chairing an NSF Task Force on Advanced Scientific Computing. Knapp also indicated that NSF was gathering information from grantees about their immediate needs for access to Class IV¹³⁷ computers and would negotiate with suppliers who could provide appropriate blocks of time. Under this plan, NSF would continue to support research in the theoretical and experimental design of computers as well as on computational mathematics, software, and algorithms. NSF indicated that its networking initiative would be part of Advanced Scientific Computing. Subsequently, the 98th Congress voted \$40 million to fund the recommendations of the Bardon-Curtis Report.

NSF established an Advisory Committee on Supercomputer Access¹³⁸ chaired by Neal Lane, then Chancellor of the University of Colorado at Colorado Springs, that would become the Advisory Committee for the Office of Advanced Scientific Computing. I staffed a subcommittee¹³⁹ on networking options, chaired by Joe Wyatt, the Chancellor at Vanderbilt. At that time, NSF expected that a supercomputing network would be developed in two phases: Phase I using conventional network technology and expanding existing viable networks, and a Phase II using satellite transponder facilities and optical fiber trunks as they became available.

In December 1983, Landweber wrote to Edward Knapp encouraging him to “proceed as quickly as possible to establish a national Science Net [and] to use existing technologies . . . such as ARPANET, CSNET, and BITNET.”¹⁴⁰ A few weeks later, Jack Schwartz (NYU) wrote¹⁴¹ to Edward Hayes asking him to have the advisory committee consider other needs for high-bandwidth communication beyond access to the supercomputing centers. Even before the Office for Advanced Scientific Computing was established, the community, in particular Landweber and Kenneth Wilson, were pushing for a national network.

As the CER and CSNET programs grew, Kent Curtis and I had been discussing ARPANET opportunities with DARPA. In mid 1983, Curtis asked Frank Kuo of SRI for advice on expanding or duplicating ARPANET technology to support supercomputer access¹⁴² in a network called “USERNET” that might support 200 academic research institutions or 2000 college and university sites. Kuo pointed out that splitting off MILNET from ARPANET would leave only a 40-node network intended to be a “research and development” testbed. He estimated that developing a 200-site USERNET using ARPANET technology would require \$7.5–11.5 million for IMP¹⁴³ hardware and \$17 million in operational costs and 10 times that much for 2000 institutions. He also raised the issue of NSF competition with commercial packet networks such as TELENET, TYMNET, UNINET, or NET1000, and suggested that NSF look instead into using commercial networks for a backbone. My thought at the time was that no commercial network supported full network services and there might be alternative “tiered” approaches.

In April 1984, Kent Curtis and I met with DARPA’s Bob Kahn, who said that the ARPANET, as an R&D network, could only expand by an additional 20 nodes and for \$1 million in capital costs and \$4.8 million in annual costs. At that time, the Division of Computing Research was already supporting ARPANET sites at RAND and Delaware for CSNET and at a few CER sites.

NSF created the Office of Advanced Scientific Computing in May 1984, with John Connolly from the Division of Materials Research as Director, Larry Lee from Mathematics as Program Director for Centers, and me (on loan from the Division of Computing Research) as Program Director for Networking. The first awards for time on Class IV supercomputers totaling almost \$19 million were made July 1, 1984, to Purdue University, University of Minnesota, and Boeing Computer Services. In 1985, OASC expanded access to existing supercomputer resources, adding centers at AT&T Bell Labs, Colorado State University, and Digital Productions (an early computer animation company).

An NSF OASC review panel¹⁴⁴ met in February 1985 to consider 22 applicants. On February 25, 1985, NSF announced¹⁴⁵ funding for four National Advanced Scientific

Computing Centers: the John von Neumann Center (JVNC) at Princeton University, the San Diego Supercomputer Center (SDSC) at the University of California, San Diego, (managed initially by General Atomic), the National Center for Supercomputing Applications (NCSA) at the University of Illinois, and the Cornell Theory Center. Later, NSF named the Pittsburgh Supercomputing Center (PSC) as a fifth center. (The National Center for Atmospheric Research (NCAR) was sometimes considered a sixth center, but was always dedicated to climate researchers and never a part of the program.) Each of these centers was associated with academic and industry partners and had developed a tentative “customer” base of scientists needing access to high performance computing.

With responsibility for the centers and a network under OASC, the OASC Networking Advisory Committee recommended the establishment of a “Sciencenet Phase 1”¹⁴⁶ using available and proven technology to implement a network as soon as possible. The preferred strategy was to expand and interconnect existing networks such as ARPANET and BITNET with selective use of commercial network services. By 1985, the Defense Communications Agency had begun to use ARPANET as an operational DoD network following the cancellation in 1983 of the new command, control, communications, and intelligence (C3I) network, AUTODIN II. DoD was looking into splitting off the “research” sites and using the ARPANET (as MILNET) only for military purposes. This action complicated any approach to leveraging ARPANET for supercomputer access and eventually accelerated the growth of NSFNET.

The planned Sciencenet Phase 1 effort involved the development of Internet protocols, access protocols, and a management strategy for the network. David Farber and Landweber defined a Phase I strategy¹⁴⁷ to quickly enable users of existing networks (ARPANET, BITNET, CSNET, MAILNET, MFENET) to run jobs on supercomputers at the national centers. ARPANET, CSNET/X.25, and MFENET users could remotely log in to supercomputers and run interactive or batch services. BITNET, CSNET/Phonenet, and MAILNET users would have to depend on electronic mail or file transfer for batch submissions. The Landweber-Farber Report also recommended that NSF should (1) add a Sciencenet manager and management team (or contract for such services); (2) establish a working group representing the centers, networks, and NSF management; and (3) establish a permanent Technical Advisory Committee (TAC). Because someone had trademarked “Sciencenet,” the network quickly became known as NSFNET.

After I left for my IR&D assignment to Berkeley in August 1984, I was still involved remotely with both CSNET and NSFNET. NSF was looking for a replacement in OASC and concurrently UCAR was looking for a permanent CSNET Executive

Director. Landweber had met Dennis Jennings, the Director of the BITNET-based European Academic Research Network (EARN) and Computing Center Director at University College Dublin. He encouraged him to apply for the CSNET directorship. Jennings visited NSF and spoke with John Connolly in August 1984. Offered both the CSNET and NSFNET positions, Jennings accepted the NSFNET directorship and began in January 1985. As he recalled, “So when I arrived at the NSF on January 2nd, 1985, the key components were in place: The demand from key researchers; a significant budget for networking—roughly 10% of the supercomputer program budget was devoted to the network; and the CSNET experience that provided the confidence in the internetworking concept and technology. What was required was a Catalyst—and that was my role.”¹⁴⁸

Jennings identified several key decisions made under his leadership.¹⁴⁹ The first was to develop a general-purpose network for all science and engineering research rather than a network only providing supercomputer access. There was considerable disagreement on this issue between the OASC networking subcommittee and John Connolly—and to some extent the centers. Connolly was reluctant to separate the network development from the centers, but Gordon Bell eventually split the networking program off as a separate division in the Computer and Information Science and Engineering (CISE) Directorate, as described below.

Another important decision was to adopt a “network of networks” approach. CSNET employed a network of networks approach by integrating Phonenet dial-up services, public network X.25 services, and ARPANET; and ARPANET, as it transitioned to an R&D network, was also integrating networks with quite different communications layers: satellite, phone lines, etc. After a visit to the Cornell Theory Center, Jennings met with Richard Mandelbaum, who was then working with Cornell and other New York state universities, corporations, and research laboratories to develop a statewide network, NYSERnet. Jennings provided some seed funding to NYSERnet, and later to other regional networks including SURAnet, BARRnet, MIDnet, Westnet, Merit, NorthWestNet, and NEARnet. The network of networks model evolved from supporting networks¹⁵⁰ with differing transport and physical layers and a common Internet layer (such as in CSNET) to also support “tiered” networks that included campus Local Area Networks (LANs), regional networks, and a national backbone. Similarly, the NSFNET program funded the center-based SDSC and the JVNC networks.

As an interim arrangement in October 1985, NSF and DARPA¹⁵¹ agreed that ARPANET could be used to access the centers hosted on ARPANET (Illinois, Cornell, Minnesota, and Purdue). This agreement opened up ARPANET hosts, typically

servers in computer science and engineering departments, to a broader set of users via campus-wide networks. The NSFNET program also funded CSNET and BITNET to develop advanced TCP/IP services to provide similar access.

In September 1985, NSF announced its intention to implement a national backbone linking the five NSF supercomputer centers and NCAR, with connections to regional and campus networks. There was some initial pushback from the centers, concerned that they could lose customers in moving from “star networks” and proprietary protocols to a broadly accessible national network with common protocols. A related decision was the selection of Dave Mill’s “fuzz-ball” PDP-11-based routers due to the high cost of ARPANET Interface Message Processors and the lack of commercial alternatives.

The decision for which Jennings may be best known is the adoption of the DoD TCP/IP and related ARPANET protocols as the standard for NSFNET. NSF had originally intended to use the International Standards Organization (ISO) Open Systems Interconnection (OSI) protocols, but they were not yet widely available.¹⁵² The scientific communities planning to use the centers had developed preferences for protocols used by specific disciplines: MFENET by the magnetic fusion energy community, DECNET by the high energy physics community, etc. TCP/IP was mostly available on Unix systems and not on the Cray Time-Sharing System (CTSS) that was running at many of the centers.

Jennings left NSF at the end of March 1986, having developed a model for NSFNET and moving it forward. He had established a Networking Technology Advisory Group (NTAG) and put a staff in place. Following a brief stint as acting president at the John von Neumann Center, he returned to Ireland and was replaced at NSF by Steven Wolff. Wolff had met Dave Farber when DARPA was arranging an ARPANET connection for the Delaware CSNET Relay. Wolff was working at the Army Ballistic Research Labs (BRL) located in the Aberdeen Proving Ground in Maryland, and BRL had provided the ARPANET line connecting the University of Delaware for the CSNET relay. Farber convinced Wolff to join NSF on a detail from BRL as NSFNET program director. Wolff brought substantial experience to the position, having served on the faculty of Johns Hopkins after receiving a Ph.D. from Princeton. His experience at BRL had included work on TCP/IP. He became Division Director for the Networking and Communications Research and Infrastructure when Gordon Bell split it off from the Division of Advanced Scientific Computing in April 1985—April Fool’s Day as Wolff¹⁵³ remembers it. He was responsible for much of the development, expansion, and eventual privatization of NSFNET. Details are in Chapter 9.

2.12 The Beginning of CISE

In 1986, the NSF programs and offices supporting computing and information research and applications were brought together for the first time since NSF was founded. The new Directorate for Computer and Information Science and Engineering (CISE) would become the organizational core NSF used to exert federal leadership in computing.

At the request of Richard Nicolson, AD/MPS, in April 1985 Kent Curtis carried out an analysis of the options for organizing computing programs within NSF. Curtis considered the NSF Office of Information Services (OIS), the IT support organization led by Connie McLindon. He concluded that, while OIS was funding some projects such as EXPRES and working with the networking program, it “had no primary research role” and should remain an administrative unit. In his analysis, he looked at programs funding “informatics” viewed broadly. These included the Computer Engineering program, the Division of Computer Research, the Division of Information Science and Technology, the Office of Advanced Scientific Computing, and various elements of Materials Research, Mathematical Sciences, Electrical and Computer Engineering and Behavioral and Neural Sciences. Curtis considered various combinations of these programs and even a new directorate encompassing Mathematical Sciences, Cognitive Science, Linguistics, Systems Engineering, and Management Science. There were “substantial benefits and faults to be expected from any decision” he noted, and added that “the Director should feel free to follow his instincts because there is no obvious wisdom to suggest a particular course.”¹⁵⁴

In the late fall of 1985 after I returned to Washington, I began to work with Chuck Brownstein on Bloch’s plans to consolidate NSF computing activities. Bloch officially announced his intentions to create a new directorate and hire Gordon Bell to run it on March 3, 1986. Bell had already begun consulting with Bloch and Mary Clutter. In February, Bell requested that Brownstein take on the role of Executive Officer of the new directorate, Jerry Daen be added as Planning and Administrative Officer, and I join half-time on loan from DCR.

Albert Bridgewater, the Deputy Assistant Director for Astronomical, Atmospheric, Earth, and Ocean Sciences (AAEO), wrote a memorandum¹⁵⁵ to Bell in February 1986 suggesting a process and schedule for a new Computer and Information Science and Engineering (CISE) Directorate. This process included meetings with the National Science Board (NSB), developing long-range plans, and presentations to Bloch and the NSB in the spring of 1986. Bell, Brownstein, Daen, and I began to develop a strategy to address the deliverables outlined in the Bridgewater memorandum.

In a second memorandum,¹⁵⁶ Bridgewater encouraged Bell to be a proactive Assistant Director Designate by: Identifying areas needing greater emphasis or support; organizing community support; organizing National Academy studies; encouraging links and cooperation with other agencies and directorates; and keeping the National Science Board, the Office of Management and Budget, the Office of Science and Technology Policy, and NSF management informed. Bridgewater was really telling Bell that, with Bloch's backing and his national credentials, he had an opportunity that had not been given to the NSF computing program leaders in the past.

In Bell's typed and hand-annotated notes of February 26, 1986, he began to sketch out ideas for CISE:

CISE encompasses fields that are predominately concerned (measured either in design effort or system cost) with the understanding (computer science) and design of computers (computer engineering). These SYSTEMS include: traditional and specialized computers, all forms of computer and communications networks, various transducer interfaces for computers and robots, specialized signal processors, and VLSI circuits and their design systems to implement the particular information processing system.

CISE is not concerned with the phenomena or processes necessary to implement the above systems . . . although it is concerned with the design and implementation of the large systems that integrate and carry out complex, manufacturing processes.

CISE supplies supercomputer resources and network access to programs in all directorates. CISE will initiate programs to facilitate more effective use of supercomputers, including: understanding vector multiprocessors, improved algorithms, software development faster networks and high speed graphics workstations for more effective and enhanced use.¹⁵⁷

Bell decided to argue that CISE encompasses all fields in which the major fraction of the intellectual discipline is computer science or engineering (e.g., robotics, VLSI, signal processing), and that other disciplines would have a "non-trivial" portion of their budget devoted to computing as an "experimental apparatus" and would be responsible for their own applications and for utilizing the computer as a simple component. CISE would "provide the scientific and engineering knowledge for these fields." Bell wondered if educational activities, including supercomputing training and multidisciplinary projects, should be included.

In a memorandum dated February 27, 1986, to Bloch, Clutter, and Engineering AD Nam Suh, Bell proposed that CISE should include the Computer Research

Division (from Mathematical and Physical Sciences), the Information Sciences and Technology Division (from Biological and Behavioral Sciences), the Office of Advanced Scientific Computing, including NSFNET, programs in real time computing applied to signals and communications systems, image understanding, and systems theory (from the Engineering Science in Electrical Communications and Systems Engineering Division), Computer Engineering and Manufacturing Engineering for Computers and Semiconductors (from the Design, Manufacturing and Computer Engineering Division), the Columbia University Engineering Research Center, the Advanced Technology program (from Science and Engineering Education), and the EXPRES Project.¹⁵⁸

Alarmed by Bell’s wide-ranging vision, Nam Suh responded in a memorandum to Bell (copied to Bloch and Clutter) dated February 28, that “the only thing that really deals with the essence of Computer and Information Sciences and Engineering that you ought to take into your new directorate is Computer Engineering. In the rest of the programs, the computer is a peripheral tool, but not the intellectual driving force behind them. You will find that this view is widely supported in engineering schools throughout the country.” Suh added, “Sometimes we have the feeling that this world evolves around computers [but the] role of computers in our society has got to be looked at in the proper context.”¹⁵⁹

Bloch issued a memorandum,¹⁶⁰ dated March 3, 1986, to all NSF staff indicating that he was officially appointing Bell as a consultant to assist him in reorganizing the NSF computing activities with the intention of naming him AD/CISE. Bloch intended to “consolidate into a new directorate several computer-related divisions and programs [including] the Division of Computer Research (MPS); the Division of Information Science and Technology (BBS); the Office of Advanced Scientific Computing (O/D); and certain engineering programs from ENG.” In an attachment, Bloch stated the following rationale—that creating CISE:

- (1) Brings together ongoing activities now spread among several NSF units;
- (2) Simplifies formulation and coordination of new policy directions;
- (3) Makes it possible to deal easily with full span of functions, from basic research through systems engineering in an area critical to national well-being;
- (4) Facilitates internal management; takes program activities out of Director’s office and puts them into a technical area; and
- (5) Will [create a] small disciplinary research directorate—in the range of \$110–130 million, [with an] approximately 50 person staff drawn largely from other parts of NSF.

The allocations of budget and personnel attached to Bloch’s memorandum are shown in Table 2.1. The budget increase from the FY 1986 current plan to the

Table 2.1 Bloch's initial allocation to CISE (in \$ millions)

Organization	FY 1985 Actual	FY 1986 Current Plan	FY 1987 Estimate	FY 1986 Staffing ^a On-Board	IPA
Division of Computer Research	\$39.13	\$38.22	\$44.44	16	2
Division of Information, Science, and Technology	\$8.95	\$8.81	\$11.91	7	1
Office of Advanced Scientific Computing	\$41.40	\$43.28	\$53.63	14	1
Engineering ^b	\$21.22	\$21.63	\$23.85	12	3
Total	\$110.69	\$111.94	\$133.83	49	7

a. Staffing did not include the approximate 6 for the AD Office.

b. Portions of DMCE and ECSE yet to be determined, with the transfer amounts estimated in the \$10–30 million range. Amounts shown are mid-range estimates.

FY 1987 Estimate is \$21.89 million (19.5%), but with almost half of the increase (\$10.35 million) going to OASC.

During March 1986, many people in the Engineering Directorate became alarmed about the potential scope and definition of CISE. Nam Suh was not happy with Bloch's initial decision and mobilized¹⁶¹ members of his Advisory Committee (Frederick Garry, Sheila Widnall, Lester A. Gerhardt, Paul C. Jennings, and Herbert H. Richardson) at NSF on March 18, 1986. Meeting attendees also included Suh's Task Group on Computing (Herbert Voelcker chair, with program directors Alan de Pennington, John Mayer, Howard Moraff, Michael Gaus, Michael Polis, Elias Schutzman, and Donald Silversmith) and Frank C. Huband, Division Director of Electrical, Communications, and Systems Engineering (ECSE). Voelcker was also the Deputy Division Director for Design, Manufacturing, and Computer Engineering (DMCE). ECSE and DMCE were the Engineering divisions most likely to be impacted.

The Engineering Advisory Committee had access to Voelcker's Task Group report,¹⁶² which considered a "broad" and a "narrow" construct for CISE; but the draft report failed to get the full support of the Task Group. The Task Group also analyzed the DMCE Computer Engineering program in a report¹⁶³ concluding that many elements of the Computer Engineering program had stronger connections to the Engineering Directorate programs than to the CISE programs. The committee members were particularly concerned about a home for the joint DARPA-NSF MOSIS VLSI fabrication facility, which became a key activity of the MIPS Division

in CISE. Norman Caplan, the Deputy Division Director of ECSE, forwarded to the attendees his memorandum¹⁶⁴ to Nam Suh concerning robotics, which raised concerns about the definition of the field of robotics being assigned to CISE.

Following the March 18th meeting of the Engineering Advisory Committee, Frank Garry, its chair, wrote to Bloch that the Committee concurred with “the consolidation of the following into the new CISE Directorate: DCR/MPS, IST/BBS, OASC/OD and the Program in Computer Engineering from the Engineering Directorate.” Garry went on to say that “the remaining Engineering programs outlined in Gordon Bell’s memorandum of February 27 have their intellectual base in the Engineering Directorate” and we “fear that their transfer to CISE would narrow their focus and eventually erode their disciplinary strength.” The recent reorganization of Engineering had been carefully constructed and based on broad input from the Advisory Committee, the National Academy of Engineering, and other members of the engineering community. It “would be precipitous to alter the [d]irectorate’s programs in a major way without a similar review.”¹⁶⁵

Other Engineering Directorate managers also pushed back against Bell. Frank Huband stated, “The creation of a computer-related directorate is an exciting event, and has the potential to create new opportunities for development in this important discipline” but research funded in CISE “must pass muster as computer-related” . . . the practitioners in non-computer disciplines “want—and I believe deserve—an independent home for their research proposals.”¹⁶⁶ “Assignment of program elements to ENG or CISE should be governed by considerations of their fit to the respective missions of the directorates, and of their relative contributions to strengthening U.S. research in computing or in engineering,”¹⁶⁷ suggested program director Howard Moraff. He also raised concerns about possible disruptions to the new programs created in Nam Suh’s recent reorganization of Engineering and urged collaboration between CISE and Engineering.

Bell wrote to Bloch defending his approach to organizing the directorate, suggesting:

The rationale for Engineering disciplines in CISE is that contemporary engineering naturally divides into two parts: those research areas based on the physical and biological sciences and those which deal with information. This provides a working boundary: on one side are research areas governed by the physical transformations, and on the other side those research areas concerned with the transformation of information from one form to another.¹⁶⁸

In Bell’s personal notes dated March 20, 1986, he characterized the Nam Suh position as being that CISE should be “pure computer,” while Bell himself thought it should be “pure information processing (and storage, transmission, switching,

transduction), robotics, and some cross-disciplinary programs that make significant use of the computer.” In these notes, Bell discussed other organizations of traditional EE and computing, and the National Academy of Engineering taxonomies.

On March 24, Frank Huband responded to Nam Suh¹⁶⁹ concerning the probable decision by Bloch to transfer Communication Systems and Signal Processing programs (CSSP) to CISE. Huband was concerned that “important parts of the current CSSP program will not be relevant to the purposes of CISE and may thus not be eligible for future funding.” This memorandum was forwarded to Bloch and Bell, and Bell forwarded it to Chuck Brownstein, Bernie Chern (who had been reassigned from the Computer Engineering program to CISE), and me. Bell commented that “we will work it out over the next couple of months along with MOSIS funding, etc.”

Bell sent his initial plan¹⁷⁰ for CISE to Bloch for approval via the Assistant Director for Administration on April 17, 1986. This memorandum proposed transferring the Division of Computer Research (DCR) from Mathematical and Physical Sciences, the Division of Information Science and Technology (DIST) from Biological and Behavioral Sciences, and the Office of Advanced Scientific Computing (OASC) from the Director’s Office intact. It proposed creating a new Division of Computer and Information Engineering (CIE) to temporarily house the computer engineering programs (Software Systems Design; Computer Systems Architecture; Vision, Robotic, and Knowledge-based Systems) and the Communications and Signal Processing programs, each from the Engineering Directorate. The proposal added a Division Director (Bernard Chern) for the new CIE Division, a CISE Planning Officer (Jerry Daen), an acting CISE Executive Officer (Charles Brownstein, who remained Director/DIST), and an acting senior scientist for planning and program development (W. Richards Adrion, who remained Deputy Director/DCR). The new directorate would have 54 positions, 49 permanent and 5 IPAs. Bloch approved, and the directorate was officially launched on May 1, 1986. At the time it was officially created, CISE had two advisory committees: Computer Research (mainly associated with DCR) and Advanced Scientific Computing (mainly associated with OASC).

In Bell’s presentation¹⁷¹ to the National Science Board in May, he described using the CISE research budget as a “balance wheel” to DARPA and industry. He defined five research areas: parallelism as applied to parallel processing; automation, robotics, and intelligent systems; ultra-large-scale integrated systems; advanced scientific and engineering computing; and networks and distributed computing. Bell focused on these areas because they had relatively clear, long-term goals; measurable output; an emphasis on maintaining U.S. leadership in computing; significant economic and competitive impact; and a demand for undergraduate and graduate training. Across these five initiatives, CISE would support basic,

front-end research throughout the entire computer research community at a time when DARPA was becoming more “mission oriented.”

Between April and August, a number of organizational structures were considered. Bell’s five initiatives helped structure the divisional organization of CISE. While the initiatives were cross-cutting, divisions were thought of as the leads: DCR for parallelism as applied to parallel processing; DIST for automation, robotics, and intelligent systems; CIE for ultra-large-scale integrated systems; and OASC for advanced scientific and engineering computing. Bell did not see OASC leading networks and distributed computing, and that eventually led to a fifth division in CISE. DIST leadership in automation, robotics, and intelligent systems led to a proposed ARIS division that eventually became Information, Robotics, and Intelligent Systems (IRIS) due to a continuing commitment to information sciences and systems. Initially CIE was to become the Ultra-Large-Scale Integration (ULSIS) Division; but with a responsibility for computing design and architecture, it was renamed the Microelectronic Information Processing Systems (MIPS) Division. The need to locate the communication and signal processing programs led to DCR taking on that responsibility as the Computer and Communications Research (CCR) Division. One other proposed restructuring would have divided OASC¹⁷² into three sections: Centers with Larry Lee as Head; Networking and Distributed Computing with Steve Wolff as Head; and a New Technologies Section with Al Harvey as Head. The plan included research and EXPRES program directors. The tension over NSFNET as a national vs. supercomputer network continued and John Connolly strongly objected.

On August 26, 1986, Bell proposed¹⁷³ restructuring CISE by reconfiguring the DCR, moving the Intelligent Systems program to DIST, renaming DCR as the Division of Computer and Computation Research, restructuring DIST and renaming it the Division of Information, Robotics, and Intelligent Systems, and restructuring CIE and naming it the Division of Microelectronic Information Processing Systems. The Office of Advanced Scientific Computing was renamed the Division of Advanced Scientific Computing. In the divisions, programs were restructured to reflect a new divisional mission. This plan was approved and became official¹⁷⁴ on October 17, 1986.

This reorganization resulted in some personnel changes. Earlier in July, Chuck Brownstein officially was named CISE Executive Officer. (He had served briefly as Acting AD/CISE until Bell was sworn in on June 17, 1986.) While I would resign officially on September 14,¹⁷⁵ remaining on a part-time basis through the fall, Gordon Bell already knew this when he wrote his August 26 memorandum. He officially transferred me into the position from DCR and argued to keep the position after my departure. Kent Curtis filled this position a year later. EXPRES was moved from

DASC into the AD's office. The last organizational change occurred on December 10, 1986, when the Networking and Communications programs in DASC became a fifth division in CISE: the Networking and Communications Research and Infrastructure Division (NCRI), with Steve Wolff as DD.

While the organizational debates were going on, the staff that were clearly moving to CISE began to address the planning process outlined by Bridgewater. I wrote¹⁷⁶ in March 1986 to Chuck Brownstein, Bernie Chern, John Connolly, and Kent Curtis about long range planning for CISE, asking them to produce planning documents covering both existing programs and potential new initiatives. This was a three-part request: an exercise to "define the base," long-range strategic planning, and issue papers on new initiatives. All of these were due March 17.

I wrote¹⁷⁷ again to the (unofficial) CISE division directors about the need for them to develop plans to address the five Bell initiatives, asking for two-page position papers. Each acting division director needed to answer four questions Bloch had asked each directorate to address: What difference has NSF support made? What is the NSF role in [discipline] research? What are the programmatic gaps? What are your priorities in the event of reductions? Table 2.2 includes excerpts from the Long-Range Planning Material submitted to the Office of Budget, Audit, and Control¹⁷⁸ in April 1986.

Each of five research initiatives included research opportunities/breakthroughs needed, current efforts, and plans and initiatives. Two-page position papers were written on (1) parallelism (lead: Rick Adrion); (2) advanced scientific computing (lead: John Connolly); (3) networking (lead: John Connolly); (4) fabrication facilities expanding the MOSIS concept (lead: Bernie Chern); (5) robotics (lead: Chuck Brownstein); and (6) experimental systems (lead: Robert Minnick).

In July, Chuck Brownstein asked¹⁷⁹ the division directors to prepare backup materials for the FY 1988 budget request. CISE was requesting a \$69.02 million increase to \$192.00 million, a 56% increase over the FY 1987 current plan. Proposed initiatives included project and instrumentation support for research on parallel techniques of computing and information processing to be expanded throughout CISE; several large group or "mini-center" awards to be made to promote experimentation with large-scale systems; additional infrastructure for use throughout U.S. academic institutions, including upgrading instrumentation and improving laboratories; a major effort to be undertaken to expand university research and teaching in the design, fabrication, and use of integrated microelectronics; and a commitment to advancing and accelerating the state-of-the-art of advanced scientific computing.

More on the development of CISE is covered in the following chapters.

Table 2.2 Answers to the “Four Questions” in the 1986 long-range planning exercise

What difference has NSF support made?	CISE programs have improved the knowledge base for research and commerce and have developed the national scientific and engineering personnel and facility infrastructure required for the maintenance of U.S. leadership. CISE also has a unique NSF role in the improvement of scientific and engineering computing and communications through shared use facilities, training, and network links among researchers.
What is the NSF role in Computer and Information Processing research?	NSF has unique responsibility for long-term and theoretical research and for the broad base of academic research. NSF is critical for the improvement of the national academic infrastructure in the CISE areas. Advanced Scientific Computing is a unique first step toward conditioning the general scientific research community to use computing as a new mode of research.
What are the Programmatic Gaps?	The academic base of computing and information processing research and the support base from both federal and private sources are too small. Too many research universities lack the necessary manpower and infrastructure in computing and information processing research to achieve the critical mass needed for quality research in this important technology.
What are your priorities in the event of reductions?	None of the major areas of CISE would be dropped; fewer awards would be made. Graduate student and young faculty support will have a high priority. Grant sizes will be maintained. Network access to the supercomputer centers will be given higher priority than maintaining all the centers. If necessary, we would reduce the CER program in favor of individual faculty research project support.

2.13 Summary and Conclusions

In a mere 12 years, computing programs at NSF transitioned from two weakened offices, OCA and OSIS, to a directorate that had positioned itself to lead the major national initiatives described in later chapters. Along the way, a number of significant initiatives and activities fundamentally changed the perception of computing as a discipline. Not only did dozens of Turing Award winners begin their careers with NSF funding, but so did hundreds of ACM, IEEE, AAI, and AAAS fellows. Theorynet led to CSNET and then to NSFNET. The Computer Science Research

Table 2.3 NSF CISE FY 1988 budget request

Organization	FY 1986 Actual	FY 1987 Request	FY 1987 Current Plan	FY 1988 Request
Computer and Communication Research	\$38.23	\$36.98	\$36.98	\$59.60
Information, Robotics, and Intelligent Systems	\$8.81	\$16.30	\$16.30	\$26.40
Microelectronic Information Processing Systems	\$10.47	\$12.20	\$12.20	\$24.20
Advanced Scientific Computing	\$36.25	\$46.60	\$46.60	\$61.80
Networking and Distributed Computing	\$6.80	\$10.90	\$10.80	\$20.00
Total	\$100.56	\$122.98	\$122.98	\$192.00

Source: NSF OBAC 1986

Equipment program laid the ground work for the Coordinated Experimental Research program. CER fundamentally altered the capacity for experimental research in colleges and universities.

I am fortunate to have had a career that spanned those 12 years and the opportunity to observe how far the field came during those years and to contribute to its growth. The narrative above names a number of important people, but it omits a great number of administrators, program managers, program assistants, and other staff who made the successes of the period possible.

Notes

1. Engineering became a separate directorate in 1978 and the Computer Science Section remained in the Mathematical and Physical Sciences (MPS) Directorate.
2. Theoretical Computer Science was a program in the Computer Science Section (CSS) of the Mathematical and Computer Science (MCS) Division within the Mathematical and Physical Sciences (MPS) Directorate.
3. National Research Council. 1982. Ad Hoc Panel to Study the Conduct of Basic Research in Computer Science and Its Interaction with Applied Research and Development. National Academies.

4. L. Fein. 1959. The role of the university in computers, data processing, and related fields. *Communications of the ACM*, 2(9): 7–14. DOI: [10.1145/368424.368427](https://doi.org/10.1145/368424.368427).
5. S. Gorn. 1963. The computer and information sciences: A new basic discipline. *SIAM Review*, 5(2): 150–155. <https://www.jstor.org/stable/2027479>.
6. G. E. Forsythe. 1967. A university's educational program in computer science. *Communications of the ACM*, 10(1): 3–11. DOI: [10.1145/363018.363038](https://doi.org/10.1145/363018.363038).
7. A. Newell, A. J. Perlis, and H. A. Simon. 1967. Computer science. *Science*, 157(3795): 1373–1374. DOI: [10.1126/science.157.3795.1373-b](https://doi.org/10.1126/science.157.3795.1373-b).
8. D. E. Knuth. 1968. *The Art of Computer Programming. Volume 1: Fundamental Algorithms* (1st ed.). Reading, MA: Addison-Wesley Publishing.
9. P. Wegner. 1968. *Programming Languages, Information Structures, and Machine Organization*, McGraw Hill.
10. T. J. Misa, ed. 2016. *Communities of Computing: Computer Science and Society in the ACM*, Morgan & Claypool.
11. J. Abbate. October 2013. Is computer science “Science”? A half-century debate. Keynote talk, 2nd International Conference on History and Philosophy of Computing, Paris. <https://hapoc2013.sciencesconf.org/27047/document>.
12. S. D. Conte, J. W. Hamblen, W. B. Kehl, S. O. Navarro, W. C. Rheinboldt, D. M. Young, Jr., and W. F. Atchinson. 1965. An undergraduate program in computer science—Preliminary recommendations. *Communications of the ACM*, 8(9): 543–552. DOI: [10.1145/365559.366069](https://doi.org/10.1145/365559.366069).
13. G. E. Forsythe. 1965. President's letter to the ACM membership: Why ACM? *Communications of the ACM*, 8(3): 143–144. DOI: [10.1145/363791.363792](https://doi.org/10.1145/363791.363792).
14. H. A. Simon. 1969. *The Sciences of the Artificial*. Cambridge, MA: MIT Press, p. 21.
15. J. Abbate. 2016. From handmaiden to “proper intellectual discipline”: Creating a scientific identity for computer science in 1960s America. In *Communities of Computing: Computer Science and Society in the ACM* (edited by T. Misa). Morgan & Claypool, p. 28.
16. Gordon Bell. January 23, 1974. Letter to Edward C. Creutz (AD/R), internal document.
17. This idea would return in the Feldman, Snowbird, and other reports.
18. National Science Foundation. 1974. *Twenty-Fourth Annual Report for Fiscal Year 1974*. Washington, DC: U.S. Government Printing Office.
19. B. W. Arden. 1976. The computer science and engineering research study (COSERS). *Communications of the ACM*, 19(12): 670–673. DOI: [10.1145/360373.360376](https://doi.org/10.1145/360373.360376).
20. B. W. Arden. 1983. *What Can Be Automated?: Computer Science and Engineering Research Study*. MIT Press.
21. J. A. Feldman and W. R. Sutherland. 1979. Rejuvenating experimental computer science: A report to the National Science Foundation and others. *Communications of the ACM*, 22(9): 497–502. DOI: [10.1145/359146.359147](https://doi.org/10.1145/359146.359147).
22. D. D. McCracken, P. J. Denning, and D. H. Brandin. 1979. An ACM executive committee position on the crisis in experimental computer science. *Communications of the ACM*, 22(9): 503–504. DOI: [10.1145/359146.362786](https://doi.org/10.1145/359146.362786).

23. P. J. Denning, E. A. Feigenbaum, P. Gilmore, A. C. Hearn, R. W. Ritchie, and J. F. Traub. 1981. A discipline in crisis. *Communications of the ACM*, 24(6): 370–374. DOI: [10.1145/358669.358682](https://doi.org/10.1145/358669.358682).
24. National Science Foundation. May 21–23, 1979. Summary Minutes of the Advisory Subcommittee for Computer Science. Washington, DC.
25. P. J. Denning. 1981. ACM President's letter: Eating our seed corn. *Communications of the ACM*, 24(6): 341–343. DOI: [10.1145/358669.358672](https://doi.org/10.1145/358669.358672).
26. J. F. Traub. 1981. Quo vadimus: Computer science in a decade. *Communications of the ACM*, 24(6): 351–369. DOI: [10.1145/358669.358677](https://doi.org/10.1145/358669.358677).
27. J. Hopcroft, B. Lampson, J. McCarthy, A. Penzias, M. Rabin, J. Schwartz, D. Scott, and A. Van Dam. December 1986. *Preliminary Draft Report on the Scientific Contributions of Computer Science*. National Science Foundation.
28. J. E. Hopcroft, K. W. Kennedy, and K. K. Curtis. 1989. *Computer Science: Achievements and Opportunities*. Philadelphia: Society for Industrial and Applied Mathematics.
29. J. McCarthy. April 13, 1996. *A Petition for the Withdrawal of "Computing the Future"*. <http://jmc.stanford.edu/commentary/petition.html>; last accessed 31 December 2018.
30. National Research Council. 1997. *Defining a Decade: Envisioning CSTB's Second 10 Years*. Washington, DC: The National Academies Press.
31. L. Kleinrock. 1998. *Toward a National Research Network*. National Academies.
32. J. F. Traub and National Research Council. 1988. *The National Challenge in Computer Science and Technology*. National Academies.
33. Links to CSTB publications can be found here: http://sites.nationalacademies.org/CSTB/CSTB_042201.
34. National Research Council, 1997.
35. CRA provided three earlier reports: (1) P. Freeman and W. Aspray. 1999. *The Supply of Information Technology Workers in the United States*. <https://eric.ed.gov/?id=ED459346>. (2) J. Cuny and W. Aspray. 2001. *Recruitment and Retention of Women Graduate Students in Computer Science and Engineering*. Computing Research Association. https://www.researchgate.net/publication/239547208_Recruitment_and_retention_of_women_graduate_students_in_computer_science_and_engineering. (3) J. Stankovic and W. Aspray. 2003. *Recruitment and Retention of Faculty in Computer Science and Engineering*. Computing Research Association. <http://archive.cra.org/reports/r&rfaculty.pdf>.
36. Y. Oyanagi. 1999. Development of supercomputers in Japan: Hardware and software. *Parallel Computing*, 25: 1545–1567. DOI: [10.1016/S0167-8191\(99\)00084-8](https://doi.org/10.1016/S0167-8191(99)00084-8).
37. National Research Council. 2005. *Getting up to Speed: The Future of Supercomputing*. National Academies Press.
38. Panel on Large Scale Computing in Science and Engineering, and P. D. Lax. 1983. *Report of the Panel on Large Scale Computing in Science and Engineering*. National Science Foundation. https://www.pnnl.gov/scales/docs/lax_report1982.pdf.

39. K. G. Wilson. 1988. *On Supercomputing*. In Panel on Large Scale Computing in Science and Engineering and P. D. Lax. 1983. National Science Foundation. Supplement to the “Report of the Panel on Large Scale Computing in Science and Engineering,” pp. 24–33.
40. J. F. Decker. 1983. “Report to the Federal Coordinating Council on Science, Engineering and Technology Supercomputer Panel on Recommended Government Actions to Provide Access to Supercomputers.” In *Supercomputers, Hearings before the Committee on Science and Technology, U.S. House of Representatives, Ninety-eighth Congress, First Session, November*, vol. 15, p. 16.
41. M. Bardou. 1983. *A National Computing Environment for Academic Research*. National Science Foundation: Report of an NSF Working Group on Computers for Research.
42. Committee on Science, Engineering, and Public Policy. September 7, 1984. *Report of the Research Briefing Panel on Computer Architecture*, National Academies of Sciences and Engineering and the Institute of Medicine. DOI: [10.17226/19410](https://doi.org/10.17226/19410).
43. Erich Bloch. September 11, 1984. NSF internal memorandum to Carl Hall, AAD/ENG, Marcel Bardou, AAD/MPS, and John Connolly, OD/OASC. Charles Babbage Institute.
44. Kent Curtis, DD/DCR, Blake Cherrington, DD/ECSE, and John Connolly, Head/OASC. September 24, 1984. NSF internal memorandum to Erich Bloch, Director. Charles Babbage Institute.
45. R. Adrion, D. Farber, F. F. Kuo, L. H. Landweber, D. Angel, and J. B. Wyatt. December 1984. *SCIENCENET: Report on the Evolution of a National Supercomputer Access Network*. National Science Foundation.
46. D. M. Jennings, L. H. Landweber, I. H. Fuchs, D. J. Farber, and W. R. Adrion. 1986. Computer networking for scientists. *Science*, 231(4741): 943–950. DOI: [10.1126/science.231.4741.943](https://doi.org/10.1126/science.231.4741.943).
47. This quote is from the introduction to the program activities of the Division of Computing Research (DCR) in the Mathematical and Physical Sciences and Engineering (MPE) Directorate. *NSF Annual Report*, 1975.
48. I applied and was recruited to NSF by Kent Curtis and joined in a “rotator” position. NSF brings academics into serve as “rotating” program officers using several employment strategies including Visiting Scientist, Engineer, and Educator (VSEE); Intergovernmental Personnel Act (IPA); and temporary excepted service assignments.
49. Congress changed the fiscal year start from July 1 to October 1 beginning in 1977, hoping the later start would find all federal budgets passed before the fiscal year began. The result was a need for a 3-month “Transition Quarter.” See J. J. Hogan. 1985. Ten years after: The US Congressional Budget and Impoundment Control Act of 1974. *Public Administration*, 63(2): 133–149. DOI: [10.1111/j.1467-9299.1985.tb00896.x](https://doi.org/10.1111/j.1467-9299.1985.tb00896.x).
50. ESL, Inc. was a defense contractor specializing in satellite and other photographic image processing. ESL was led by William Perry, later Secretary of Defense under Clinton and earlier Under Secretary of Defense for Research and Engineering from 1977 to 1981, and Robert Fossum, who was DARPA Director from 1976 to 1981.
51. Sponsored Projects Office, University of California, Berkeley. *Quick Guide to Working with NSF FastLane and Research.gov*. <https://spo.berkeley.edu/guide/fastlanequick.html>; last accessed 19 March 2019.

52. Theoretical Computer Science was the largest program and received only about 100 proposals per year.
53. Oral history, Rick Weingarten, interviewed by Peter Freeman, July 11, 2017. Charles Babbage Institute.
54. R. Adrion and S. Mahaney. 2003. Shooting inward. *Computing Research News*, 15(1). archive.cra.org/CRN/articles/jan03/adrion.mahaney.html; last accessed 15 January 2019.
55. Nomination information and winners can be found at <https://amturing.acm.org/>.
56. The Turing Award is presented each June at the ACM Awards Banquet and is accompanied by a prize of \$1,000,000 plus travel expenses to the banquet. Financial support for the award [currently] is provided by Google Inc.
57. D. E. Knuth. 1969. *The Art of Computer Programming, Vol. 1*. Reading, MA, Addison-Wesley Publishing Company.
58. D. Knuth. 1986. *Computers & Typesetting*. Reading, MA: Addison-Wesley.
59. Ted Codd, Dennis Ritchie, Ken Thompson, John Cocke, Ivan Sutherland, Butler Lampson, Dick Sterns, Doug Engelbart, Jim Gray, Alan Kay, Vint Cerf, Bob Kahn, Fran Allen, and Chuck Thacker were primarily in industry or government, while Tony Hoare, Steve Cook, Niklaus Wirth, Robin Milner, Ole Dahl, Kristen Nygaard, Peter Naur, Joseph Sifakis, Silvio Micali, and Tim Berners-Lee carried out their Turing-related research in Europe, Canada, or Israel.
60. R. K. Yin. 2011. *Applications of Case Study Research*. Sage Publishers.
61. D. Leavitt. May 3, 1972. NSF-sponsored research could lead to national science computer network. *Computerworld*, p. 10.
62. D. D. Aufenkamp and E. C. Weiss. 1972. NSF activities related to a national science computer network. *Computer Communications: Impacts and Implications*, p. 226.
63. M. Greenberger, J. Aronofsky, J. L. McKenney, and W. F. Massy. 1973. Computer and information networks: The movement in research and education toward national resource sharing via networks is accelerating. *Science*, 182(4107): 29–35. DOI: [10.1126/science.182.4107.29](https://doi.org/10.1126/science.182.4107.29).
64. J. Abbate. 2000. *Inventing the Internet*. MIT Press.
65. Telenet was developed by Larry Roberts, the original ARPANET program manager.
66. S. R. Hiltz. June 1981. *The Impact of a Computerized Conferencing System on Scientific Research Communities*. FINAL REPORT NSF-MCS-77-27813.
67. M. E. Smid and D. K. Branstad. 1988. Data encryption standard: Past and future. *Proceedings of the IEEE*, 76(5): 550–559. DOI: [10.1109/5.4441](https://doi.org/10.1109/5.4441).
68. W. Diffie and M. E. Hellman. 1977. Special feature exhaustive cryptanalysis of the NBS data encryption standard. *Computer*, 10(6): 74–84. DOI: [10.1109/C-M.1977.217750](https://doi.org/10.1109/C-M.1977.217750).
69. FIPS Publication 46-3. 1999. Data encryption standard (DES). National Institute of Standards and Technology 25(10): 1–22. <https://csrc.nist.gov/csrc/media/publications/fips/46/3/archive/1999-10-25/documents/fips46-3.pdf>.
70. W. Diffie and M. Hellman. 1976. New directions in cryptography. *IEEE Transactions on Information Theory*, 22(6): 644–654. DOI: [10.1109/TIT.1976.1055638](https://doi.org/10.1109/TIT.1976.1055638).

71. R. L. Rivest, A. Shamir, and L. Adleman. 1978. A method for obtaining digital signatures and public-key cryptosystems. *Communications of the ACM*, 21(2): 120–126. DOI: [10.1145/359340.359342](https://doi.org/10.1145/359340.359342).
72. J. H. Ellis. 1970. The possibility of secure non-secret digital encryption. *UK Communications Electronics Security Group*, vol. 6. <http://cryptocellar.org/cesg/possnse.pdf>.
73. C. C. Cocks. 1973. “A note on non-secret encryption.” *CESG Memo*. <https://www.semanticscholar.org/paper/A-note-on-non-secret-encryption-Cocks/fab40f7645d931a6094e1d8c0113143091942ad3>.
74. D. Shapley and G. B. Kolata. 1977. Cryptology: Scientists puzzle over threat to open research. *Science*, 197(4311): 1345–1349. <http://science.sciencemag.org/>; last accessed 2 March 2018. Also: D. Shapley. 1977. Cryptography meeting goes smoothly. *Science*, 198(4316): 476. <http://science.sciencemag.org/>; last accessed 2 March 2018.
75. J. C. Cherniavsky. 1985. Case study: Openness and secrecy in computer research. *Science, Technology, & Human Values*, 10(2): 99–103. DOI: [10.1177/016224398501000214](https://doi.org/10.1177/016224398501000214).
76. R. C. Buck. 1982. The public cryptography study group. *Computers & Security*, 1(3): 249–254. DOI: [10.1016/0167-4048\(82\)90043-8](https://doi.org/10.1016/0167-4048(82)90043-8).
77. P. J. Denning, D. H. Brandin, D. C. Schwartz, and G. I. Davida. 1981. Report of the public cryptography study group. *Communications of the ACM*, 24(7): 434–450. DOI: [10.1145/358699.358710](https://doi.org/10.1145/358699.358710).
78. Denning et al., 1981, *op. cit.*
79. G. B. Kolata. 1980. Cryptography: A new clash between academic freedom and national security. *Science*, 209(4460): 995–996. DOI: [10.1126/science.209.4460.995](https://doi.org/10.1126/science.209.4460.995).
80. B. R. Inman. 1980. Cryptography research funding. *Science*, 210(4466): 134. DOI: [10.1126/science.210.4466.134](https://doi.org/10.1126/science.210.4466.134).
81. D. A. Langenberg. 1980. Cryptography, NSF, and NSA. *Science Magazine*, 210(4473). DOI: [10.1126/science.210.4473.960](https://doi.org/10.1126/science.210.4473.960).
82. I left NSF in the fall of 1978 and returned to manage a different set of programs in January 1980, so I am not familiar with the process used by the program directors who followed me. I am quite confident that something similar to this process continues today.
83. Institute of Medicine, National Academy of Sciences, and National Academy of Engineering. 1982. *Scientific Communication and National Security*. Washington, DC: The National Academies Press. DOI: [10.17226/253](https://doi.org/10.17226/253).
84. Minutes of the May 28–29, 1981, meeting of the NSF Advisory Subcommittee on Computer Science. Jack Minker, Chair. Charles Babbage Institute.
85. J. Guttag et al. May 1981. *The Role of the NSF in Supporting Cryptological Research: A Report to the National Science Foundation by Its Mathematical and Computer Sciences Advisory Committee*. National Science Foundation. Charles Babbage Institute.
86. Thelma Estrin was married to Gerald Estrin and both were UCLA computer science professors. They had three daughters. Margo Estrin is a medical doctor. Deborah Estrin is a leading computer scientist, who moved from USC to UCLA after her parents retired. She led the NSF Science and Technology Center for Embedded Networked Sensing (CENS). Recently, she joined Cornell Tech. Judith Estrin is a corporate executive and entrepreneur,

- having founded Zilog, Bridge Communications, Network Computing Devices, Precept Software, Packet Design, and JLABS. Judy served as chief technology officer and senior vice president of Cisco Systems until 2000 and CEO of Eventlive in 2013.
87. Ettore F. “Jim” Infante had served as program director for Applied Mathematics in MCS a few years earlier.
 88. K. Curtis. May 21–23, 1979. *Summary Minutes, Advisory Subcommittee for Computer Science*. National Science Foundation. Charles Babbage Institute.
 89. Curtis, 1979, *op. cit.*, p. 2.
 90. Curtis, 1979, *op. cit.*, p. 3.
 91. P. Young. December 9, 1979. “ReDRAFT: NSF support for experimental computer science,” Internal document. Charles Babbage Institute.
 92. Oral history, Rick Adrion, interviewed by William Aspray, October 29, 1990, in Amherst, Massachusetts. Charles Babbage Institute. <https://conservancy.umn.edu/handle/11299/104300>; last accessed 15 January 2019.
 93. R. Riesenfeld and L. Hollaar. 1984. *NSF Grant 8406682-1984 NSF CER Conference to Be Held February 23–24, 1984, at the University of Utah, Salt Lake City, Utah (Computer Research)*. https://www.nsf.gov/awardsearch/showAward?AWD_ID=8406682&HistoricalAwards=false; last accessed 14 February 2019.
 94. J. Schwartz. November 1983. *A Taxonomic Table of Parallel Computers, Based on 55 Designs*. Ultracomputer Note #69. Courant Institute, N.Y.U.
 95. Ultracomputer Project, <https://cs.nyu.edu/cs/projects/ultra/home-page.html>; last accessed 15 February 2019.
 96. D. Gajski, D. J. Kuck, D. H. Lawrie, and A. S. Sameh. 1983. CEDAR—A large scale multiprocessor. *Proceedings of the 1983 International Conference on Parallel Processing*, Belaire, MI. DOI: [10.1145/859526.859527](https://doi.org/10.1145/859526.859527).
 97. M. Sejnowski, E. T. Upchurch, R. N. Kapur, D. P. S. Charlu, and G. J. Lipovski. 1980. An overview of the Texas reconfigurable array computer. *AFIPS, Proceedings of the May 19–22, 1980, National Computer Conference*, p. 631. DOI: [10.1145/1500518.1500623](https://doi.org/10.1145/1500518.1500623).
 98. J. R. Goodman and C. H. Sequin. 1981. Hypertree: A multiprocessor interconnection topology. *IEEE Transactions on Computers*, 20(12): 923–933. DOI: [10.1109/TC.1981.1675731](https://doi.org/10.1109/TC.1981.1675731).
 99. R. M. Keller, G. Lindstrom, and S. Patil. 1979. A loosely-coupled applicative multi-processing system. *AFIPS Conference Proceedings*, 48: 613–622. DOI: [10.1109/AFIPS.1979.8](https://doi.org/10.1109/AFIPS.1979.8).
 100. D. J. DeWitt, R. Gerber, G. Graefe, M. Heytens, K. Kumar, and M. Muralikrishna. 1986. “Gamma—A high performance dataflow database machine.” University of Wisconsin–Madison Department of Computer Sciences.
 101. C. Rieger, R. H. Trigg, and B. Bane. 1981. ZMOB: A new computing engine for AI. *IJCAI*, pp. 955–960. <https://www.ijcai.org/Proceedings/81-2/Papers/071.pdf>.
 102. J. A. Fisher. 1983. Very long instruction word architectures and the ELI-512. *ACM*, 11(3). DOI: [10.1145/800046.801649](https://doi.org/10.1145/800046.801649).

103. R. Wagner. 1983. The Boolean vector machine [BVM]. *Proceedings of the 10th International Symposium on Computer Architecture*, Stockholm, Sweden, June 13–16, pp. 59–66. DOI: [10.1145/800046.801639](https://doi.org/10.1145/800046.801639).
104. L. Snyder. 2011. Blue CHiP. *Encyclopedia of Parallel Computing*, pp. 165–175. DOI: [10.1007/978-0-387-09766-4_237](https://doi.org/10.1007/978-0-387-09766-4_237).
105. G. T. Almes, A. P. Black, E. D. Lazowska, and J. D. Noe. 1985. The Eden System: A technical review. *IEEE Transactions on Software Engineering*, 1: 43–59. DOI: [10.1109/TSE.1985.231536](https://doi.org/10.1109/TSE.1985.231536).
106. D. J. Dewitt, R. A. Finkel, and M. H. Solomon. 1984. *The CRYSTAL Multicomputer: Design and Implementation Experience*. University of Wisconsin–Madison Department of Computer Sciences.
107. K. P. Birman, T. A. Joseph, T. Raechle, and A. El Abbadi. 1985. Implementing fault-tolerant distributed objects. *IEEE Transactions on Software Engineering*, 6: 502–508. DOI: [10.1109/TSE.1985.232242](https://doi.org/10.1109/TSE.1985.232242).
108. J. L. Bates and R. L. Constable. 1985. Proofs as programs. *ACM Transactions on Programming Languages and Systems* (TOPLAS) 7.1: 113–136. DOI: [10.1145/2363.2528](https://doi.org/10.1145/2363.2528).
109. Minutes of the May 28–29, 1981, meeting of the NSF Advisory Subcommittee on Computer Science. Jack Minker, Chair. Charles Babbage Institute.
110. W. R. Adrion and B. H. Barnes. 1987. Trends in National Science Foundation funded research and their impact on software engineering education. In *Software Engineering Education*, pp. 322–329.
111. D. Gries, R. Miller, R. Ritchie, and P. Young. 1986. Imbalance between growth and funding in academic computing science: Two trends colliding. *Communications of the ACM*, 29(9): 870–878. DOI: [10.1145/6592.6596](https://doi.org/10.1145/6592.6596).
112. B. Simons and J. Yudken. December 1987. Federal funding in computer science: A preliminary report. *SIGACT News*, 19(1): 54–63. DOI: [10.1145/38309.38310](https://doi.org/10.1145/38309.38310).
113. W. R. Adrion, E. D. Lazowska, and A. van Dam. 1992. From discipline in crisis to mature science: Evolving needs for computing research infrastructure. *Computer*, 12: 18–24.
114. The term “computer science department” is meant to be inclusive of the academic units covering computer science and engineering, computer and information science, electrical and computer engineering, as well as variations and differing organization unit types (departments, schools, colleges).
115. NSF Program Solicitation 19-512, CISE Community Research Infrastructure (CCRI). <https://www.nsf.gov/pubs/2019/nsf19512/nsf19512.pdf>.
116. The ARPANET Directory published in December 1979 by the Network Information Center at SRI International for the Defense Communications Agency includes most of the email addresses of individuals with ARPANET accounts. Approximately 3,600 individuals are listed.
117. P. J. Denning, A. Hearn, and C. W. Kern. March 1983. History and overview of CSNET. *ACM SIGCOMM*, pp. 138–145. DOI: [10.1145/1024840.103526](https://doi.org/10.1145/1024840.103526).
118. Attendees besides Landweber, Curtis, and Kahn included Rich DeMillo (Georgia Tech), Peter Denning (Purdue University), Ed Desautels (University of Wisconsin), Phil Enslow

(Georgia Tech), Dick Fateman (University of California, Berkeley), Mike Fisher (University of Washington), Randy Frank (University of Utah), Tony Hearn (University of Utah), Bob Ritchie (University of Washington), and Larry Rowe (University of California, Berkeley).

119. Denning, Hearn, and Kern, 1982, *op. cit.*, p. 139.
120. Denning, Hearn, and Kern, 1982, *op. cit.*, p. 139.
121. D. H. Crocker, E. S. Szurkowski, and D. J. Farber. November 1979. An internetwork memo distribution capability—MMDF. *Proceedings of the Sixth Symposium on Data Communications*. ACM, pp. 18–25. DOI: [10.1145/800092.802976](https://doi.org/10.1145/800092.802976).
122. D. A. Nowitz and M. E. Lesk. 1978. “A dial-up network of UNIX™ systems.” Bell Laboratories, Murray Hill, New Jersey.
123. D. Comer. 1983. The computer science research network CSNET: A history and status report. *Communications of the ACM*, 26(10): 747–753. DOI: [10.1145/358413.358423](https://doi.org/10.1145/358413.358423).
124. J. Minker. May 28–29, 1981. Summary Minutes, Advisory Subcommittee for Computer Science. Washington, DC: National Science Foundation. Charles Babbage Institute.
125. Dave Farber. January 16 (10:29:05) 1982. Posting on the UUCP bulletin board fa.arpa-bboard.
126. CSNET CIC. May 1983. *CSNET News*, vol. 1.
127. UCAR was the management organization for the NSF National Center for Atmospheric Research (NCAR).
128. CSNET CIC. October 1983. *CSNET News*, vol. 2.
129. The members of the CSNET Executive Committee were (* indicates they are shown in Figure 2.7): Peter J. Denning*, Chairman; Joel Birnbaum, Hewlett-Packard Laboratories; J. C. Browne, University of Texas at Austin; David J. Farber*, University of Delaware; Jerome A. Feldman*, University of Rochester; Bernard A. Galler*, University of Michigan; Michael A. Harrison, University of California, Berkeley; Anthony C. Hearn*, The Rand Corporation; Alan Kay*, Atari Corporation; Kenneth W. Kennedy, Rice University; Lawrence H. Landweber, University of Wisconsin-Madison; Richard LeFaivre*, Tektronix, Inc.; Raymond E. Miller*, Georgia Institute of Technology; Richard F. Riesenfeld, University of Utah; Jerome H. Saltzer, Massachusetts Institute of Technology; Patricia Selinger*, IBM Research Laboratories; Andries van Dam*, Brown University; Vic A. Vyssotsky, Bell Laboratories; Leonard C. Romney*, University Corporation for Atmospheric Research (ex officio); Richard D. Edmiston*, BBN Laboratories Inc. (ex officio); W. Richards Adrion*, National Science Foundation (observer); Kent Curtis*, National Science Foundation (observer); Barry Leiner, DARPA (observer).
130. L. Landweber. 2012. *InternetHistory.Asia*. <http://InternetHistory.Asia>. <https://sites.google.com/site/internethistoryasia/book1/3-2-csne>; last accessed 4 March 2018.
131. To expand membership, CSNET had renamed itself the Computing + Science Network.
132. W. R. Adrion. January 28, 1989. “Detailed Business and Transition Plan for ONENET.” ONENET Transition Team. Charles Babbage Institute.
133. Panel on Large Scale Computing in Science and Engineering and P. D. Lax, 1983.

134. The members included Don Aufenkamp, John Connolly, Kent Curtis (Chair), Jerry Daen, Thelma Estrin, Stephen Gould, Richard Isaacson, Arthur Kowalsky, Larry Lee, Edward McCullough, Michael J. McGill, Jane Stutsman, and Al Thaler.
135. Bardon, 1983, *op. cit.*
136. *Hearings before the Committee on Science and Technology*, U.S. House of Representatives, Ninety-eighth Congress, First Session, vol. 15, p. 16 (November 1985).
137. In the 1980s a class IV computer was capable of millions of floating-point operations per second (MFLOPS), in the classification scheme: microcomputers; minicomputers; mainframes; and supercomputers.
138. The members included Joan M. Centrella (University of Texas at Austin), Margaret A. Drake (NCAR), Oscar N. Garcia (University of South Florida), James F. Gilbert (University of California, San Diego), Anthony C. Hearn (RAND), Cecil E. Leith, Jr. (Lawrence Livermore Laboratory), Cyrus Levinthal (Columbia University), David C. Nagel (NASA/AMES), Steven A. Orszag (MIT), Christopher A. Sims (University of Minnesota), Mark A. Stadtherr (University of Illinois at Urbana-Champaign), Frank H. Stillinger (AT&T Bell Laboratories), and Kenneth G. Wilson (Cornell University).
139. The members included Oscar N. Garcia (University of South Florida), David C. Nagel (NASA/AMES), and Kenneth G. Wilson (Cornell University) from the main advisory committee, along with Dave Farber (University of Delaware), Bob Kahn (DARPA), Frank Kuo (SRI), Larry Landweber (University of Wisconsin), and Jack McCredie (Digital).
140. Letter to Edward Knapp, National Science Foundation Director, from David J. Farber (Delaware), Ira H. Fuchs (CUNY), Lawrence H. Landweber (Wisconsin), Kenneth G. Wilson (Cornell), Jerome A. Feldman (Rochester), Bernard A. Galler (Michigan), Gene Golub (Stanford), Jerome H. Saltzer (MIT), Joe F. Traub (Columbia), Andy Van Dam (Brown), Mischa Schwartz (Columbia), Seymour V. Parter (Wisconsin), Kenneth M. King (Cornell), Chuck Dickens (SLAC/Stanford), Joshua Lederberg (Rockefeller), Melvin Ferentz (Rockefeller), Gerald Estrin (UCLA), and Juris Hartmanis (Cornell). December 1983.
141. Jacob T. Schwartz of NYU. December 20, 1983. Letter to Dr. Edward Hayes of the National Science Foundation. Charles Babbage Institute.
142. Franklin Kuo of SRI International. July 15, 1983. Letter to Kent Curtis. “ARPANET expansion.” Charles Babbage Institute.
143. The BBN-developed Interface Message Processor (IMP) was the forerunner of the Internet routers that came later.
144. Contestants vie for NSF Funds. *Computerworld*, February 11, 1986.
145. *A Timeline of NSF History*. <https://www.nsf.gov/about/history/timeline80s.jsp>; last accessed 21 January 2019.
146. R. Adrion, D. J. Farber, F. F. Kuo, L. H. Landweber, D. C. Nagel, J. B. Wyatt. December 1984. *SCIENCENET: Report on the Evolution of a National Supercomputer Access Network*. National Science Foundation.
147. L. Landweber and D. Farber. August 20, 1984. SCIENCENET Phase I Implementation Plan.
148. Oral history, Dennis Jennings, interviewed by Andreu Veà Baró, November 28, 2007. Charles Babbage Institute.

149. D. M. Jennings, L. H. Landweber, I. H. Fuchs, D. J. Farber, and W. R. Adrion. 1986. Computer networking for scientists. *Science*, 231(4741): 943–950. DOI: [10.1126/science.231.4741.943](https://doi.org/10.1126/science.231.4741.943).
150. H. Frystyk. July 1994. *The Internet Protocol Stack*. <https://www.w3.org/People/Frystyk/thesis/TcpIp.html>; last accessed 20 January 2019.
151. Jennings et al., 1986, *op. cit.*, p. 946.
152. Jennings et al., 1986, *op. cit.*, p. 945.
153. Oral history, Steve Wolff, interviewed by Rick Adrion, July 20, 2017. Charles Babbage Institute.
154. Kent Curtis, DD/DCR to Richard Nicolson, AD/MPS. April 18, 1985. “NSF Organization of Computing Services, Computer and Information Science, and Computer Engineering.” Internal memorandum. Charles Babbage Institute.
155. Albert Bridgewater, DAD/AAEO. February 25, 1986. Internal memorandum to Gordon Bell, ADD/CS. “Thoughts on establishing Computer Sciences Directorate.” Charles Babbage Institute.
156. Albert Bridgewater, DAD/AAEO. February 26, 1986. Internal memorandum to Gordon Bell, ADD/CS. “Thoughts for proactive ADD/CS.” Charles Babbage Institute.
157. Gordon Bell. February 26, 1986. Typed and hand-annotated notes. Charles Babbage Institute.
158. Gordon Bell. February 27, 1986. NSF internal memorandum to Erich Bloch, Director; Mary Clutter, O/D; and Nam Suh, AD/Engineering Subject. “Thoughts on the composition of the Computer and Information Sciences and Engineering (CISE) Directorate.” Charles Babbage Institute.
159. Nam Suh. February 28, 1986. NSF internal memorandum to Gordon Bell, AD/CISE; Erich Bloch, Director; and Mary Clutter, O/D. “Response to your memorandum of February 27, 1986, re: composition of CISE Directorate.” Charles Babbage Institute.
160. Erich Bloch. March 3, 1986. Consolidation of computer-related activities. Staff memorandum O/D 86-14. Charles Babbage Institute.
161. Nam Suh, AD/ENG. March 10, 1986. Internal memorandum to members of the NSF Engineering Advisory Committee. Charles Babbage Institute.
162. Herb Voelcker et al. March 12, 1986. *Two Rational Approaches for Integrating ‘CISE’ into NSF’s Pantheon of Disciplines*. A report of the Engineering Task Group on Computing. Charles Babbage Institute.
163. Herb Voelcker et al. Undated but likely March 1986. *Computer Engineering in the DMCE Division and Its Relation to the New CISE Directorate*. A report of the Engineering Task Group on Computing. Charles Babbage Institute.
164. Norman Caplan, DDD/ECSE. March 10, 1986. Internal memorandum to Nam Suh, AD. “Robotics research issues and definition.” Charles Babbage Institute.
165. Frederick Garry, Chairman, Advisory Committee for Engineering. March 18, 1986. Letter to Erich Bloch, “Rough draft of the views of the Advisory Committee for Engineering at a meeting convened by Dr. Nam P. Suh.” Charles Babbage Institute.

166. Frank Huband, DD/ECSE. March 10, 1986. Email to Gordon Bell and Erich Bloch. “The appropriate locus of ECSE programs.” Charles Babbage Institute.
167. Howard Moraff, DMCE. March 12, 1986. Internal memorandum to Task Group on Computer Reorganization. “Principles for evaluating program element transfer.” Charles Babbage Institute.
168. Gordon Bell. March 12, 1986. Internal memorandum to Erich Bloch. “Rationale for the engineering content of the CISE Directorate.” Charles Babbage Institute.
169. Frank Huband, DD/ECSE. March 24, 1986. Internal memorandum to Nam Suh, AD/ENG. “Transfer of the Communication Systems and Signal Processing (CSSP) program to the Computer and Information Sciences and Engineering Directorate (CISE).” Includes handwritten notes on it by Gordon Bell. Charles Babbage Institute.
170. Gordon Bell, AD/CISE. April 17, 1986. Internal memorandum to Erich Bloch, Director, via Geoffrey Fenstermacher, AD/A. “Creation of the Directorate for Computer and Information Science and Engineering and relocation of constituent elements.” Charles Babbage Institute.
171. Rick Adrion. Drafts of presentation on long-range plans. Personal files.
172. Rick Adrion, Chief Scientist/CISE. July 28, 1986. Internal memorandum to John Connolly, D/OASC. “Reorganization revision.” Personal files.
173. Gordon Bell, AD/CISE. August 26, 1986. Internal memorandum to Geoffrey Fenstermacher, AD/A. “Reorganization of the Directorate for Computer and Information Science and Engineering.” Rick Adrion’s personal files.
174. C. Gordon Bell. October 17, 1986. “Reorganization of the Directorate for Computer and Information Science and Engineering,” in AD/CISE Bulletin No. 87-01.
175. Memorandum. Rick Adrion’s personal files.
176. R. Adrion, DDD/DCR. Likely before March 17, 1986. Undated memorandum to C. Brownstein, DD/IST, B. Chern, DD/DMCE, J. Connolly, D/ASC, and K. Curtis, DD/DCR. “Long range planning for CISE.” Rick Adrion’s personal files.
177. R. Adrion, Chief Scientist/CISE. March 27, 1986. Internal memorandum to C. Brownstein, DD/IST, B. Chern, DD/CIE, J. Connolly, D/OASC, K. Curtis, DD/DCR. “Long range planning for CISE.” Rick Adrion’s personal files.
178. LRP material submitted to OBAC. April 25, 1986. CISE Administrative Records. National Science Foundation.
179. Charles Brownstein. July 30, 1986. Internal memorandum to Division Directors, CISE. “Preparation of subactivity sections for FY 1988 OMB budget request.” Charles Babbage Institute.



1986–1998: The New Directorate in a Period of Computer Science Expansion

Peter A. Freeman¹

This chapter covers the years from the founding of CISE² in 1986 through 1998. As will be evident, the new directorate quickly established its structure and importance in NSF and within the U.S. government, in spite of some pushback from other areas of NSF. The period witnessed a succession of short-term ADs (Gordon Bell, William Wulf, Nico Habermann, Paul Young, and Juris Hartmanis), and the continued questioning of the validity of computer science as a fundamental discipline. By the end of this period, though, CISE started receiving increased respect, greater funding, and sustained leadership from its scientific community.

From the beginning, the formation of CISE brought welcome change. “For the first time we had all those concerned with fundamental questions in computing gathered together without the distraction of those with other concerns being in the room,”³ to paraphrase Gordon Bell’s comment in a recent interview⁴ about the start of CISE. Bell credits Erich Bloch (NSF Director, 1984–1990) for the overall vision for CISE and for understanding the importance of computing throughout society. Bell viewed Bloch as a superb manager who understood that people primarily concerned with computing often shared more with each other than with their colleagues in other disciplines. Bloch’s vision and insight, coupled with his desire to rationalize the organizations reporting directly to him,⁵ provided the force to create CISE, effective July 1, 1986. Over the next year there were a number of changes in the internal structure of CISE, including increased use of rotators.⁶

Five people served as Assistant Director (AD)⁷ of CISE in its first twelve years. About a fourth of the time there was an Acting AD, who was a permanent NSF employee whose background was not in a core computing area. These ADs each served approximately two years, for a variety of personal reasons and perhaps also because of the lack of a service tradition in computer science.

In spite of this turnover, CISE's support for core research and related topics yielded some outstanding results: NSF helped create the Internet; provided development support for common software structures and services; and supported projects that in the pursuit of other objectives created the first widely distributed, easily usable web browser (Mosaic) and the Google search engine. Sustained funding from NSF also brought about the expansion and deepening of the U.S. research community in computing.

3.1 Initial Structure and Leadership of CISE: 1986–1987⁸

This short period at the beginning of CISE's existence can be characterized as one of firming up structure, operations, and divisional leadership.

The idea of pulling together the disparate programs and efforts in NSF addressing computing had been percolating for some years prior to 1986.⁹ As with many successful actions, multiple people played key roles, but in this case four key people stand out: Erich Bloch (in the background guiding the effort), Gordon Bell, Chuck Brownstein, and Rick Adrion. In addition, Jerry Daen was invaluable in dealing with budgetary matters, and Kent Curtis added essential programmatic and operational knowledge. While those named were clearly in favor of a new directorate, some pushback (as noted below) was encountered.¹⁰

Erich Bloch was probably the one person without whom the creation of CISE could not have happened, and perhaps might never have happened since the impact of computing on society was about to explode. He was a long-time computer designer at IBM,¹¹ leading the design of the early supercomputer Stretch and responsible for leading the hardware design of the highly successful IBM 360. When he was tapped by President Ronald Reagan to head NSF, he had most recently been vice president of technical personnel development at IBM. His broad vision of computing, deep technical experience, ability to look ahead, excellent leadership skills, and informed view of the world from his position in industry made him the right person at the right time.¹²

Gordon Bell, a noted computer architect of minicomputers and highly parallel machines, implemented Bloch's vision and added to it substantively in important ways, for example by insisting that networking be an independent activity in the new

organization.¹³ Although fundamentally an engineer, Bell had significant exposure to and participation in the forefront of computing research and development of computer science as an intellectual discipline when he was a faculty member at Carnegie Mellon University (CMU). He had been brought into NSF by Bloch in late 1985, first as a consultant and then as a rotator to be the first AD/CISE. The mutual respect and compatibility of worldviews that he and Bloch shared meant that together they, with the assistance and contributions of Adrion, Brownstein, and others, were able to bring CISE into being in short order.¹⁴

Chuck Brownstein, trained as a political scientist who used computer modeling in his research, had joined NSF from a university position in the 1970s. His knowledge of how NSF worked internally aided Bloch in organizing CISE. At the same time, his own grounding in one of the early uses of computing in a field not normally thought of as numerically based, along with his understanding of what was happening more broadly in the uses of computers, enriched the makeup of CISE by including programs in the social and organizational impacts of computing and in education. He was already working with Bloch when Bell joined NSF in early 1986.¹⁵

The fourth person in developing the details of CISE was Rick Adrion, who had been affiliated with NSF on and off since 1976. With a strong educational background, including a Ph.D. in computer engineering, he had already served in various positions at NSF including program manager and deputy division director (DDD) in the various divisions in the Mathematical and Physical Sciences Directorate (MPS), which had been supporting core computer science, and also as a program officer in the Office of Advanced Scientific Computing. When programs and divisions were being considered for inclusion in CISE, he was the intellectual architect of many of the programmatic details. In that role, he was made Senior (Chief) Scientist in the new AD/CISE office. Although he left shortly after the official start of CISE, he continued service on several CISE advisory committees and returned to NSF as Division Director (DD) in 2000–2003.¹⁶

3.2 Organizational Initiation

CISE officially came into existence on July 1, 1986. It was created by combining entire divisions, programs, and parts of programs from across NSF. Table 3.1 shows how this developed. CISE executive staff included Gordon Bell, Chuck Brownstein (Executive Officer), Rick Adrion, and Jerry Daen.

Bell organized CISE into five technical areas: Parallel Processing; Automation, Robotics and Intelligent Systems; Advanced Scientific Computing; Networking and Distributed Computing; and Ultra-Large-Scale Integrated Systems.¹⁷ Initially,

Table 3.1 The organizational formation of CISE in 1986.

	SC	NRI	FOCUS CS	IT	SYSTEMS
January 1, 1986 Before CISE	OASC (OD)	OASC (OD)	DCR (MPS)	DIST (BBS)	Multiple units ¹⁸ (ENG)
July 1, 1986 Announcement	DASC	DASC	DCCR	DIST	DCIE
January 1, 1987 After refinement	DASC	DNCRI	DCCR	DIRIS	DMIPS

Legend: SC=supercomputing, NRI=networking research & infrastructure, CS=computer science, IT=information technology, SYSTEMS=experimental systems, D=Division, O=Office. (OD=Office of Director; for the rest, see “Abbreviations and Acronyms” section in the end matter of this book.)

the programs were in divisions similar to their previous home directorates. Soon enough, the CISE leadership (Bell, Brownstein, Adrion, Daen, and the division directors) began to refine the internal organization.

There were some internal concerns, as might be expected. Kent Curtis did not want to lose any programs in the Division of Computer and Computation Research (CCR), and more vociferous resistance came from John Connolly over losing the networking research program.¹⁹ In December, the Division of Advanced Scientific Computing (ASC) shed its networking research program to create a fifth division in CISE, the Networking and Communications Research and Infrastructure Division (NCRI). Steve Wolff, who had been the program director of networking in ASC, became DD/NCRI. Bell believed that NSFNET (just starting in 1986) should serve the entire scientific community, not just supercomputer users. This was Connolly’s fundamental objection, and he soon left NSF. Throughout the formation of CISE, Bloch completely supported Bell’s actions.

There was also a certain amount of pushback from the directorates losing programs (and thus budget). Bell was not able to convince the Directorate of Math and Physical Sciences (MPS) and the Directorate for Engineering (ENG), as well as the Office of Advanced Scientific Computing (OASC), of the logic of creating a “broad” CISE; however, Bloch’s strong support settled the issue, and Bell’s clear vision and operational plan prevailed.²⁰ This style of top-down, rational management of programmatic substance was not typical of NSF and challenged the dominance of physics, which had reigned since the earliest days of the Foundation.

Organizationally, the first year of CISE was largely consumed with turning a collection of diverse divisions and programs into a coherent directorate. The hope

was for most programs and people to share more with others in CISE than with programs or people in other directorates—not an easy task. Following Bell’s clear vision and Bloch’s clear authority, Adrion and Brownstein, with the able assistance of Daen and others, developed the details. Programmatically, the main issue was what topics would be covered by which programs. Procedurally, as a new organization within NSF, the rules, guidelines, and operational processes had to be defined (there were no standard NSF templates). Operationally, the routine proposal evaluation and grant-making activities of the component programs continued.

3.3 Initial Actions

Years later, Bell recalled a Balkanization in which users of particular computer centers were being tied to that center, rather than forming a true network. By removing networking research from the supercomputer centers, the program staff could then focus on networking for a broader audience (that Bell could help guide).²¹

In retrospect, it was another visionary move, embedded in an organizational structure. Because of the rapidly growing popularity of NSFNET among scientists and the rapid development of networking technology as a broad service to science and soon the public at large, a true network independent of particular end-nodes soon developed. Its success and advantages, coupled with the success of the underlying technology developed earlier for ARPANET, formed the prototype of today’s Internet. Bell admitted that, at the time they were creating this more general networking concept, they had no idea of what it would evolve into. Only when he first saw a demonstration of the University of Minnesota’s Gopher²² did he understand some of the deeper technical implications of what they had done.²³

Two other significant, NSF-wide events took place in mid- and late 1986; CISE was involved in both and led one. The Science and Technology Centers (STC) program²⁴ was an effort to encourage technological transfer and innovative approaches to interdisciplinary problems, while supporting basic research and education. It was patterned after the recently started Engineering Research Centers (ERCs)²⁵ and embodied Bloch’s vision that all of science could benefit from both the investment of larger sums of money over longer time periods and closer collaboration between investigators in multiple universities and with industry. To this day, STCs are still the “premier” NSF funding program; they are highly competitive and credited with numerous important results.²⁶

The program attracted strong attention in the scientific community in spite of unfounded concerns that the funds devoted to STCs would reduce funding for single investigators. When the first cohort of grantees was announced in 1989, two

of the six winners were CISE-related: a high-performance machine at the Center for Research in Parallel Computing (CRPC) at Rice University, headed by Ken Kennedy; and a theoretical CS center, DIMACS, hosted at Rutgers University and headed by Daniel Gorenstein and Fred Roberts.²⁷

The second major event involved networking infrastructure to support major science efforts in all fields. Led by CISE, NSFNET backbone service went online at the end of 1986, connecting the five NSF Supercomputer Centers (San Diego Supercomputer Center, National Center for Supercomputing Applications at the University of Illinois, Cornell Theory Center, Pittsburgh Supercomputing Center, and the John von Neumann Center at Princeton) plus the National Center for Atmospheric Research (NCAR) in Boulder. The NSF backbone architecture, due primarily to Dennis Jennings and other NSF staff, was a three-tier structure joining regional and campus networks. The NSFNET and its new backbone were operated directly by NSF primarily through cooperative agreements and contracts—an unusual action for NSF.²⁸

Stephen Wolff, the DD for Networking and Communication Research and Infrastructure, cites²⁹ three important consequences of creating NSFNET: empowerment of the regional networks, permission to use NSFNET commercially, and the concept of access points. The first was a matter of necessity because of insufficient NSF funds to pay for everything. The second was a result of what is often called the Boucher Amendment.³⁰ The third recognized the value of access points—a concept that came out of work at the commercial network exchange on the West Coast. By mid-1987, the need for improved campus-level networking had become critical. A solicitation for improving NSFNET capacity was issued with an award made in late November to a consortium of Merit Network Inc., IBM, MCI, and the State of Michigan.³¹

3.4 Changes in CISE Leadership

1987 also began an important transition within CISE from internal leadership to people with closer connections with the computing community.³² Having experienced line management from the start permitted CISE to gain traction quickly. For the first year, most of the leadership positions (DDs and senior staff), with the exception of Gordon Bell, had been filled by permanent NSF employees. While the use of rotators in leadership positions in CISE later became the norm, the existence of long-time, competent leadership for the newborn directorate was beneficial. While rotators may bring fresh insights and knowledge from outside, awareness of the processes, structures, and politics of NSF (as in any large organization) re-

quires time to learn. Further, it often takes several years to see an effort through to fruition, something that rotators may not appreciate. Adrion left NSF at the end of 1986 to assume a university position. Because of deteriorating health, Kent Curtis transitioned out of CCR to be the Senior Science Advisor in OAD/CISE in mid-1987, eventually retiring in mid-September and passing away in December of that year.³³

I joined NSF as a rotator (from the University of California, Irvine) in mid-September 1987 to be DD/CCR, the first line manager brought to CISE from the outside. Over the short period between when I was asked by Bell in early May 1987 to consider the position and when I joined, I had a number of interactions with Bell and others that provided me a good overview of the new strategic and tactical objectives. Due to my 20-year prior relationship with Bell,³⁴ we shared a number of common understandings. Similarly, my 15-year history as an NSF Principal Investigator (PI/Co-PI) meant that I already knew several of the people in CISE and understood some of the most important NSF processes. My pre-existing relationship with Bloch³⁵ and his hands-on management style resulted in an in-depth, personal meeting with him during my first week at NSF.

After I arrived, Bell and Curtis provided me with additional strategic and operational advice. Beyond these interactions, most of my guidance came from Brownstein, Daen, and Jan Gatton (the experienced Administrative Manager in CCR) as well as wisdom and operational advice from the knowledgeable Harry Hedges³⁶ (who headed a largely autonomous section within CCR that handled the Coordinated Experimental Research (CER) program, large proposals such as STCs, and equipment grants). At the end of October 1987, Gordon Bell left NSF to join a company he had founded before coming to NSF.

CCR was descended from the offices, sections, and visions that had served as the primary source of NSF funding for computer scientists for some years.³⁷ The original programmatic structure of CCR reflected these origins, with seven different programs in core computer science (CS) subjects.³⁸ John Hopcroft of Cornell was chair of the CCR Advisory Board (AB) at that time. As with other AB chairs, he willingly devoted additional time to representing the community on specific issues when asked, as well as helping to shape the composition of the AB and other matters.³⁹

One of the activities I undertook almost immediately was outreach to others in the federal government to establish collaborative relations. These included Jack Schwartz, Saul Amarel, Bill Scherlis, and Mark Pullen at DARPA; Ralph Wachter and John O'Hare at ONR; and many others. One principal means for accomplishing this goal was to attend multiple meetings of FCCSET (Federal Coordinating Committee for Science, Engineering, and Technology).⁴⁰ Bell (along with and reinforced by

Bloch) initiated this outreach and guided me initially; Brownstein continued that push after Bell left.

Research support for computer science and closely related fields was occurring in multiple agencies across the federal government, but without much substantive coordination. Similarly, efforts to establish computer science as an academic discipline were recognized in various places in academe but little was being done within the government. Funding was small compared to other disciplines and, in contrast to the 1960s and 70s when DARPA heavily funded a few leading computer science programs (Carnegie, MIT, Stanford, Utah, and a few others) in their efforts to establish a viable discipline, most government support was focused on contributing to the mission of specific agencies.

In late 1987 an effort that Bell had helped initiate and led⁴¹ produced a government-wide strategy focused on maintaining U.S. computing leadership in computing at a time when Japan was seriously challenging the United States with its Fifth Generation Project. This strategy was of high interest to the Department of Defense (DoD) because of the importance of advanced computing to defense, as well as to other agencies who recognized the importance of high-performance computing (HPC) to their missions.

In early 1989,⁴² the President's Science Advisor appointed a task force to produce an implementation plan of the strategy. Bloch asked me (as a representative of computer science research), Wolff (as a representative of advanced networking), and Mel Ciment (as a representative of computational science and engineering and related fields) to represent NSF on the task force. The group was led by David Nelson from the Department of Energy (DoE) Office of Science and involved representatives from many agencies. After a few general meetings, a plan drafting committee led by Nelson was formed. In turn, an even smaller writing group led principally by Steve Squires and Bill Scherlis from DARPA, and including me and Ciment, pounded out the details and drafted the wording of the Plan. After revisions and editing, this resulted in a public report, entitled "The Federal High Performance Computing Plan," which was sent to Congress on September 8, 1989, by D. Allan Bromley, the Director of Office of Science and Technology Policy (OSTP).

The report called for additional appropriations over current spending of \$150 million in Year 1, rising to \$600 million in Year 5. Increased expenditures were called for in four areas: High Performance Computing Systems, Advanced Software Technology and Algorithms, a National Research and Education Network, and Basic Research and Education (an area of special interest to Ciment and me). The last of these four components was intended to receive more than 20% of the total, based on a percentage of the amounts in other categories. NSF efforts in networking were also key to support requested in that category.

This plan ultimately provided the basis for Congressional efforts led by Sen. Al Gore (D-TN) that resulted in the HPCC (High Performance Computing and Communications) Act and the establishment of the National Coordination Office (NCO) to coordinate efforts across agencies and report annually to Congress. The detailed story of those actions after 1989 is beyond the scope of this chapter, but the origins show the direct impact of NSF actions both organizationally (through Bell's efforts in initiating the strategy effort in 1986–1987) and substantively in the key roles we (Ciment, Freeman, and Wolff) played in formulating the Implementation Plan sent to Congress.

A more direct way in which the early CISE leadership helped shape the field was Bell's efforts in the area of computational parallelism. As a pioneering computer architect, he often saw the future direction of technical development sooner and more clearly than almost anyone else. He understood well in the mid-1980s that parallel computing systems were the future in many types of computing systems, from laptops to massively parallel high-performance machines consisting of millions of individual processing units.⁴³ Although not a software specialist, he understood that without algorithms and supporting software that would easily permit the utilization of parallel machines, much of their power would be wasted.

Apparently stemming from his interactions with the CISE Advisory Committee, including the widely respected Donald Knuth,⁴⁴ Bell in 1987 created the Gordon Bell Prize using his own personal funds.⁴⁵ It was intended to spur the development of the algorithms and software to facilitate the practical use of parallel computers. Initially offered for 10 years, it remains key to advancing and benchmarking high-performance computing over 30 years later.⁴⁶

Yet another way in which Bell left a lasting impression on CISE was his insistence that networking needed to be a major, top-level concern of CISE. In 1986, when CISE was formed, networking research was not widely recognized by the CS community, even though Theorynet and CSNet had been supported since the early 1980s.⁴⁷ As research projects, they produced communication tools for CS researchers. As their usage grew rapidly, the demand for these communication tools expanded rapidly, with researchers from other fields wanting to use them. This resulted in the formation of NSFNET while CISE was being formed.

Even though NSFNET was initially a means of interconnecting the NSF Supercomputer Centers, Bell saw from the usage data that it was being used much more broadly. He understood that robust networking research and operation of high-speed networks was needed to serve all of science and engineering. As a result, he removed networking activity from ASC to form a new Division of Networking Research. Had he not done that, Steve Wolff might not have been able to garner the support, recognition, and freedom of action needed to combine NSFNET and other

IP-based networks to serve a broad community; make sure that networking received a prominent place in the Implementation Plan for the Federal High Performance Computing Program; and eventually oversee the spinoff of the NSF-supported networking operations to form what is today's Internet.⁴⁸

Bell was also involved in other activities that are not so well remembered today, but that nonetheless had great impact. One was the initial co-funding with DARPA for MOSIS, a fast-prototyping service for chips designed by researchers and students.⁴⁹ It was a highly popular and effective service for academic research and education and is still operating today on a self-sustaining basis. It has been key to groundbreaking research as well as the education of generations of students, helping to ensure the leadership of American industry in design of computer-based products.

Another contribution was Bell's insistence that the NSF-supported supercomputer centers develop and use a common version of UNIX, thus unifying the centers' software platform to the benefit of the users.⁵⁰ Another example was his working to reenergize the moribund Computer Science and Telecommunications Board (CSTB) of the National Academies. It became an important voice for the field during the 1990s. These actions had the eventual result of raising the profile of and respect for CS and, more generally, for computing in the federal government and the academy.

Bell has been described as an excellent manager by those who worked closely with him,⁵¹ and he describes himself⁵² as focusing on the big picture and trusting those whom he has identified to work with him as capable of following his lead without his micromanagement.⁵³ His vision, both technologically and organizationally, is amply illustrated by his actions in helping create CISE and forming its initial structure and character. The traditional view at NSF had been, and in many ways continues to be, that what NSF did was entirely driven by the needs of the scientific community. This implied that NSF staff, at all levels, essentially followed the lead of the community.

While in many ways it is true that the best ideas originate in labs and universities, a fundamental difference in computer science and computing more generally is that many of the ideas are generated in the context of creating the technology and using it, which happens largely in industry or wherever the new technology is being used. Bloch and Bell were used to making things happen; they imparted this ethos to the newborn CISE. Whether by design or tacitly, they attracted and hired similar-minded activists, such as Wolff and myself. At the same time, both Bloch and Bell understood the importance of peer review and careful research (predominantly found in universities) to the eventual maturation of the field and the production of future generations.

That activist outlook has continued in CISE to the present, modulated by the changes in personnel over time. It manifests itself in CISE not in the direct way that DARPA manages research, but by harvesting and combining the best ideas from the community, applying them with a vision of a research path that will support the best research, and often forming a broad initiative that may extend beyond computing. Bell did this as illustrated above; Wulf later did the same thing in providing early support for the crucial computing component of the Human Genome Project and in developing the idea of “collaboratories.”

3.5 Transition to Routine Operation

Anticipating Bell’s departure, Bloch asked me in the fall of 1987 for my thoughts on a new AD/CISE. When I mentioned Bill Wulf, Bloch indicated he had heard of him but did not know him and asked me to make an initial contact.⁵⁴ As a long-time personal friend, Wulf grilled me on why he should consider the position. In early January 1988, he accepted the position starting that May.⁵⁵

In the interim, Chuck Brownstein was named Acting AD. He took over from Bell and continued the policies and activities that Bell had initiated until Wulf came on board: promoting the emergence of NSFNET as the major, multi-institutional network; splitting networking off from ASC; and appointing Steve Wolff the DD for the new division. Brownstein steered CISE through mandated budget cuts by using a strategy⁵⁶ to convince Bloch to preserve CISE funding; and Brownstein generally represented CISE well both internally and externally, judging by the successful continuation of actions Bell had started.⁵⁷

3.5.1 William A. (Bill) Wulf (1988–1990)

Wulf, with a Ph.D. in computer engineering, had been a long-time faculty member and department chair in computer science at Carnegie Mellon University (CMU). He and his wife, Anita Jones, had also founded a successful software company, Tartan Labs. When Wulf was approached and eventually accepted the NSF offer, he and his wife approached several universities in the Washington area about faculty positions. They eventually joined the University of Virginia (UVA, Wulf’s Ph.D. alma mater). He needed to complete a term of service at UVA and, as a result, his first day at NSF as an Intergovernmental Personnel ACT assignment (IPA) was May 16, 1988. Although well-funded by DARPA, he had never received NSF funding and was not really familiar with the NSF processes. Like Bell (with whom he had worked at CMU and in industry), he understood advanced computing technology in depth and enjoyed the support of Brownstein, Daen, and the in-place division directors. This permitted him to advance the activities already underway and initiate new ones.

NSFNET was expanding rapidly as the technology improved and demand soared. Wulf provided political support that Steve Wolff, whose group was expanding and operating NSFNET, needed in order to expand and eventually commercialize the network. Wulf did this externally by establishing strong relationships with key members of Congress and their staffs, including Senator Al Gore, who latched on to the idea of a network to support many areas of science and technology and became its champion in Congress.⁵⁸ This effort was aided by Wulf's personality and his record both in research and as an entrepreneur and consultant to industry. Those characteristics also permitted him to establish a strong relationship internally with Bloch and with the division directors in CISE who respected his leadership, as well as with the ADs in other directorates.

As Wulf came up to speed on the programs and processes of NSF, he soon focused on the supercomputer centers.⁵⁹ With his in-depth knowledge of computing systems, he was able to recognize fatal flaws (noticed earlier by Gordon Bell) in the equipment at the John von Neumann Center at Princeton—one of the original supercomputer centers—and did not hesitate to shut it down. He became a strong proponent of the centers, to the consternation of some in the CS community who feared that the centers were taking money that otherwise would have gone to CS research. Wulf developed a cogent argument that computational science using the supercomputers was an important application of the computing technology based on CS research (a precursor to later ADs attitudes, especially that of Ruzena Bajcsy).⁶⁰ He encouraged us division directors to make that point in our interactions with the community.

Changing a culture takes a long time, and the CS community often did not understand that it is important to connect research to outcomes when possible. The CS community now understands more broadly that additional support comes to those who can make a positive case to the society at large for the value of their research. Today, the competition internally for NSF funds is more by biology and other areas; externally, the competition for federal funds sometimes uses the mistaken notion by some (including some members of Congress) that Google, Microsoft, Intel, and other major industry giants no longer need federally funded, basic research and can develop whatever they need without it.

In his discussions with the leaders of the other science and engineering disciplines at NSF, as well as with the proponents of a broad range of applications of computing, Wulf came to believe that interactions between computer scientists and those in other disciplines was essential to progress. He coined the term *collaboratory* to represent a “center without walls, in which the nation's researchers can perform their research without regard to physical location, interacting with col-

leagues, accessing instrumentation, sharing data and computational resources, [and] accessing information in digital libraries.”⁶¹ An invitational workshop was held at Rockefeller University in early 1989, co-chaired by the Nobel laureate Joshua Lederberg and Keith Uncapher, a long-time leader of computing activities at the RAND Corporation. The concept has been applied slowly, but as the power of remote interactions via the Internet has expanded, it is increasingly used in a variety of fields.⁶²

In May 1990, Wulf stepped down from his position as AD/CISE and returned to the University of Virginia. Asked why he left at that time, just when NSFNET and the supercomputer centers were starting to take off, he answered, “. . . I decided earlier that I would come for two years because that is what I was asked to do.” It is worth reiterating that his involvement in NSFNET and the supercomputer centers was critical to their eventual success. One year later, he was asked to be the interim president of the National Academy of Engineering (NAE) and later was elected to that position, where he served for a total of 11 years. As the first AD/CISE who joined an already fully formed and functioning directorate, he set an outstanding bar of accomplishment for future ADs in CISE to meet.

Three months after Wulf’s departure, Bloch’s term as Director of NSF ended.⁶³ He had come to NSF six years earlier, describing himself as an agent of change.⁶⁴ As the first, and to date only, NSF Director without a Ph.D. (a fact of which he was proud) he changed the agency forever. First and foremost, he foresaw the need for interdisciplinary research in general and specifically for removing the barriers between science and technology that could be applied to societal needs.⁶⁵ In addition to creating CISE, he oversaw the start and rapid success of the Engineering Research Centers and the Science and Technology Centers, the creation of the forerunner of the Internet, and the broad use of supercomputers; and he encouraged numerous new efforts in all aspects of NSF activities, including education. In 1985, his accomplishments at IBM before coming to NSF garnered him, along with E.O. Evans and Frederick Brooks, the first National Medal of Technology and Innovation.⁶⁶ After NSF, Bloch was engaged in a variety of science and research policy activities, co-founding the Washington Advisory Group in 1996. In 2002, he was awarded the prestigious Vannevar Bush Award.⁶⁷ He died on November 25, 2016, at the age of 91.⁶⁸

When Wulf left NSF in May 1990, the change in AD/CISE was apparently unexpected. Brownstein ended up serving as the Acting AD/CISE for 16 months. This long transition was undoubtedly due to the fact that, after Bloch’s term ended, it was six months before the next director, Walter E. Massey, was sworn in. As Acting AD during that period, Brownstein again continued the activist mode established

by Bell, Bloch, and Wulf; at the same time, he was working on the national scene in helping shepherd the HPCC Initiative of Senator Gore, which was the instantiation of the HPC Strategy and Implementation Plans mentioned above.

By the time Wulf left, the structure of CISE was largely stable. With a few perturbations, it stayed the same until 2003. NSFNET was well on its way to becoming today's Internet, and the support of supercomputing by CISE was a fact of life. The next eight years (mid-1990 to late 1998) saw a succession of three ADs, a new Deputy AD (who was also Acting AD for ten months), and an Executive Officer.

3.5.2 A. Nico Habermann (October 1991 to August 1993)

Wulf set an admirable standard in terms of accomplishment, but he may have also inadvertently set an example that AD service by non-NSF employees brought from the outside, as had long been the case for program directors (PDs), only required two years of service. While that is perfectly legal, and in some respects advantageous in a fast-moving field such as computing, it does not fit well with the realities of the federal budget cycle and the time required to conceive and start a new, major activity. As a result, AD service of less than three years has tended to result in diminished impact on an ongoing operation, where the time it takes from the initiation of a budgetary idea to the start of its implementation is a bare minimum of almost two calendar years. For a major initiative that may be programmed to last for five years or longer, it may take much of the first year of an appropriation to make adjustments to the initiative.

Once the new Director was in place, it was then another seven months until the next AD/CISE was chosen. Professor Nico Habermann⁶⁹ was sworn in on October 1, 1991.⁷⁰ He had come to Carnegie Mellon University (CMU) in 1968, shortly after obtaining his Ph.D. in the Netherlands under the well-known computer scientist Edsger Dijkstra. He was chair of the CMU Department of Computer Science from 1980 to 1988 and was the first dean of their School of Computer Science; he was also a co-founder of CMU's Software Engineering Institute in 1985. He was a serious scholar, in the more formal European model, and later a successful organizational leader at CMU in addition to his research and teaching.⁷¹

As Habermann was taking over at CISE, NSFNET was moving to ever higher-speed connections; and the 102nd United States Congress passed the HPCC Act of 1991,⁷² often referred to as the "Gore Bill" because the bill was created and introduced by Senator Gore. Much of the bill was based on the earlier DARPA/IPTO Strategic Computing Initiative, but by calling for a multi-agency initiative, this reduced the emphasis on IPTO as the lead supporter of advanced computing.⁷³

In December 1991, the HPCC bill was signed into law.⁷⁴ The bill created the President's Information Technology Advisory Committee (PITAC) to provide independent advice. Then early in 1992, the National Science Board (NSB) commissioned a blue-ribbon panel to "investigate the way science will be practiced in the next decade and recommend an appropriate role for NSF. . . ." ⁷⁵ The final report ⁷⁶ issued in August 1993 included five appendices that surveyed the state of HPC. The panel was chaired by Lewis Branscomb and included several noted leaders including Neal Lane. Habermann organized the work of the panel and saw it through to completion (just prior to his untimely death). The report lays out challenges regarding how NSF can best advance all fields of science and engineering with HPC and makes recommendations on policy, implementation, NSF Centers programs, and relationships to state programs.

The NSFNET story that started in the late 1980s continued apace under Habermann. On June 29, 1992, HR 5344 102d Congress, 2d Session, passed the House. It was popularly known as the "Boucher Amendment"⁷⁷ after its author, Rick Boucher (D-VA), who was Chair of the House Science Committee. Wulf, during his time as AD, interacted extensively with his fellow Virginian, Boucher; and by 1992 Wulf, then at the National Academy of Engineering, held an even more prominent science policy position in Washington. This legislation authorized the first commercial use of NSFNET, which later transitioned to the National Research and Education Network (NREN) and then the Internet.⁷⁸ Wulf had established a strong working relationship with Boucher in the late 1980s, briefing him on the importance of NSF activities in networking.

In January 1993 the National Center for Supercomputer Applications (NCSA) released the first versions of "Mosaic for X" developed by Marc Andreessen and Eric Bina;⁷⁹ by September, they had released working versions of Mosaic for three common platforms (X, PC/Windows, and Macintosh).⁸⁰ Both of the original developers were staff members at NCSA, which derived its major support from NSF through CISE. Mosaic's usability and incorporation of multiple networking protocols, along with its liberal licensing arrangements from the University of Illinois, made it an instant hit, rapidly becoming the preferred means for accessing the World Wide Web. Andreessen soon left Illinois for Silicon Valley and helped create a commercial version called Netscape Navigator. This was a signal event in the introduction of the Internet to the world and sparked the frantic development of online applications that continues to this day.

The Internet needed to be managed. In 1993, after an open, competitive process, NSF entered into a five-year cooperative agreement with Network Solutions, Inc.

(NSI)⁸¹ to provide Internet domain registration services for the non-military part of the Internet, primarily composed of the research and education community. This was a key step in making the expanding network publicly available. Later, as Internet growth exploded, the fees shared with CISE grew into the millions; and after external pressure and lawsuits, NSI lost its monopoly on domain registration and NSF no longer received royalties. This was affected by amending the NSF-NSI cooperative agreement in September 1995. Thirty percent of the registration fees were to be set aside in “an interest-bearing account for the preservation and enhancement of the Intellectual Infrastructure of the Internet.” By 1997, the set-aside contained \$30 million and was growing at the rate of several million dollars per month.⁸²

Another external event, the “Encryption Wars,” was much broader than NSF and focused more on other federal agencies. Started in April 1993, CISE’s CS Theory program had a long history of support of encryption research. In April 1993, the Clinton White House announced the “Escrowed Encryption Initiative, a voluntary program to improve security and privacy of telephone communications while meeting the legitimate needs of law enforcement.” The initiative included a chip for encryption (Clipper), to be incorporated into telecommunications equipment, and a scheme under which secret encryption keys would be escrowed with the government. Keys would then be available to law enforcement officers with legal authorization. The National Security Agency (NSA) designed the system and the underlying cryptographic algorithm SKIPJACK, initially classified but later made public. Despite substantial negative comment, ten months later the National Institute of Standards and Technology approved the Escrowed Encryption Standard (EES) as a voluntary federal standard for encryption of voice, fax, and computer information transmitted over circuit-switched telephone systems.⁸³

Habermann unexpectedly passed away on August 8, 1993.⁸⁴ We will never know his intentions about staying more than two years at NSF, but given his chronological age and professional record of significant organizational service, it is likely he would have served longer. Because he was only AD for 22 months and there were a number of activities already underway, he did not initiate any major actions within CISE. He did, however, provide outstanding service as Executive Secretary to the NSB Blue Ribbon Panel on HPC and represented CISE well to the rest of NSF.

With Habermann’s unexpected demise, there was another period of almost a year while a new AD/CISE was recruited. This time, Mel Ciment, an NSF career employee and Deputy AD/CISE at the time, was asked to be Acting AD, again providing effective oversight of CISE until a replacement for Habermann could be found.⁸⁵

3.5.3 Paul R. Young (July 1994 to September 1996)⁸⁶

Paul Young started his career at Purdue, then moved to the University of New Mexico and later to the University of Washington, where he was a professor and chair from 1983 to 1988 and then Associate Dean of Engineering. As a member of the Computing Research Association (CRA) Board of Directors from 1983 to 1991 and its chair from 1989 to 1991, he understood the importance of the AD/CISE position to the CS research community and stepped up to fill the shoes of Habermann. He was sworn in on July 1, 1994, and served until September 15, 1996. After returning to the University of Washington, he later moved to the University of Wisconsin-Madison.

Early in Young's tenure as AD, a large grant was made to the Stanford Integrated Digital Library Project to investigate multimedia, online libraries of information.⁸⁷ Although the impact of this project was not visible until several years later, two students working on this project developed the algorithms and systems that they later commercialized as the Google search engine: easily one of the largest commercial success stories of CISE-funded research.⁸⁸ The history of this project provides an excellent example of how fundamental research investigations (in this case, how to organize online information) may ultimately have major and unexpected practical results.⁸⁹ A related story⁹⁰ from a computer science theory researcher is how his work with a graduate student 15 years earlier, supported by an NSF grant on abstract formal properties of algorithms, resulted in a published paper long before Google; it was scientifically valid but of no great import at the time. It was found later in a literature search by development engineers at Google and became a key part in some highly valuable mechanisms for allocating advertisement space.

A major internal action Young undertook in September 1995 was to empanel three broad programmatic reviews of CISE, with the assistance of the CISE Advisory Committee. The membership of these panels included many senior people from the field, including several past and future ADs.⁹¹ The first of the panels reviewed CISE programs in Computer Systems and recommended that CISE should place increased emphasis on heterogeneous, distributed systems; scalability; application-level fault tolerance; composable, predictable performance measures; and embedded systems. The committee suggested that the theory and programming semantics community may have become disconnected from mainstream issues and proposed that NSF encourage self-assessment efforts. The panel suggested that NSF program officers should consider sunsetting some ongoing activities to enable investments in new areas. It also encouraged funding of workshops and review boards, and some members strongly encouraged increasing cross-disciplinary grants. Individual program reviews can be found in the panel report.⁹²

The second panel reviewed CISE programs in Human-Centered Systems. It recommended increased emphasis in general on new areas, but three in particular: Electronic Communities, Improved Devices for Human-Computer Interaction, and Digital Libraries and Electronic Commerce. Panel priorities included distributed computing and collaborative technologies; speech, language, graphics and other interactive modalities; the signal-symbol problem; databases; models of intelligence and knowledge-based systems; and autonomous robots. Other issues were also addressed: grant size (too small); unrefereed grants (be bolder); industrial involvement (continue); budget balance between programs and divisions (more involvement of division directors and front office in assessing quality of funded/unfunded proposals).⁹³

The third review panel considered CISE programs in Networking, Communications, and the Convergence of Computing and Communications. It observed that “more cross-fertilization is needed within CISE,” drawing particular attention to the compartmentalization of communications research. It raised several concerns including: peer inertia (subfields tended to perpetuate themselves); obsolete views of technology (communities lacked a contemporary view of technology trends); architecture gaps (insufficient attention to middleware, systems services, and operating system kernels); isolated research programs (insufficient interdisciplinary research); critical infrastructure gaps (software infrastructure and experimental platforms); ineffective communication (program officers did not always communicate and coordinate well); grant sizes (mostly single investigator and large multi-investigator grants without adequate support for small teams; shrinking grant size). The panel recommended promoting team grants, “venture funding” (e.g., SGIRs),⁹⁴ strategic program statements, better internal communication, and better use of technology for managing grants and communicating with the community.⁹⁵

In December 1995, the CISE Advisory Committee formed the CISE Organizational Review Committee (CORC) to assess the reports and the evolution over the next five years of CISE. In addition to the three panel reports, they reviewed the Hayes Report⁹⁶ on the supercomputer centers, an internal “work-flow” study, COVs⁹⁷ and GPRA⁹⁸ indicators, and external reports. CORC found that CISE reflected the original five initiatives defined by Gordon Bell; it was instrumental in the HPCCI; it was the appropriate home for the NSF-wide infrastructure programs because of the interplay with research; it had grown significantly and must exert more cross-agency leadership, manage an increasing load, and work within budget constraints and other agency programs; it needed to be flexible about its organizational structure; CISE management needed to address morale issues and

stress, the hierarchical structure, and recruiting; and CISE should improve its own use of information technology. The CORC recommendations included: maintain CISE as a directorate; consider fewer divisions and offices (five instead of the current six); encourage cross-disciplinary activities and closer interaction between infrastructure and research; employ a team-oriented and flexible administrative structure; and experiment with other technology tools beyond FastLane to manage programs.⁹⁹ The consequences of these recommendations occurred after Young left NSF.

The most significant external event while Young was AD was the spinoff of NSFNET to become a private entity.¹⁰⁰ By this time, the course of events had been well set and agreed upon, so the type of political coverage provided by Wulf at an earlier stage was not needed. A related event occurred in August 1996, when NSF recommended the first set of 13 awards for innovative high-performance network connections. Awards were given to consortia including Illinois-Chicago, Northwestern, University of Chicago, and Carnegie Mellon; Colorado-Boulder, NCAR and the Pittsburgh Supercomputing Center; Virginia Tech and the Virginia Broadband Education Network (VBEN); and Minnesota.¹⁰¹

The reviews that Young carried out did inform actions of the next AD almost immediately. For the first time since CISE was started, there was no gap between ADs.

3.5.4 Juris Hartmanis (September 1996 to November 1998)

Juris Hartmanis is a highly honored co-founder of computational complexity theory and of the computer science department at Cornell, where he was the first chair and professor for many years (now emeritus). As the winner of the 1993 Turing Award (along with Richard E. Stearns)¹⁰² and a member of the National Academy of Engineering (1989), he was widely known and respected in the computer science community; as a result, he brought significant experience, recognition, and respect to CISE. In 2013, he was elected to the National Academy of Science, a rare double accolade.

Wulf, at that time president of the National Academy of Engineering and chairing the search committee for the next AD, approached Hartmanis about being AD/CISE. After visiting NSF to discuss the position, Hartmanis accepted. When asked why he accepted, he replied, “My feeling was that NSF had supported me through the course of my research career, and that I in some sense would like to repay in some minor way, by accepting the invitation to come and work at NSF.”¹⁰³ He also served only slightly more than two years. He indicated that he had been

quite willing and interested in staying longer, but family issues had required him to return home.

Nonetheless, Hartmanis made considerable contributions during his short tenure. His stature in the field and broad knowledge of computer science permitted him to raise the bar of focus on quality and significant results in the funding decisions CISE made, perhaps for the first time since CISE had been formed. He did this through careful, sustained, personal focus on what each of the programs was achieving, not just on their funding. He interviewed all but one of the program directors in CISE at that time (only the person running the theory program was never available, for some reason!), and this direct and low-key leadership set a tone. While results from the grants made during his time could not be assessed until after he left, at a minimum he raised the quality bar for everyone, a tradition that continued after he left.

This focus on quality results carried over into his interactions with the other parts of NSF, as well as with other parts of the government—notably Congress, the Office of Management and Budget (OMB), and key players such as Tom Kalil, a staff member on the Council of Economic Advisors at the time, i.e., those who were in a position to make things happen in the executive branch. It was during Hartmanis's tenure that the proposals for significantly increased funding for CISE became a reality. While there were other factors involved, his leadership clearly was one of them. It is notable that Hartmanis did not have a personal agenda beyond doing the right thing for the field as a whole.

Hartmanis's perspective was informed by his early 1990s leadership of the National Research Council study on the future of computing. That study recommended that Computer Science and Engineering (CS&E) continue to support fundamental research on its intellectual foundations while “looking outward as well as inward” and “encourag[ing] greater interactions between research (especially theoretical research) and computing practice.”¹⁰⁴ Hartmanis's personal stature and whole hearted support of the study's results provided enhanced understanding and support by other parts of NSF and Congress.

Hartmanis's broad view and desire to do the right thing for the field most likely also led him to put into place the first significant reorganization of CISE since 1986; it certainly informed the changes he made. The studies of organizational changes that Paul Young had initiated were read by Hartmanis¹⁰⁵ and perhaps served as a basis for the individual interviews he did with the CISE program directors as well as for his reorganization.

The Hartmanis reorganization can be described at a high level as the renaming or removal of some divisions, the removal or downplaying of hardware-focused

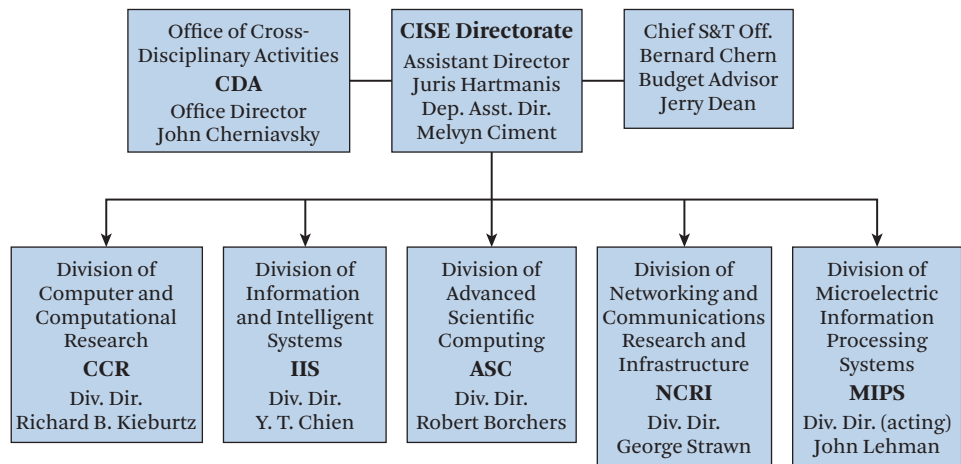


Figure 3.1 CISE organization when Hartmanis began as AD/CISE.

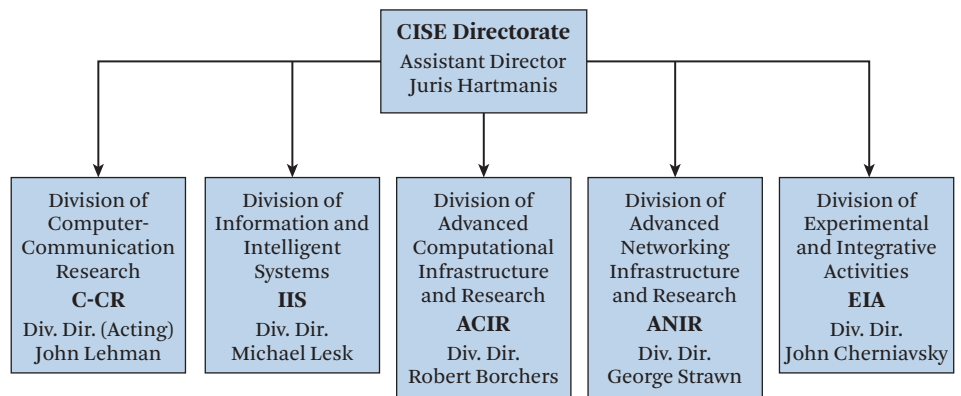


Figure 3.2 CISE organization when Hartmanis ended his term.

programs in others (e.g., robotics), and the strengthening of core research. This can be seen in Figures 3.1 and 3.2.

By the time Hartmanis returned to Cornell, the budget of CISE was being significantly expanded,¹⁰⁶ the level of attention to results had been raised, and the structure of CISE was more in tune with what was happening broadly in computing. It is appropriate to close this chapter on the early years of CISE by noting that the directorate was now ready to take its place as a major player within NSF and on the national stage, at the same time that the public at large was awakening to the many uses of computing.

3.6 Analysis

The early years of CISE could be viewed by analogy to when a young person leaves home to enter college and experiences a number of short-lived activities before graduating into a more adult world of longer, sustained activities. Similarly, there had been a longer, earlier period during which support of computing and computer science in particular was just beginning at NSF. During the period covered here (1986 to 1998), there was a succession of short-term ADs and Acting ADs, who for the most part maintained the status quo. This limited the number of new activities undertaken during the period, especially since there was an Acting AD nearly one-third of the time. (To extend the college student analogy, the initial choice of a major limits the courses one takes.) There were also four NSF Directors and two periods of an Acting Director serving part or all of their normal six-year terms during this period, further dampening new activities.

By no means am I suggesting that the activities of CISE, or any part of NSF, are solely determined by the executives—including division directors. The program directors and staff are the ones who carry on the day-to-day work to which most people in the field pay attention. On the other hand, the executives set overall directions, initiate new programs and cross-cutting initiatives, and represent NSF to the rest of the federal government and the public at large.

CISE came into existence with many programs already in place and functioning well, along with four major objectives in place by the middle of 1987:

1. Develop computer science in the broad sense (i.e., computing) as a field of study.
2. Pursue the development of networking as far as practical for a funding agency.
3. Develop and provide high-performance computing resources for all of science.
4. Bring in more leadership for limited terms from active positions in the field.

With the exceptions of Wulf's introducing the idea of collaboratories and Hartmanis undertaking a reorganization (based in part on analysis done under Young), CISE operated well with a succession of short-term ADs without any additional active management. Several factors played into the history of CISE during this time: the early 1990s saw an economic downturn in the country, meaning that budgets were tight; there was a succession of ADs and Directors; the computing field was still young and relatively small compared to what it would be by the end of the decade; Brownstein (as Acting AD) and Wulf were comfortable carrying on the di-

rections established by Bloch and Bell; and the major “new” activities (networking operations and high-performance computing) had lives and constituencies of their own, somewhat separate from the usual academic research.

None of the above should be construed, however, as implying that the ADs at the time sat around reading journals or leaving the office early! The second and third major objectives named above required considerable political cover and direct management review, which the ADs had to supply. Simply providing the routine oversight of the divisions and making sure they have good leadership has always taken a significant portion of an AD’s time. In addition, an AD has duties both vertical (leading and managing CISE) and horizontal (participating in the general leadership cohort of NSF and interacting with others outside).¹⁰⁷

Considering the four directions listed above, by 1997 considerable progress had been made on all of them; the first two were supported and led by CISE actions to a large extent, while the second two were fundamentally internal issues:

1. Computer science as a valid and accepted academic discipline had grown significantly; as one example, in 1986 there were 111 accredited Ph.D. CS programs in the U.S.,¹⁰⁸ while in 1996 there were over 130 programs.¹⁰⁹
2. The Internet was born in 1995 when NSFNET was commercialized and NSF stepped away from operational control;¹¹⁰ while there were a number of critical developments, the activities of CISE and its immediate predecessors were central to the early application of the TCP/IP protocols, development of networks open to all science researchers, indirect support of many related aspects (including Mosaic and the work that led to Google), and the foresight to open the networks up and permit community control of the underlying technology.
3. The activity of providing HPC resources to the science research community nationally was and remains one that ultimately belongs to individual organizations, including NSF. Within NSF, CISE was born with this responsibility when the former Office of Computing Activities, located in the Office of the Director, was moved under the control of CISE. That piece of the story is complex and is described elsewhere in this study.¹¹¹ CISE tried to do a responsible job in spite of many external pressures.
4. Bloch was really the one that initiated the focus on bringing in new leadership on a regular basis (across the Foundation), starting with Bell; it has continued to this day with every AD, most of the DDs, and a substantial number of the PDs in CISE coming from outside.

Risking some generalizations, which might be overdrawn, about the first 12 years of CISE:

- It had significant impact on the world through its actions and some of its research.
- It developed reasonable internal coherence and cross-divisional cooperation.
- It went from what had initially been viewed as a collection of fringe or service activities of only incidental importance to NSF, to one that was growing in importance in research to other disciplines as well as to the outside world.

Asking what drove the decisions made and directions taken by CISE in this period, I again will venture some generalizations:

- The vision of a few people (Bloch, Bell, Brownstein, Adrion, and Wolff) of what computing is (or should be), the opportunities that it afforded, and a realistic view of the world outside provided the spark and the direction.
- Important technical ideas clearly came from the community, initially as embedded in the experiences and knowledge of the decision-makers, but perhaps more in their ability to understand new developments in the field and then go beyond them to fashion new visions.
- All of these leaders not only had the ability to envision futures that were not yet concrete, but all were action-oriented.
- They were not averse to taking risks with their decisions.
- Their decisions, while subject to much discussion and some modification internally, were largely unaffected by anything external to NSF other than budgets and the mission of NSF, as broadly interpreted.
- The small size of CISE and the relative insignificance of computing in the early years meant that what the CISE Directorate did was below the radar of most outside influences that could have thwarted its development.

By the end of 1998, many things had changed—external conditions, personnel, the importance of computing, the nature of research in many areas of computing, and even NSF itself. Again, we come to a transition in the history of CISE.

Notes

1. I have used the first person because I was at NSF from late 1987 to late 1989 as division director of CCR. During the remainder of the period covered here, I was a board member of

the Computing Research Association (CRA) and sometimes directly aware of major actions at NSF because of frequent briefings to the CRA Board by the ADs and others. I was not involved in the key decisions and actions.

2. See the list of acronyms and abbreviations in the end matter of this book.
3. Short biographies of many of the people named in this chapter can be found in a “Biographies” section in the end matter of this book.
4. Oral history, Gordon Bell, interviewed by William Aspray, July 14, 2017. Charles Babbage Institute.
5. Erich Bloch. October 2016. Personal communication to Peter Freeman.
6. “Rotator” as used at NSF is a generic term that may refer to one of several varieties of limited term employees, including IPAs. IPA refers to the Intergovernmental Personnel Act of 1974 that authorizes the federal government to exchange personnel with non-profit and other governmental organizations for a limited period of time. It is a program that has been used extensively by NSF to attract working scientists and others, primarily from universities, to fill scientific and leadership posts as a way of helping NSF stay at the forefront of science and education efforts in the country.
7. NSF is primarily structured into directorates, with divisions within directorates. Offices usually report to the Office of the Director (e.g., public relations). The “Assistant Director (AD) of CISE” refers to the head of CISE. The position reports directly to the Director of NSF and thus has horizontal duties such as participating with the Director in discussions that affect NSF-wide policies or decisions, as well as the vertical responsibility for the operation of CISE (and similarly for other directorates). A more precise title would be “Assistant Director for CISE.”
8. See Chapter 2 in this volume for additional detail on this period.
9. See Chapter 1 in this volume for additional detail.
10. There was resistance from the Engineering DDs, particularly from Electrical, Computer, & Systems Engineering (ECSE) Director Frank Huband, but also from Science Base Development in Design, Manufacturing, and Computer Engineering Director Bernie Chern, and Nam Suh, AD/Engineering. John Connolly, Director of OCA, was not anxious to be “downgraded” to a directorate-level division. Kent Curtis was helpful, but concerned that his division might be diluted. (Personal recollection of Rick Adrion.)
11. At the time, IBM had a strong market position in computing technology, driven partly by Bloch’s technical efforts.
12. The 1980s were also a critical time in the development of computer science as a defined discipline: numerous Ph.D. programs at top universities began to be accepted and expand (often with NSF support), fundamental research outcomes were resulting in new and major industry developments, and the strategic importance of computer science and computing more broadly was beginning to be understood by policymakers.
13. Another of the enduring legacies of Gordon Bell, which has inspired over 30 years of improvements, was his establishment in 1987 of the Bell Prize for improvement in parallel processing. See “Gordon Bell Prize, three decades: Motivating and measuring

High Performance Computing progress.” Keynote presentation, Supercomputing Frontiers, National Computing Centre of Singapore, March 2017. Charles Babbage Institute.

14. Oral history, Rick Adrion, interviewed by William Aspray, March 14, 2017. Charles Babbage Institute. Oral history, Steve Wolff and Charles Brownstein, interviewed by William Aspray, June 23, 2017. Charles Babbage Institute. See also Bell 2017 interview above, in which Bell noted that his, Bloch’s, and Nam Suh’s interest in expanding the CISE and Engineering portfolios to contain more applied research made a real difference in moving the fields ahead.
15. Oral history, Charles Brownstein, interviewed by William Aspray, June 23, 2017. Charles Babbage Institute.
16. Oral history, Rick Adrion, interviewed by William Aspray, March 14, 2017, and January 2, 2018. Charles Babbage Institute.
17. Erich Bloch. March 3, 1986. “Consolidation of computer-related activities.” Staff memorandum. Charles Babbage Institute.
18. These were the Science Base Development program and the Communications and Signal Processing program.
19. Rick Adrion personal files.
20. Oral history, Brownstein, interviewed by William Aspray, June 23, 2017; Oral history, Adrion, interviewed by William Aspray, March 17, 2017.
21. Again, there is an irony here. While Bell’s broad vision enabled the development of the Internet, its spread in the late 1990s has permitted large companies to sometimes tie users of their equipment and systems to using *only* those items and not similar items from other companies.
22. P. L. Frana. 2004. Before the web there was gopher. *IEEE Annals of the History of Computing*, 26(1): 20–41. DOI: [10.1109/MAHC.2004.1278848](https://doi.org/10.1109/MAHC.2004.1278848); last accessed 23 September 2019.
23. A similar comment was made by Bob Kahn, one of the developers of the TCP/IP protocol, to a small group (of which I was one) almost two decades later. Networks had been built to support computational and communications resource sharing (for bomb resistance, in the case of ARPANET). By the time Bloch arrived at NSF, he (and many others) could see that email and file transfer were going to replace paper. Bloch stopped sending OD (Office of the Director) memos in favor of email, which quickly converted the senior NSF staff. Very few at that time could see that the power of networks was in information sharing and access. Gopher hinted at the possibility; Mosaic (discussed later in this chapter) really made it clear.
24. <https://www.nsf.gov/od/oia/programs/stc/graduated-centers.jsp>; last accessed 23 September 2019.
25. https://www.nsf.gov/funding/pgm_summ.jsp?pims_id=5502; last accessed 23 September 2019.
26. See S. J. Fitzsimmons, O. Grad, and B. Lal. June 1996. *An Evaluation of the NSF Science and Technology Center (STC) Program, Volume I: Summary*, Quantum Research Corporation and National Science Foundation, https://www.nsf.gov/od/oia/programs/stc/old_reports/abt.pdf; last accessed 23 September 2019; S. Mason. December 2010. *The NSF Science and*

- Technology Centers Integrative Partnerships Program, 2000–2009*, Report of the AAAS Blue Ribbon Panel, https://mcmprodaaas.s3.amazonaws.com/s3fs-public/STC_BRP_Report.pdf; last accessed 23 September 2019; D. E. Chubin et al. December 2010. *AAAS Review of the NSF Science and Technology Centers Integrative Partnerships (STC) Program, 2000–2009*, Final Report, Charles Babbage Institute, http://www.aaas.org/sites/default/files/reports/stc_aaas_full_report.pdf; last accessed 23 September 2019.
27. Two of the class of 1991 STCs were also CISE-related: The Brown-UNC-Utah Graphics Center and the University of Pennsylvania Cognitive Science Center.
 28. https://www.nsf.gov/news/news_summ.jsp?cntn_id=103050; last accessed 23 September 2019.
 29. Wolff interview, July 20, 2017. Charles Babbage Institute.
 30. https://en.wikipedia.org/wiki/Rick_Boucher; last accessed 23 September 2019.
 31. K. D. Frazer. 1996. *NSFNET: A Partnership for High-Speed Networking: Final Report, 1978–1995*, Merit Network Inc. https://www.merit.edu/wp-content/uploads/2019/06/NSFNET_final-1.pdf; last accessed 23 September 2019.
 32. There are pros and cons to the tradeoff. For example, dedicated permanent employees typically have a wider knowledge across a field, but may have lost touch with the reality of the life of an academic researcher; at the AD level, the advantage may go to an experienced, senior academic administrator, even though they are often farther from direct research activity.
 33. Curtis was honored at a dinner in mid-September 1987 that included many of the senior computer scientists from the community, as well as NSF staff. (Peter Freeman, personal notes.)
 34. Bell and I had back-to-back weekly meetings at CMU with Allen Newell starting around 1967, and because Newell was always running late, Bell and I shared a good bit of time in a small waiting room. Later, as a research faculty member at CMU, I worked with Bell on an ARPA Project and wrote my first published paper with him.
 35. I had been on an IBM external advisory committee reporting to Bloch at the time he joined NSF.
 36. See Appendix C.
 37. For more detail see Adrion interview, 2017, *op. cit.*
 38. Because those programs set much of the tone for core CS funding, it is relevant to note who the PDs were at that time: CS Theory (Carl Smith), Software Systems (Tom Keenan), Computer Architecture (Zeke Zalcstein), Software Engineering (Sanat Basu), Numerical Computation (Bobby Caviness), Instrumentation (Al Thaler), and Coordinated Experimental Research (Harry Hedges, J. Mack Adams); Larry Oliver was a staff associate for special programs and outreach. Carl Smith, Bobby Caviness, and J. Mack Adams were IPA rotators and Sanat Basu was a temporary government employee. The others were regular NSF employees.
 39. This, of course, can be a two-edged sword: On the one side it can be a very useful source of candid information from the field, but on the other side one must always make sure that the information being provided is not unduly biased.

40. FCCSET was a committee by the OSTP (Executive Office for Science and Technology Policy) to ensure various policy initiatives were working together across the federal government. Bell played a key role in revitalizing it. See Brownstein interview, 2017.
41. See Brownstein interview, 2017, *op. cit.*; “A research and development strategy for high performance computing,” Executive Office of the President, Office of Science and Technology Policy, November 20, 1987; last accessed 23 September 2019.
42. Freeman, personal notes, book VIII.
43. As an example, in the late 1960s Bell was asked by DARPA to design the most powerful computer he could to support AI calculations. His design was a multi-processor system connected by an unusual (for the time) cross-bar communication system; the system was called C.ai, but it was never built. As a post-doc at CMU, I led the software design for C.ai, working with him. It was later followed by a similar architecture, C.mm, designed by Bill Wulf (second AD/CISE). (Personal recollection by Peter Freeman.)
44. Bell interview, 2017, *op. cit.*
45. See <https://awards.acm.org/bell/> and https://en.wikipedia.org/wiki/Gordon_Bell_Prize; last accessed 15 May 2018.
46. Gordon Bell Prize, *ibid.*
47. See Chapter 2 in this volume.
48. *NSFNET Final Report*, 1996. Fraser, 1996, *op. cit.*
49. See <https://www.mosis.com>; last accessed 23 September 2019.
50. Brownstein interview, 2017, *op. cit.*
51. Brownstein interview, 2017, *op. cit.*; Freeman personal recollection.
52. Bell interview, 2017, *op. cit.*
53. The truth is probably in the combination of these two observations; Bell was a leader, not a manager in the narrow sense.
54. He could easily have contacted him directly, but as I learned much later in working with Bloch as a close colleague from 2007 to 2012, he never told you everything he knew but often preferred to work through you to achieve his goals.
55. Apparently, I did an acceptable job because in telling me of his acceptance of the position, he said, “It’s all your fault!” [Freeman, personal notes, book VI].
56. Simply stated, the strategy was to say, “If CISE is cut by X%, we will have to cut out A, B, and C: but you know from your experience that those are critical areas for the advancement of the field.”
57. Brownstein interview, 2017, *op. cit.*
58. In a speech at a symposium celebrating the 20th anniversary of the Internet in 2015, Gore noted that it was easy for him to see the possibilities of the Internet because his father, also a senator, had championed the Interstate Highway System in the 1950s, which had greatly improved commerce and many other aspects of American life. [Recollection of Peter Freeman, who was attending the symposium.]
59. Oral history, William Wulf, interviewed by William Aspray, July 28, 2017. Charles Babbage Institute.

60. See Chapter 4.
61. W. Wulf. March 1989. "The national collaboratory." In *Towards a national collaboratory*. Unpublished report of a National Science Foundation invitational workshop, Rockefeller University, New York. Also: W. Wulf. 1993. The collaboratory opportunity. *Science*, 261: 854–855. DOI: [10.1126/science.8346438](https://doi.org/10.1126/science.8346438); last accessed 23 September 2019.
62. <https://en.wikipedia.org/wiki/Collaboratory>; last accessed 15 May 2018.
63. <https://www.nsf.gov/about/history/bios/ebloch.jsp>; last accessed 11 September 2018.
64. "Pushing ivory-tower scientists into the high-tech race." *The New York Times*, February 15, 1987.
65. See the video at <https://www.nsf.gov/about/history/leaders/bloch.jsp>; last accessed 11 September 2018.
66. <https://www.uspto.gov/learning-and-resources/ip-programs-and-awards/national-medal-technology-and-innovation-nmtia>; last accessed 11 September 2018.
67. Bloch, "Biographies."
68. "Erich Bloch, who helped develop IBM mainframe, dies at 91." *New York Times*, December 3, 2018.
69. https://en.wikipedia.org/wiki/Nico_Habermann; last accessed 23 September 2019.
70. One might ask how it was that the first three CISE ADs were from CMU. There was always a strong ethos of public service at CMU, encouraged by Herb Simon, Allen Newell, and Alan Perlis. There was also the "family immigration" effect: I suggested and helped recruit Wulf. He likely suggested Habermann, who was experienced and at a good transition point in his career. Also, in the 1980s other computer science schools were still struggling to make a name and get on the map, and the Director did not routinely use a search committee, as is the case now.
71. Peter Freeman, personal recollection; he served on my Ph.D. thesis committee in 1968–1969.
72. U.S. Congress. Senate. 1991. High Performance Computing Act 1991. 102nd Congress, 2nd session, s.272. <https://www.congress.gov/bill/102nd-congress/senate-bill/272>; last accessed 23 September 2019.
73. A. Norberg and J. O'Neill. 1996. *Transforming Computer Technology: Information Processing for the Pentagon, 1962–1986*. Baltimore: Johns Hopkins University Press.
74. <https://www.gpo.gov/fdsys/pkg/STATUTE-105/pdf/STATUTE-105-Pg1594.pdf>; last accessed 30 August 2018.
75. NSF Blue Ribbon Panel on High Performance Computing. August 1993. *From Desktop to Teraflop: Exploiting the U.S. Lead in High Performance Computing*, Appendix B. <https://www.nsf.gov/nsb/publications/1993/nsb0893.pdf>; last accessed 23 September 2019.
76. NSF Blue Ribbon Panel on High Performance Computing. 1993. Charles Babbage Institute.
77. <https://www.congress.gov/bill/102nd-congress/house-bill/5344/text/eh>; last accessed 11 September 2018.
78. <http://thomas.loc.gov/cgi-bin/bdquery/z?d102:HR00656:@@P%7C/bss/d102query.html%7C>; last accessed 23 September 2019.

79. <http://1997.webhistory.org/www.lists/www-talk.1993q1/0099.html>; last accessed 19 May 2018.
80. <http://www.ncsa.illinois.edu/enabling/mosaic>; last accessed 23 September 2019. X' was the name of Unix's X Window System.
81. https://en.wikipedia.org/wiki/Network_Solutions; last accessed 11 September 2018.
82. Testimony of Joseph Bordogna before the House Science Committee's Basic Research Subcommittee. September 25, 1997. <https://www.nsf.gov/about/congress/105/jbdomain.jsp>; last accessed 25 February 2018.
83. These "wars" have periodically broken out since, including in 2018, without a lasting resolution.
84. Obituary of A. Nico Habermann. *New York Times*, August 11, 1993. <http://www.nytimes.com/1993/08/11/obituaries/a-n-habermann-62-computer-school-dean.html>; last accessed 22 February 2018.
85. Oral history, Melvyn Ciment, interviewed by William Aspray and Rick Adrion, June 20, 2017. Charles Babbage Institute.
86. We were unable to interview Young for this project, and, with the exception of Rick Adrion, we do not have much contemporaneous information about his actions as AD or afterward.
87. See https://www.nsf.gov/awardsearch/showAward?AWD_ID=9411306. last accessed 23 September 2019. For those not familiar with the details of grant-making by NSF, only the very largest grants made by a program ever come to the attention of the AD, and his/her signature is not required as part of the approval process. Further, ADs are not involved in making the funding decision unless there is some special issue to be adjudicated that cannot be decided between the program and division directors. The final, formal approval for a disbursement of funds is actually made by the NSF Budget Office after legal, administrative, and other regulatory determinations (not involving scientific criteria) are made.
88. See <https://en.wikipedia.org/wiki/PageRank>; last accessed 23 September 2019.
89. See https://en.wikipedia.org/wiki/Google#cite_ref-milestones_11-0; last accessed 23 September 2019.
90. Private communication of Vijay Vazarani to Peter Freeman, ca. 2000.
91. Panel 1—Computer systems: Al Aho, Fran Allen, Forrest Baskett, David Patterson, Larry Snyder, Bob Sproull, Jeannette Wing; Panel 2—Human-Centered Systems: Ruzena Bajcsy, Alan Biermann, Fred Brooks, Hector Garcia-Molina, Jim Hollan, Michael Rabin, Doug van Houweling, David Wohlers; Panel 3—Networking, Communications, Convergence: Dave Clark, Dave Forney, Jim Gray, H.T. Kung, Ed Lazowska, Dave Messerschmitt, Jon Turner; Organizational Review Committee (CORC)—Rick Adrion, Fran Allen, Forest Baskett, Dave Forney, Sue Graham, Sid Karin, Ken Kennedy, H.T. Kung, Barbara Liskov, Michael Rabin, Raj Reddy, Bill Wulf.
92. CISE Computing Systems Review, draft internal report to CISE and the CISE Advisory Committee. November 9, 1995. Charles Babbage Institute.
93. CISE Research Review on Human Centered Systems, draft internal report to CISE and the CISE Advisory Committee. October 22, 1995. Charles Babbage Institute.
94. SGIR = Small Grant for Innovative Research.

95. CISE Research Review on Networking, Communications, and the Convergence of Computing and Communications, draft internal report to CISE and the CISE Advisory Committee. October 23, 1995. Charles Babbage Institute.
96. <https://www.nsf.gov/pubs/1996/nsf9646/nsf9646.htm>; last accessed 23 September 2019.
97. COV = Committee of Visitors, empaneled periodically to provide external review for the operations and results of individual divisions and/or programs.
98. The Government Performance and Results Act of 1993 (GPRA) (Pub.L. 103-62) is a United States Law enacted in 1993, and is one of a series of laws designed to improve government performance management. See https://en.wikipedia.org/wiki/Government_Performance_and_Results_Act; last accessed 23 September 2019.
99. CISE Reorganization Report, internal report to CISE and the CISE Advisory Committee, cover memorandum from Rick Adrion. August 8, 1996. Charles Babbage Institute. FastLane is the NSF's official website for application submission and grant management. See T. Misa and J. Yost. 2016. *FastLane: Managing Science in the Internet World*. Baltimore: Johns Hopkins University Press.
100. *NSFNET: Final Report, 1996*.
101. https://www.nsf.gov/news/news_summ.jsp?cntn_id=101779; last accessed 23 September 2019.
102. Oral history, Juris Hartmanis, interviewed by William Aspray, July 26, 2009. ACM Oral History Interview No. 19. <https://dl.acm.org/citation.cfm?id=1775727>; last accessed 23 September 2019. This extensive interview not only covers many details of Hartmanis's professional career, but has a fascinating, first-person account of his early life in Central Europe before, during, and immediately after the Second World War.
103. Oral history, Juris Hartmanis, interviewed by Rick Adrion, October 24, 2017. Charles Babbage Institute.
104. National Research Council. 1992. *A Broader Agenda for Computer Science and Engineering*. National Academy Press.
105. Personal recollection of Rick Adrion.
106. Primarily as a result of the Information Technology Research (ITR) program that was just being created. See Chapter 8 in this volume.
107. See Chapter 12 for additional comments on the duties of an AD/CISE.
108. D. Gries and D. Marsh. 1988. The 1986–1987 Taulbee survey report, in *The Computing Research Board's Survey on the Production and Employment of Ph.D.s and Faculty in Computer Science*. New York: Cornell University. <https://hdl.handle.net/1813/6737>; last accessed 12 November 2018.
109. D. Kozen and S. Zweben. March 1998. 1996–1997 CRA Taulbee survey, in *Computer Research News*. <http://archive.cra.org/statistics/survey/97/97.pdf>; last accessed 14 November 2018.
110. There are many descriptions of the history of the Internet, but J. Abbate, *Inventing the Internet*, MIT Press, 1999, and “Brief history of the Internet,” Internet Society, are good places to start. <https://www.internetsociety.org/internet/history-internet/brief-history-internet/>; last accessed 15 November 2018.
111. See Chapter 10 in this volume.

1999–2006: Broadening Computer Science with New Initiatives

Peter A. Freeman¹

A remarkable confluence of technical progress, leadership, and socioeconomic factors in the late 1990s took the Computer and Information Science and Engineering Directorate (CISE) to new levels of importance within the National Science Foundation (NSF), the government as a whole, and the nation. This fortunate occurrence was composed of remarkable technical progress and societal changes fueled in large part by previous research supported by CISE, farsighted national political leadership, a new NSF Director who understood the significance of extraordinary advances in computing technology, and the choice of two consecutive Assistant Directors (ADs) willing and able to serve longer terms than previous CISE/ADs.

This confluence enabled greatly expanded funding for CISE, the longer AD terms (33 and 57 months) permitted the initiation and oversight of seminal activities on a range of subjects, and the political events of the period further strengthened CISE's expanded efforts. The specific research results of this expansion are only beginning to unfold, but the impacts on computing as an activity and on science more generally have already been pervasive.

It is useful to remember the general technological and political context. The late 1990s was a time of rapid development, largely in the realm of computer power and connectivity, rarely seen before in any area of science or technology. Little more than a decade earlier, personal computers were only beginning to be widely known and individually owned, remote access to a computer over telephone lines was barely past the novelty stage, and what one could do with the software on personal computers was still primitive. The exponential growth of computing power that

had been predicted by Gordon Moore² in 1965 continued unabated, so that novel applications were suddenly realizable.

At the same time, the national political leadership was uncommonly attuned to the future possibilities of technological developments and willing to invest in the basic R&D necessary to attain them. At the same time, economic policies produced the federal budget surpluses that permitted substantial new investments.

Basic physical science R&D—supported largely by the Department of Energy, NSF, and the Defense Department—provided the foundation for continued advances in computing and communications technology. Research and other activities supported by CISE had resulted in or enabled many system- and application-level developments, including highly parallel machine clusters, usable browsers, the Google search engine, high-speed optical switching, and the Internet, among others. After many years of R&D in a range of disciplines, modern computing resulted in economically useful results to a broad segment of society. In the mid-1990s an explosion of online, remotely accessible applications over the World Wide Web appeared. Known as the “dot-com boom” (before it became the “dot-com bust” around 2001), this rapid and broad expansion in the use of computing was fueled by the usual human desire to cash in on quick riches and by a venture capital system hungry for profits.

In August 1998, Rita R. Colwell³ became the eleventh Director of NSF (and the first female director). A widely respected and honored microbiologist, she described in our recent interview her understanding of the importance of computation:⁴

. . . I was totally sympathetic and understood fully the potential that computation would have [from that time]. . . So that predisposed me to being very enthusiastic and supportive of the initiative[s]⁵ that we carried out [in computing] while I was Director. Very important initiatives. I’m very proud of what we accomplished.

In another fortuitous development, Ruzena Bajcsy became the AD/CISE just months after, in December 1998. She also was an experienced and senior researcher in her field (AI, robotics, and related areas). She and Rita Colwell became instant friends, and for a variety of reasons⁶ they quickly became a mutually supportive team. In early 2002, I was chosen to be the next AD/CISE, after Bajcsy stepped down somewhat unexpectedly in the summer of 2001. More on this transition below.

Prior to either Colwell or Bajcsy coming to NSF in 1998, efforts had been underway in the government,⁷ closely supported by senior leadership in the computing community, to provide a massive injection of new funds into NSF over a five-year

period.⁸ Colwell, Bajcsy, and then I utilized the funds in ways that benefited the CISE community, ushering in an eight-year period in which the importance and leadership of CISE grew in ways unimaginable in 1986, including the range of research supported, size of budget, respect from other parts of NSF, and importance to Congress and the executive branch.

On top of this convergence of favorable circumstances, the tragic events of September 11, 2001, ushered in a period in which the importance of computing became not just a commercial opportunity, but also a matter of national priority because of its importance to defense, intelligence operations, and economic prosperity. Thus began the middle years in this history of CISE.

4.1 Ruzena Bajcsy (December 1998 to August 2001)

Ruzena Bajcsy continued the tradition of CISE ADs being senior professionals, with strong and recognized research careers. She held two Ph.D.s (electrical engineering from her native Slovakia and computer science from Stanford under John McCarthy) and was a dedicated scientist. Her research for over 30 years, starting when she was a Ph.D. student at Stanford and then as a faculty member (1972–2001) at the University of Pennsylvania, had been focused on robotics, especially problems of perception and understanding; she had also served as chair of the Computer and Information Science Department at Penn.

She was born the year Hitler was elected Chancellor of Germany and escaped the Holocaust only with her sister. She grew up in a Catholic orphanage (to hide her Jewish heritage) and was educated in Bratislava, first under the Nazis and then under the Communists, earning a Ph.D. in electrical engineering from the Slovakia Technical University in 1947. As she noted in an interview:⁹

In 1948 Czechoslovakia became Communist, and I experienced the Stalin dictatorship and oppression again. All these experiences made me who I am today. They taught me that while I love mathematics, engineering, and science in general, the most important thing in life is caring for people. I made it my life mission to create the best possible environments I know how for colleagues and students to do the best science they can.

Her European heritage, focus on high-quality performance, love of science in general, and dislike of organizational matters lent her a demeanor somewhat different from her predecessors as AD/CISE.

She accomplished a great deal in slightly less than three years, ably managed a large increase in the CISE budget, raised the reputation within NSF of computer

science (as a serious science and not just as a technology), and served as an effective ambassador for the field to Congress and in the broader community. With her dedication to science, hard work, willingness to cooperate with other scientific communities and directorates, forthright manner, and personal maturity, she arguably accomplished more to raise the stature of the CISE Directorate than any previous AD/CISE.

A comment she made while being interviewed for this project illustrates this.¹⁰ In a discussion between the ADs and the Director on how to divide a budget increase, each was asked how they would spend an extra million dollars. When Bajesy's turn came to speak after hearing what important problems could be solved with the extra money in each field, she said: "Well, if I don't get support for my community, you folks will not be able to do any of that."

Bajesy came to NSF holding a view common among computer scientists: that the supercomputer centers were draining money from CISE that could otherwise be spent on computer science research.¹¹ As she came to better understand the arguments that other fields were making and refined her own views, other directorates at NSF began to understand that computer science is a true science, grappling with deep questions and not merely programming applications for others. In turn, she came to appreciate and support the importance of the supercomputer centers—not only for their importance to other fields but as drivers for advances in computer science as well.

The supercomputer centers provided extra money to CISE and bolstered its stature in the government, as well as providing the opportunity to make sure that other sciences were involving computer scientists when appropriate. As an indicator of this changing attitude, when the long-time heads of the two largest centers stepped down or moved on they were replaced by card-carrying computer scientists, Dan Reed at NCSA and Fran Berman at SDSC.

While the supercomputer centers have been an important part of CISE for most of its history, Bajesy was the one who proposed an unprecedented NSF-wide panel on high-performance computing. It was approved by the National Science Board and begun under her direction. Employing a rarely used term (*cyberinfrastructure*),¹² she empaneled the NSF Advisory Panel on Cyberinfrastructure (CI) to consider a number of specific questions.¹³ The panel included members from many scientific disciplines and heard testimony from approximately 75 people. The final report¹⁴ of the panel was officially presented to NSF on February 3, 2003.

This report received mixed reviews, ranging from praise for focusing on the importance of high-performance computing (HPC) for science overall to criticism of insufficient emphasis on some areas such as communications and software.

Now (in 2019), the major contribution of the panel's study and the report might be viewed as focusing NSF's institutional efforts on the provisioning of advanced computational resources for the support of advanced research.¹⁵ This is still an issue today because of the rapid advance of the technology, the change in the ways in which computational resources are utilized by science, and the cost of providing these resources.

Bajcsy's stewardship of resources for cyberinfrastructure serving science broadly (typically around 40% of the CISE budget at that time) is arguably one of the two most influential actions she took. The other action, which has been much more influential on computer science specifically, was her stewardship of the start of the expanded funding mandated by President Clinton.

Soon after Bajcsy joined NSF, President Clinton requested an additional \$150 million for "Information Technology Research (ITR)."¹⁶ This request built on the earlier IT Research Initiative for the Twenty-first Century (IT2).¹⁷ Bajcsy was asked by Rita Colwell to lead the implementation of this initiative for NSF if the money was allocated.¹⁸

To do so, she had to fend off the efforts of other directorates and disciplines that thought they should get the money or large portions of it, most especially the Engineering Directorate (ENG). Most of the funding programs for computer engineers had been moved into CISE barely a decade earlier, and that memory lingered with some of the program directors in ENG; also, the AD/ENG at the time when ITR was being formulated was Eugene Wong,¹⁹ who had been an Associate Director at the Office of Science and Technology Policy (OSTP) in the early 1990s and a faculty member in Electrical Engineering and Computer Science at Berkeley.

Colwell played a key role in supporting Bajcsy. At one point, when Colwell learned that a computer science community representative had not been invited to a meeting of professional society representatives seeking funding from the ITR money, she personally called the Executive Director of the Computing Research Association to invite him to the meeting.²⁰ It is clear from our interviews with both Colwell and Bajcsy that Colwell's level of support for and respect of what Bajcsy was doing was key in ensuring that CISE managed the majority of the ITR. One mechanism for affecting that, besides explicit decisions by Colwell in allocating funds to CISE, was to establish an NSF-wide steering committee chaired by CISE: in the first year of ITR, all panels were managed by CISE; in future years, CISE's role was one more of coordination among the other participating directorates.

ITR was announced in mid-1999²¹ with proposals due in early 2000 (FY 2000). Subsequent solicitations were made thru 2003 (FY 2004).²² Several hundred million dollars were awarded over the lifetime of the program, with the last awards being

made in late FY 2004 and fully expended in FY 2007. A more detailed description of the program can be found in the solicitation referenced above, and it is mentioned several times in the broad 2007 report to the President.²³

From the standpoint of many observers, perhaps the single most significant result of ITR was to promote a broader view of computing: a view that computer science concepts and mechanisms are the core but not the only relevant discipline in computing. This was one of Bajcsy's objectives and was seen in what CISE (and to some extent what the other directorates) funded, but more importantly it was an objective in the computing and other scientific and engineering communities broadly. A second result was to enlarge the modality of most CISE funding from small, single-investigator grants to include medium (3 to 4 principal investigators) and large (more money for longer periods) grants.²⁴ This too was an objective of Bajcsy's, although she feels that the large grant category did not turn out to be successful because most of the large grants were simply conjunctions of a number of principal investigators' smaller grants rather than enabling larger efforts.²⁵ A third result of ITR, and part of Bajcsy's overall strategy, was to facilitate a change in the computer science (CS) community's view of cyberinfrastructure.

At the end of the initial round of ITR funding, just as the FY 2002 budget request was being prepared for submission to Congress at the end of August 2001, Bajcsy stepped down as AD/CISE; she then resigned from the University of Pennsylvania and accepted a faculty position at the University of California, Berkeley. George Strawn, her experienced Executive Director from the time she began at NSF, was named Acting AD/CISE, serving in that position until I was sworn in as AD/CISE in May 2002.

The initial ITR funding was limited specifically to five years, causing considerable concern in the computer science community. The hoped-for result was that, if the enhanced level of funding remained in the base budget of CISE after the official end of the program, there would then be funds available for new and/or modified CISE programs. Indeed, this happened in the preparation of the FY 2008 budget request (described below), which took place toward the end of my time as Bajcsy's successor as AD/CISE.

4.2 Peter Freeman (May 2002 to January 2007)

Following Bajcsy, I was the second AD/CISE with a Ph.D. in computer science. All the previous AD/CISEs were certainly computer scientists in practice, although their advanced degrees were in other fields such as electrical engineering or mathematics. I earned my undergraduate degree in physics in 1963 from Rice University, and a

master's from the University of Texas at Austin in math and psychology; I entered the Carnegie Institute of Technology (later CMU) with the first class of CS Ph.D. students in 1965. After receiving my Ph.D. in 1970 under Allen Newell, I took a tenure-track position at UC Irvine, where I served from 1971 to 1990. I then moved to Georgia Tech as founding dean and professor in the College of Computing, where I served until 2002.

I served as the first outside NSF division director in CISE from 1987 to 1989 and also served on the Computer Research Association (CRA) Board of Directors from that time until I rejoined NSF in mid-2002. I was fortunate to have known and worked directly with most of my predecessor Assistant Directors; for example, I had written my first published paper with Gordon Bell, Nico Habermann had been on my thesis committee, and I had been well acquainted with all of the others.

At the end of January 2002 when I accepted the position of AD/CISE, offered by NSF Director Rita Colwell, the Deputy Director (Joe Bordogna) immediately began involving me in budgetary and other issues.²⁶ When George Strawn moved to be CIO and after an open search, Deborah Crawford was named Deputy AD/CISE in late October 2002.

In the almost five years I was AD, there were a number of important decisions and activities. Rather than provide a strictly chronological description of my time as AD,²⁷ I begin with some general remarks and then describe six specific areas loosely in chronological order.

4.2.1 Context

I believe there are five major factors that enabled any success I may have had as AD/CISE: 1) prior experience; 2) length of service as AD; 3) the national situation during 2002–2006; 4) advice I received, especially from the Deputy AD, senior science advisors, and division directors; and 5) starting with a set of objectives and setting things in motion early in my term to achieve them.

As with all of the previous ADs, I had been actively involved with research and education at the highest level of a university (or a company in the case of Bell) up to when I joined CISE in May 2002. There is a qualitative difference in the responsibilities and activities of a dean or AD compared to those of a department chair or division director;²⁸ for anyone who has not served in a position at that organizational level, it may take significant time to grow into the new position. With the exception of Bell and Habermann (whose tenure as AD was cut tragically short by his death), none of my predecessors had had higher-level experience; fortunately, all adapted well. None of them had had direct, full-time experience within NSF or even the government, and my prior NSF experience was as a division director.

Closely related to having prior experience is the time one has to apply it. Even if you have abundant senior experience, in the timeframe dictated by the federal government it is hard to make a substantive impact on an organization as large as CISE in less than four full years. The most obvious reason is the federal budget cycle: The executive branch (of which NSF is a part) presents a budget request to Congress each year (year N) for the fiscal year that will begin the next October (year N+1). Generation of new ideas for the budget request for fiscal year N+2, to be presented to Congress the next February, begins almost literally the day after the President presents a budget request for fiscal year N+1. Full implementation of the budget appropriated by Congress, even if the appropriation is made by October 1, can take several months or more if it includes an entirely new program. By then the ideas you may have advanced, if actually included, are close to three years out of date. Any modifications of your plans based on experience with a new or changed program can easily take yet another year. In short, the AD and the staff are often dealing with three budgets simultaneously (years N, N+1, and N+2)!

A related factor—perhaps not fully appreciated by everyone coming in from the outside as an AD—is that because of the personnel standards of the United States Government (USG), it can take up to a half year (or longer, depending on budget) to recruit a new division director or a deputy AD. As a result, people already holding senior positions within the organization may not be fully aligned with a new AD; that may reduce the combined team’s effectiveness. It is common wisdom in most civilian organizations that permanent staff may quietly take the position relative to a new leader that they’ll nod “OK” to new directions but take their time in following them because they know they’ll outlast the limited time of the newcomer if he or she is a rotator. Once it is known that the newcomer will be in place for a potentially indefinite timeframe, this passive resistance may lessen.²⁹

The national context of the early 2000s is vividly remembered by all who were adults at that time. Within the USG, there was a widespread sense of urgency to push back on what had become a real threat to national security. This was especially true in the military and the intelligence community (IC). NSF had never supported classified research³⁰ and it has continued not to do so. Despite whatever pressure there may have been from Congress and other parts of the government, Rita Colwell (with great assistance from Joe Bordogna) steered NSF in a way that enabled it to be as relevant to that pressing national security need as it could be without changing its non-classified, non-mission-oriented character.

Colwell had long served as a senior advisor to the IC, so she had impeccable connections and insight.³¹ The ADs and some of the other senior professional staff were encouraged to add emphases to appropriate programs that might result

in advances that would be directly relevant to the efforts of the mission-oriented parts of the federal government. A few of us rapidly obtained high-level security clearances so that we could communicate with the IC and the Department of Defense on classified matters and then translate those insights into unclassified research initiatives. Some of us were also asked to convene meetings of relevant university-based principal investigators (PIs) and interested members of the IC to discuss relevant topics (e.g., language translation, sensors, data analysis). These meetings almost always served to make the IC better aware of results already in the public domain, as well as introducing them to experienced PIs. Some of those contacts may have resulted in individuals from the NSF community working directly with the IC, but those were always individual decisions. Because of our security clearances, a few of us also served on IC review committees for specific research topics.

The NSF peer review system is considered the “gold standard” based on *advice* given to the Program Officers by reviewers. Program Officers, in turn, make recommendations to support or decline individual proposals. Division Directors (DDs) review these recommendations and authorize the creation of grants for funding, subject to compliance checks. It is thus very important that an AD has DDs who not only will follow the general directions set by the AD but who can also reliably feed information up to the AD in a timely and coherent manner so as to help her or him steer an overall course for the directorate.

In the middle of this process is the Deputy AD, whose job is to oversee the detailed operations of the directorate and mediate the flow of information. If the AD is an Intergovernmental Personnel Act assignee (IPA)³²—which I believe strongly is one of the strengths of NSF—the Deputy AD ideally will be someone with deep and lengthy experience within NSF, including at senior levels. The Deputy AD (DAD), as opposed to an Executive Officer, is usually able to step in at any time and do anything that the AD does, albeit perhaps with less experience and less familiarity in the community. In this area, I was extremely lucky in being able to hire Deborah Crawford as my Deputy AD within a few months of joining as AD.³³ She had been a program director in multiple NSF divisions and offices and a senior advisor to the NSF Deputy Director (Joe Bordogna) for several years, with over a decade as an employee of NSF; her experience was wide and deep.

My experience as a founding dean, coupled with some excellent advice from a long-time mentor, meant that I already had a set of objectives in mind when I accepted the offer from Rita Colwell in late January 2002 to become the next AD/CISE. As noted above, I was immediately brought in as a consultant for the three months before my Georgia Tech service ended in early May. This transition

period afforded me the information and time to refine my list of objectives by the time I was sworn in as AD in early May (two days after giving a final exam in the class I was teaching).

My list of objectives contained six broad items: strengthen cybersecurity research; strengthen education and outreach efforts; start a new program in the broad fundamentals of software; rationalize the internal organization of CISE to better serve the computer science community; strengthen networking research; and better manage the NSF supercomputing centers for the benefit of all areas of science including computer science. Two defensive measures were taken to maintain the quality and quantity of research and to continue to grow the CISE budget and ensure that the ITR program funds would remain in the CISE base budget after the program officially ended.

Let's turn now to some specifics, organized around my own objectives and concerns. The order below is not chronological.³⁴ All of the activities listed below (with the exception of the reorganization) existed in some form when I started as AD; by the end of my term, all were in play, and some of the narrative below indicates when I was able to initiate my efforts in specific areas.

4.2.2 Reorganization of CISE³⁵

During my 12 years at Georgia Tech, I had very little direct contact with or internal knowledge of CISE. As a result, I had no idea before 2001 of the need for its reorganization. However, in the process of interviewing for the position of AD/CISE and during the intervening three months before I was sworn in, it became clear that there were a number of operational issues with the current organization of CISE: the soon-to-be-finished Cyberinfrastructure Report³⁶ was expected to put pressure on CISE from the other directorates (it did), there was growing interest in and pressure from Congress for more cybersecurity research, the generally small size of individual programs made programmatic flexibility difficult,³⁷ there was too little administrative support, and the narrow definitions of some programs were increasingly inconsistent with where the field was headed.

While operational in nature, these issues ultimately led to long-term strategic problems by making it difficult to focus on the various subfields and their interactions. In addition, the CISE budget was slated to grow rapidly because of ITR (it did), which would only exacerbate the other issues.

Through attendance at some major conferences and discussions with senior researchers, it became even clearer that these organizational issues existed. Internal discussions also made it evident that undertaking a reorganization was not an easy process: as in a university, everyone within CISE (and other directorates because of

their interest in CI) had an opinion on what needed to be done, so it was necessary to avoid immediately moving programs and people around. In addition, because NSF is part of the federal government, a highly structured organization, we needed to keep upper management informed, involve union representation to the extent required, and eventually justify a reorganization to the Office of Management and Budget and the House Science Committee.

Although there had been a few changes in program definitions over the 15 years since CISE was created and a modest reorganization was carried out by Juris Hartmanis, there had been no wholesale reorganization of the Directorate. Given the complexity of the process, it was almost as easy to reorganize everything at once than simply alter a few programs. It was also obvious that a reorganization needed to be undertaken soon and in a systematic, controlled manner that revealed proposed changes only when they had been well developed.³⁸

It had been clear to some members of the CS community for many decades that computing was on a classical exponential growth curve.³⁹ As applications useful to a wide, popular audience began to appear and catch on, a few people realized that the growth of the underlying technology's power, coupled with its decreasing size and cost, would soon explode onto the public stage in the form of revolutionary applications. This happened with full force in the 1990s, largely as a result of the opening up of the Internet (with CISE as the steward of non-military networks) to the development of an easily usable interface to online resources via browsers (starting with Mosaic), and to the development of highly efficient search algorithms (primarily the PageRank algorithm on which Google's search technology is based). All of these resulted from or were enabled by CISE support.

As that happened, a difference of opinion in the CS community developed between what might be called a narrow or classical view of what a CS unit should be (theory, systems software, data structures, and so on) and a broader or expansive view (of activities that depend on CS but also include other disciplines, such as human-computer interfaces, information structures, social informatics, robotics, bio-informatics, and so on).⁴⁰ Having come from CMU and UC Irvine, and having had experience within NSF in the late 1980s followed at Georgia Tech by building one of the first broad academic computing colleges, it was clear to me that, in the future, CISE needed to utilize the broader range of research funded through the ITR Program to support similar work. Large numbers of computer scientists were involved in building new, CS-based industrial and educational structures that were changing the world and *at the same time exposing great basic research questions for CS and related disciplines*. Viewed from the vantage point of today (2019), the choice to continue the "broad" support for computing research was clearly the right choice.⁴¹

By May 2002, when I was sworn in as AD/CISE, the terms of several DDs had ended (or were about to) and there were Acting DDs in several of the five divisions. Thus, I was in a position to make my own choices for several of these positions. At the start of January 2003, two of the DDs were newly arrived IPAs and three were regular NSF employees with moderate to extensive experience within NSF. Rick Adrion, a long-time NSF employee, IPA, and Advisory Committee (AC) member, as well as an experienced researcher and department chair, moved from being Division Director of Experimental and Integrative Activities (EIA) to become Senior Advisor in the Office of the AD (OAD). He was joined there by Larry Landweber, a long-time professor, department chair, and Internet pioneer as Senior Technology Advisor. Deborah Crawford had moved to CISE from the Office of the Director (OD) (as mentioned earlier), joining Steve Mahaney, Senior Financial Advisor in OAD and a former professor and program director (PD). This provided me an excellent, talented, and experienced executive team, making the reorganization easier than it might have been.

In early January 2003, I announced the Task Force on CISE Strategic Activities and held an all-hands meeting to explain the intent and process to all CISE personnel. The overall purpose was “to explore options and provide advice to the OAD/CISE on a number of strategic and organizational issues.” I announced the members of the task force, my expectation that they would finish their work within one year (in time for the FY 2005 Budget Request), and that the process would proceed in phases.⁴²

In the same memo, I noted: “The first phase activity [is] to define a role and mission for CISE with respect to cyberinfrastructure and to propose organizational and management strategies to achieve this mission.” That focus was explicitly driven by the expanding interest in CI across NSF, but implicitly by the broadening effect of ITR beyond traditional CS and the intent to make sure that CISE continued to lead NSF in all NSF activities involving research *on* computing (as distinct from *using* computing).

The task force performed its assigned task well, starting with a clean slate and general categories of current, and possible future, CI-oriented research and education (R&E) activities. The group took a fresh look at the breadth of activities that CISE might lead in without worrying about specific organizational changes. By April, the task force presented me a well-documented set of findings and possible options.⁴³ The findings began by strongly emphasizing that the highest priority for CISE was to address “long-term, fundamental research in computing, information systems, and communications.” They supported the notion that CISE should lead

the development and deployment of CI for all of NSF but warned against allowing that to obscure the primary responsibility of CISE.⁴⁴

The proposed structure had two major features. The first was the introduction of *clusters* to CISE—an organizational structure used in the Biology Directorate for some time. This system grouped together multiple topics and people into a single cluster of related topics (e.g., database systems, operating systems, and security) and involved several program directors and support staff. The concept of *themes* that would cut across multiple clusters, divisions, and even directorates made up the second new feature. This would permit a particular theme to be supported by multiple organizational units. One example was the theme of Education and Workforce Development, a topic that all research programs were expected to contribute to in some manner.

The team proceeded to develop and propose a reorganization of topics and professional staff into divisions aligned with the basic areas of responsibility of CISE—Basic Research, Research Infrastructure, Education and Human Resources, and Cyberinfrastructure for all of science. After the task force discussed the proposed reorganization with me in late April, I began briefing the Office of the Director (OD), the Office of Management and Budget (OMB), and Congress on an informal basis. From these interactions, I benefited from a great deal of constructive feedback. The result was an almost final reorganization with which we moved forward.

Two all-hands meetings were held in July at which I revealed and discussed in detail the motivations, objectives, and proposed details of the reorganization; this provided excellent feedback and suggestions. By September, we had a revised and detailed plan that we shared with the staff and Congress. By mid-October, we had the approval of OD, OMB, and the House Science Committee, so that we could then announce the details to the CISE community.⁴⁵

During all of the activities, I followed a detailed communication plan prepared by Deborah Crawford as we briefed OD, other ADs, the union that represented some staff and PDs, OMB, and the House Science Committee. Crawford's deep knowledge of the administrative side of NSF and the USG, as well as her strategic and common sense, were invaluable.

During FY 2004, starting in October 2003, we issued almost no solicitations and focused on the final year of ITR funding, ensuring that those projects received the funds they had been allocated. This provided CISE with some breathing room to work through the details of implementing an entirely new organization, set of solicitations, and supporting materials such as a revised website and continued communication with the community.

Let me offer some observations on the reorganization. First and most immediately, it addressed both of the defensive measures mentioned above (quality/quantity of research and securing and growing the CISE budget). The rational nature of the new structure permitted me to hire new DDs and PDs who were more in tune with the focus of a division, and to be in a better position at the program level to help reallocate budget to support fast-changing directions of research. In the financial realm, having a smaller set of broad efforts and organizational units permitted us to fashion more compelling arguments for increased funding than just asserting “we need more money.”

The simpler structure provided more transparency for high-level reviews by the CISE AC, Committees of Visitors (COVs), and others. From a management standpoint, it provided the Deputy AD and me with the transparency and focus essential for operational and strategic management. This was especially true for improving outreach and education. Even though those activities were dispersed to some extent, it made clearer where additional efforts were needed.

Overall, most staff accepted and learned to work with the new structure; within a year there appeared to be little dissatisfaction. Of course, initially there was a good bit of grumbling by some because they felt they were losing “their” programs. Organized cooperation among PDs was, for the most part, something new to CISE. Some of the administrative innovations to NSF, such as employing a business officer for each division to relieve the DDs and DDDs of some of the administrative detail of running their division, have been copied by other directorates. While there has been the movement of CI (first out of and then back into CISE) in the intervening years, as well as moves dictated by OD, the structure adopted in 2004 is mostly still in place today, over 15 years later.

In recent years, there have been three divisions and an Office of Advanced Computing (OAC) in CISE; each has had one or more Core Program clusters, and all have had an Education and Workforce cluster and participated in a wide variety of CISE-wide and NSF-wide cross-cutting programs spanning all aspects of research, education, workforce, and research infrastructure. Of late, the cluster concept has not been as completely utilized as in 2003, and some of the administrative functions we innovated have changed or gone away—not surprising for an organization that continues to grow with the arrival of new ADs every three or four years. In the meantime, the CISE budget has grown by more than 75% and the staff has doubled. For some time, CISE has apparently been considered the best place to work in NSF; organizationally and intellectually (in terms of innovating new programs), CISE has been considered by some to be one of the most forward-looking and active parts of NSF.

4.2.3 Strengthening Networking Research

When I started considering the CISE/AD position in 2001, my top programmatic objective was to strengthen networking research. This resulted from my long-time personal use of Internet predecessors (dating back to the early 1970s) and remote access to computational facilities (dating back to the mid-1960s),⁴⁶ as well as direct recent experience with a large networking project.⁴⁷ These experiences, coupled with what was happening commercially and with conversations with a number of senior networking researchers, made clear there was going to be a continuing need for high-quality, basic research in networking.

The situation in 2002 in the community of advanced, networking researchers appeared to me to be one of malaise and loss of direction, even though there were some individual projects of high merit. This sense was captured well by Scott Shenker in his keynote speech accepting the 2002 SIGCOMM Award in August that year.⁴⁸ His speech reviewed the past and present situation of SIGCOMM and commented on the future. He noted that the community he was addressing that day had not created the Internet but had brought it to its then-current state (“We were the Internet’s teenage babysitter”). He went on to argue that the work SIGCOMM members had been engaged in had three vital components: it was “intellectually deep, transformed the world, and was community-driven.”

Shenker noted that in spite of the Internet being a commercial success, there were a shrinking number of opportunities for (networking) researchers to have further impact. When he turned to comment on the future, he posed the question, “How can we retain the three vital components—intellectual depth, transformative impact, and community?” His answer came down to “Focus on transformative community projects that can engage a community and transform the world.”⁴⁹ He noted that this ideal was not restricted to any particular project. This struck me as encapsulating and reinforcing what I was seeing at NSF, and it inspired me in my objective of revitalizing NSF support for networking research. Before describing what we did, it may be useful to review some background.

The earliest computer network technology, notably including the concepts of packet switching, distributed and open networks, and the TCP/IP protocols, had been developed largely under Department of Defense (DoD) contracts with the RAND Corporation and BBN Technologies (originally Bolt, Beranek, and Newman) through (D)ARPA support. Yet it was with NSF support around 1980 that Theorynet and soon NSFNET⁵⁰ showed the superiority of these concepts in practice.⁵¹ When DARPA decided in the mid-1980s to cease direct support of basic networking research, NSF, under the leadership of Steve Wolff, led the conversion of NSFNET to an open, public network in mid-1995.⁵²

This community of researchers, while able to hail great success, was soon left without the ability to experiment at anything approaching the scale of a commercial production network. From my perspective, the researchers were in effect pushed aside by industry in its rush to capitalize on the “new” technology, or in some cases were swallowed up by industry.⁵³ I knew that NSF couldn’t single-handedly change that situation, but at the same time my developing understanding of the situation—initially from the outside and then from the inside when I joined NSF in 2002—was that perhaps NSF also had lost its way in networking. This was exacerbated by the events of 9/11, the growing realization of the Internet’s power for good and ill, and NSF’s desire to aid efforts aimed at national security without abrogating its long-held principle of open research.⁵⁴ At the same time, work by DARPA that touched on networking was either classified or focused on immediate results that favored industry players, further reducing financial support for the academic research community.

Fortunately, there was a core of experienced personnel at NSF who understood this situation, including but not limited to George Strawn, Aubrey Bush (a rotator), and Darleen Fisher. In the fall of 2002, I was fortunate to bring Larry Landweber on board as a Senior Technology Advisor (a rotator); in addition to understanding networking at a deep technical level⁵⁵ and having impeccable worldwide contacts in networking, he had been a long-time department head at a major CS research department. With the assistance and detailed observations of these people and others, we forged a plan for a way forward.

One of the most obvious problems was that networking had been a part of the division that also managed the NSF Supercomputing Centers until the reorganization under AD Juris Hartmanis in 1998. Even after that the title of the division and much of its budget was operationally focused. As is usually the case, the urgent issues of production work tended to dominate concerns for long-term basic research. This meant that there was an insufficient group of leaders in networking who could initiate solicitations and enable the creation of new funding programs.

The second clear need was for additional funds and personnel devoted to networking research and leadership devoted to that end who were free of near-term infrastructural or commercial needs. The efforts toward homeland security favored research in CISE (and to some extent ENG in the area of devices and complex systems). The Cybersecurity R&D Act of 2002⁵⁶ explicitly called for NSF through the CISE Computer and Networking Systems (CNS) Division to be the lead federal agency in advanced cybersecurity research. While the Act was only an authorization, not an appropriation of funds, the pressure of the House Science Committee for us to increase our funding for cybersecurity, of which networking was a large part, made obtaining more funds easier than it might have been.

The reorganization described in the previous section separated networking research from the supercomputer centers and their operational needs. Placing networking in a new division that incorporated other basic systems technology, including databases and operating systems, made it easier for programs to coordinate their efforts. The concept of clusters further improved communication between different areas of research.

Strawn, already in a senior position in OAD, moved to a different part of NSF to become CIO in late 2002, and Aubrey Bush's IPA assignment was soon ending. Landweber initially devoted much of his time not to advising me, but to advising Bush and the program directors on developing new solicitations. These efforts laid the foundations for later work, as well as providing support to some important activities already underway in the community, such as Planet Lab⁵⁷ at Princeton and a number of other sites, including EMULAB⁵⁸ at Utah, and ORBIT⁵⁹ at Rutgers.

After Bush departed in 2003, the Acting DD/CNS was Mari Maeda, an experienced employee, who served while we searched for someone from the community to fill the post. The addition of Greg Andrews as DD/CNS in 2003 (initially EIA) and Wei Zhao in 2005, both for two-year terms as rotators, brought experienced, active computer science researchers with significant leadership experience to head CNS. Several of the PDs whose rotation terms had ended or who chose to take different assignments created opportunities for new people. Through Landweber's connection with David Farber (also a colleague of mine dating back to 1971), we became aware of Guru Parulkar.

Parulkar, a former student of Farber's, had been a professor at Washington University in St. Louis (WU). With Jon Turner, also at WU and a noted senior networking researcher, he had created a company with new networking technology, which had just been sold to Cisco. When they both left Cisco, they returned to academe but Parulkar also wanted to give back to his adopted country by serving at NSF. We gladly hired him as a PD in CNS starting in 2003. His connections to the leading edge of networking research and his entrepreneurial instincts soon led him to build on several existing lines of research led by Jon Turner, Larry Peterson, Scott Shenker, and Tom Johnson.

In April 2004 Parulkar requested time with Deborah Crawford and me to present several ideas and a proposal for new research. The presentation was compelling, resonated strongly with us, and provided new ideas to energize networking research broadly. Before he left my office, we had signed on to creating a new initiative. Crawford's experience in NSF was key to understanding the degree of latitude we had in doing this. As a result, Parulkar left my office that day with a mandate and strategic guidelines to create what became the GENI Project.

The remainder of this period in networking activity was almost entirely dominated by a growing community of, and grants to, academic researchers engaged in planning GENI. That story is told in Chapter 9.⁶⁰

4.2.4 Starting a New Program in Software Fundamentals

The process of design, and the results it can produce in many different fields of endeavor, has been and remains my core intellectual interest. Early in 2003, I wrote a short vision statement entitled “Science of Design,”⁶¹ which I shared with a few people at NSF. Later that spring, I expanded on the vision⁶² and proposed some actions to the DDs, which resulted in some workshops and small planning grants. As noted in this narrative, there were many other, higher-priority initiatives underway; not until 2004 did we seriously begin a new program.⁶³

The basic point I was stressing is that, after almost 40 years of improvement in software engineering techniques, developing complex software was still mostly done in an ad hoc manner, too often resulting in unforeseen but serious consequences. The primary result I was aiming at was to spur long-range research to provide a scientific foundation for significant improvement in software development capabilities that would be fundamental and far-reaching.⁶⁴

The solicitation elicited a modest response, with several good proposals. While it resulted in some useful research, the program was soon discontinued due to a change in personnel—including my leaving NSF in early 2007 as well as limited interest in the community. I think this short attempt to spark interest in putting software development on a more fundamental basis failed for two basic reasons. First, the excitement (and rewards) of actual design efforts in any field will always dominate theorizing about the process; and second, while at a high level there appear to be similarities (e.g., choosing among alternatives, techniques for evaluation) across fields as diverse as bridge design and software creation, when one tries to carry these high-level observations down to specifics in a particular field, very little practical progress can be made. I first encountered this fundamental truth in a seminar in the 1960s that included a chemical engineer, a noted computer hardware architect (Gordon Bell), a designer of computer instruction sets, and a polymath (future Nobel laureate Herbert A. Simon). Sadly, nothing had changed in almost 40 years, but I forgot that early lesson!

4.2.5 Strengthening Cybersecurity Research

Fortunately, although I came into CISE understanding the need for more cybersecurity research from my previous activities,⁶⁵ the heightened sense of vulnerability of the United States in the wake of 9/11 and emphasized by the Cybersecurity R&D Act of 2002⁶⁶ provided a compelling case, allowing us to argue successfully for more

funds in that area. It also aided in recruiting serious and well-regarded experts in cybersecurity as PDs. In early 2003 I prepared a memo outlining the actions we might take in CISE to strengthen advanced research and education in cybersecurity.⁶⁷ Over the next several years this served as a guide for undertaking a number of activities.

One can parse cybersecurity research into two broad categories: 1) encryption and other theoretical/formal means of securing information directly so that even if divulged to “bad actors” it is of no use without very clever techniques or massive computing power, and 2) making systems (platforms and networks that connect them) more secure from compromise. Hardware engineers would add a third category that includes physical shielding and hardware embedding of some processing. While important, most of the efforts in this third category were beyond our purview and, in any event, irrelevant for the most part without security at the levels noted above in most systems. An area that CISE did have a hand in—but that we didn’t fully exploit—is the understanding of the human actors in a technical system.

The first category however, is one that CISE and its predecessor organizations have supported for a very long time. For example, by the late 1980s our CS Theory program already had a long history of support and cooperation with other agencies, notably the NSA.⁶⁸ Beyond encouraging strong support for these efforts in the Computing and Communication Foundations (CCF) Division and providing some good contacts for its DD and PDs, there is little else to say about that area. Of course, as is often the case, emphasis in solicitations (and the availability of funding) from CCF attracted additional, qualified researchers and students. Some very good research results came out of those efforts, but this is not intended to be a review of those.

The second category had not been particularly prominent in CISE-supported systems research prior to 2002 (with some very important exceptions already noted above). In the new organization, the DDs and PDs in CNS were better able to expand and report on their activities in cybersecurity research, even though it often occurred in the context of networking research, notably in the preliminary work on GENI. The supercomputer centers together became a powerful player in the cybersecurity arena because of their need for security and equally importantly because of their large and experienced staffs.

Beginning in 2002, we undertook as much as we could to strengthen CISE’s activities in cybersecurity. These endeavors included a more integrated approach among several existing efforts in databases, networking, and operating systems; extension to new areas (e.g., power grid security); and even a modest cybersecurity centers program.

4.2.6 Managing the Supercomputer Centers and Cyberinfrastructure

NSF support for the instrumental use of computers began in the early 1950s, before support for the study of the behavior of computations and the theory underlying that behavior (i.e., for computer science), which began around 1960. As activity in both the instrumental use of computing and the study of it has increased dramatically in the decades since then, the relationship between them has been an uneasy one. The two realms are sometimes complementary and sometimes competitive when it comes to obtaining limited funds. This fundamental tension is reflected in internal discussions and organizations within NSF.

Demands for more, larger, and faster systems from the user communities ultimately depend on advances in computer science. Only relatively recently has it become fairly widely understood that industry alone cannot make all the needed basic advances. At the same time, at NSF and in similar organizations, there are always limited funds. The duality in this relationship is at the core of the issues between the two groups. It has posed an organizational and management problem from the Director of NSF down to program directors over the lifetime of NSF.

Support for instrumental usage, now largely embedded in the NSF supercomputer centers and cyberinfrastructure efforts, has moved into and out of CISE several times. This has plagued every CISE AD, beginning with Gordon Bell in 1986. The centers were in the process of formation and expansion essentially at the same time that CISE was created; both developments were the idea of and overseen at the highest level by Erich Bloch. The importance of the supercomputer centers and their predecessors deserves a more focused recounting of that history.⁶⁹ The bottom line is that issues surrounding supercomputing have frequently taken a large amount of the AD/CISE's time.

The most important events in this area during my term were bound up with the Atkins Report and the strongly emerging interest in all areas of science for additional large data and computational resources. The supercomputer centers tried to meet not just this demand, but also the increased interest in broadening participation in computing. The fundamental issue, however, remained and is still unresolved: computer science research has a lot to offer to the long-term advancement of the power of computational facilities, while the rest of science is impatient for more immediate functional resources.⁷⁰

4.2.7 Strengthening Education and Outreach

NSF has been a leading force for improvement in education at least since the late 1950s (Sputnik and grants to develop computing education) and for outreach

to under-represented groups at least since the late 1970s (a grant for a women's re-entry program).⁷¹ It is perhaps a truism that outreach often begins with education, and education (in a field) enhances outreach. As a result, most of the efforts undertaken in CISE effectively promoted both better education and outreach to underserved communities.

CISE and its predecessors had long participated in some of the more ambitious and effective programs in this domain, including Integrated Graduate Education and Research Training (IGERT)⁷² and ADVANCE.⁷³ The strongest CISE effort, which in effect reached all areas of science, included the education programs undertaken by the supercomputer centers and later other large centers supported by CISE. While they were very effective and impacted a large number of students, the programs at the Centers did not directly serve to attract many students into CS educational programs; instead, they brought students into touch with the use of computers or programming in other disciplines, and less into CS proper.

Because of my background, I was well aware of the need for educational improvements and for encouraging more diversity when I started as AD. I was also somewhat aware of the recent CISE efforts in this area, which had produced some measure of useful results by studying and validating different approaches to education and outreach. I felt that research of that sort should be continued, but that it was time for actions that might lead to major changes.

At my urging, we tried to raise the importance and visibility of education. We established an award for outstanding CS educator of the year, but unfortunately that effort that did not take root beyond a couple of years. We undertook an initiative to create a new computing curriculum better suited to recent developments such as the Internet, but it didn't really take root, either; the lesson learned was that the professional societies (ACM, IEEE, SIAM, and others) are better suited to doing this.⁷⁴

The inclusion of women in computing, especially in advanced education and research, was an area of outreach/education that I had long been interested in as a professor and dean.⁷⁵ As a member of the Computing Research Association (CRA) Board of Directors from 1999 to 2002, I had the opportunity to witness outreach efforts that worked in practice—and some that didn't. One that was very successful for the CRA was its efforts to include women on the Board and through a committee (CRA-W) to advance women in the field. While there were several leaders in this endeavor, Jan Cuny,⁷⁶ a CS professor at the University of Oregon at the time, stood out. When I attended a reception sponsored by the National Center for Women & Information Technology in 2003, I ran into her. Knowing her success in helping younger professional women learn to lead, I immediately cornered her to seek her

advice on what we could do at CISE. By the end of a fairly long conversation, I knew that she was the one to replicate at CISE what she had done for CRA-W, but on a broader front including all under-represented groups. Thus, I broached the subject of her coming to NSF for a few years.

Cuny demurred, but I convinced her to at least visit us. She did, and eventually I was able to hire her to create a program to broaden participation in computing. Starting in 2004 with what funds we could muster, she started the Broadening Participation in Computing (BPC) Program (see also Chapter 11).⁷⁷ She created a portfolio of projects focused on getting results through cooperation. It worked beyond anything that I had imagined and is still in operation today,⁷⁸ although it has evolved substantially. Its success has extended well beyond CISE and served as a model for NSF and other agencies. It also has spawned other programs that are making their mark in incredibly valuable ways, including efforts to train teachers and develop interest in computing among children from preschool through high school.⁷⁹ The results of hiring Cuny and the programs she started have been transformative in this highly important area of encouraging all people to enter computing as a profession and to ensure that they have access to a quality education to prepare them.

In response to an unsolicited proposal in 2004, a grant was made (after peer review) to fund the fledgling National Center for Women & Information Technology (NCWIT),⁸⁰ enabling it to get off the ground with serious funding via a large, multi-year grant.⁸¹ Prior to this award, it had been supported by the University of Colorado, where it was headquartered, and a few small grants from NSF and industry. The awarded funds (approximately \$4 million over five years) provided it with financial stability and, perhaps more importantly in the long run, a major vote of confidence. This reassured major IT corporations that NCWIT was a serious organization. As of early 2019, NCWIT “is a non-profit community of more than 1,100 universities, companies, non-profits, and government organizations nationwide working to increase girls’ and women’s meaningful participation in computing.”⁸²

The success of the BPC program and the one-off grant to NCWIT clearly speak to the difference that one person and a fairly small, single action can make. Both also speak to the fact that CISE has become a leader within NSF by developing new programs in emerging areas and finding the people to lead them; as described in previous sections of this chapter, the same can be said of other areas such as networking.

This description of what has turned into an enduring effort in NSF also provides a good capstone to the description of this period of CISE’s history. Between 1999 and 2007, CISE went from being a small directorate with fairly routine pro-

grams and activities to one known for innovations that are felt throughout NSF, the scientific community it serves, and beyond.

4.3 Closing Observations

The years 1999–2006 were years of profound change, driven by two things: budget and active leadership. The ITR program almost doubled the CISE budget and that alone changed the research opportunities for the community. In 2000, when ITR began, the dot-com explosion of connectivity and computer-based applications across society was peaking, but that was followed just as rapidly by the dot-com crash. The events of 9/11 that soon followed appropriately caused other agencies, especially DARPA, to rededicate their efforts toward the immediate needs of protecting the nation. Largely due to the efforts of Rita Colwell to involve NSF in more immediate activities without seriously impacting our long-term mission of basic research, NSF was able to preserve the promised ITR funding.

It then fell to me, starting in early 2002, to use the growing funds in ways that continued to support core areas of computer science while expanding our efforts into new and innovative areas of research and education. Beginning in 2004, as the last round of ITR-funded grants was made, we were able to demonstrate two things: the efficacy of those efforts and the ability of a reorganized CISE with a new cohort of leaders to responsibly manage an expanded budget as well as to produce new and exciting results. In 2006 as we were planning for later years after the last ITR grants were completed, it was a foregone conclusion that CISE would retain the added funds in its base budget—the NSF leadership and Congress were already convinced.

This gave us the unparalleled opportunity to recommend major changes in the CISE budget we proposed for FY 2008. The recommendations we made in 2006 were largely appropriated by Congress and the subsequent internal allocations by the NSF Director; these included almost doubling the BPC Program, funding of GENI, continuation of the CCC, concept and budget for the Expeditions in Computing Program, expanded international activity by CISE, and what has become known as Cyber-enabled Discovery and Innovation (CDI). The overall result was to embed the broadening of computer science in major CISE programs and to not only continue major initiatives but plant the seeds for future developments.

By 2006 the profound changes that were underway in the late 1990s (which CISE had helped spark by shepherding the creation of the Internet, the first easily used browsers, and Google) and which were continuing, played the key role in showing what modern computing could do and how investment in research might

help it to continue. Without the leadership initially of Rita Colwell and Ruzena Bajcsy, followed by the outstanding work of the team I was able to assemble, the ascendancy of computing research at NSF and beyond might have either not occurred at all or fallen to other agencies where mission-oriented research (by definition, their responsibility) might have crowded out the basic research and advanced education that was and is the responsibility of NSF.

Notes

1. Even though I was AD/CISE from mid-2002 to early 2007, I have used the first person in this chapter. While my personal involvement may have biased some of my observations below, I have tried to present them objectively and where possible supported them by third-party and/or primary sources.
2. G. Moore. April 19, 1965. Cramming more components onto integrated circuits, *Electronics Magazine*, 38(8). <https://mashable.com/2015/04/19/what-moores-law-doesnt-say/#TzUCqI7Zaaqs>; last accessed 23 September 2019.
3. Rita Colwell had a long history with computing. She had programmed an IBM 650 (in assembly language!) for her Ph.D. research around 1960, had been on the Advisory Committee for the Science Information Office in 1977, and served on the National Science Board. (Oral history, Rita Colwell, interviewed by William Aspray, July 31, 2017. Charles Babbage Institute.)
4. Colwell interview, 2017, *op. cit.*
5. By context to her remarks later in the interview where she was referring to the building of cyberinfrastructure for all of science and the Information Technology Research (ITR) Initiative, both described later in this chapter.
6. It is worth noting that both were married to physicists, had faced the extra challenges posed to female scientists, had daughters who were also scientists, and were of similar ages.
7. See “Information Technology Research Program” (ITR), Chapter 8 in this volume.
8. President’s Information Technology Research Committee. February 1999. Information Technology Research: Investing in Our Future. Report to the President. https://www.nitrd.gov/pitac/report/pitac_report.pdf; last accessed 14 February 2019.
9. Ruzena Bajcsy. October 9, 2016. *Acceptance Remarks*. Simon Ramo Founder’s Award. <https://www.nae.edu/Activities/Projects/Awards/FoundersAwards/FoundersWinners/162256/162272.aspx>; last accessed 14 February 2019.
10. Oral history, Ruzena Bajcsy, interviewed by William Aspray March 19, 2017. Charles Babbage Institute.
11. Oral history, George Strawn, interviewed by William Aspray, July 19, 2017. Charles Babbage Institute.
12. Today (2019) the term is in broad use with many overlapping connotations. According to Adrion, who was a division director at the time, Bajcsy apparently had an almost all-encompassing definition including everything from hardware and software to the techniques used to design systems. The report of the panel (known as the Atkins

Report) narrowed that somewhat and since I was in charge of implementing the report's recommendations, I made it even more specific.

13. The charge to the panel is reprinted in the Atkins Report and asked them to 1) evaluate current major investments in cyberinfrastructure, and 2) recommend new areas of emphasis relevant to cyberinfrastructure; it also 3) proposed an implementation plan. The details of the charge fill three pages.
14. <https://www.nsf.gov/cise/sci/reports/atkins.pdf>; last accessed 16 September 2018.
15. Even the question of where within NSF the funding of CI should reside has continued to change. One result of this is that the direction of CI has not been located in CISE continuously, probably to the detriment of the overall effectiveness of the funding.
16. https://www.nsf.gov/news/news_summ.jsp?cntn_id=103044; last accessed 23 September 2019.
17. <https://www.nitrd.gov/historical/it2/it2-ip.pdf>; last accessed 23 September 2019. The IT2 program boosted the NSF budget by \$146 million (Adrion private communication).
18. Colwell interview, 2017, *op. cit.* See also Bajcsy interview, 2017, *op. cit.*
19. AD Service Dates Table. Charles Babbage Institute.
20. Colwell interview, 2017, *op. cit.*
21. NSF 99-167, *Program Solicitation: ITR*, 1999. Charles Babbage Institute.
22. <https://www.nsf.gov/pubs/2004/nsf04012/nsf04012.htm>; last accessed 23 September 2019.
23. *Leadership Under Challenge: Information Technology R&D in a Competitive World*, President's Council of Advisors on Science and Technology (PCAST). 2007. Charles Babbage Institute.
24. An exception was the CER Program that was started in the late 1970s and is described in Chapter 2 in this volume. In that program, the first several rounds of funding went to large, multi-year, multi-investigator grants. Some of those grants, however, were more of the form of support for similar, not focused, activities.
25. Bajcsy interview, 2017, *op. cit.*
26. It was made clear to me that the position of AD was a true executive position. My job was meant to be primarily *outward-facing* to the computing community, to Congress, and even internationally in my field, with a strong senior staff that could follow my strategic lead in implementation and operations.
27. To further understand the organization and chronology of CISE during this period, please see the "Computing Organizations at NSF" addendum in this volume.
28. The close parallel between the organization of NSF and of a major research university should be obvious (Director=President, Deputy Director=Provost, ADs=Deans, and so on).
29. Erich Bloch, when addressing staff thinking he was only a short-timer, is quoted as saying: "I might outlive you. So, think about that!" In T. J. Misa and J. R. Yost. December 27, 2015. *FastLane: Managing Science in the Internet World*. JHU Press.
30. There may have been a few instances of classified research in the early 1950s as NSF was developing the details of its mission.
31. Colwell interview, 2017, *op. cit.*

32. “IPA” stands for the Intergovernmental Personnel Act of 1974, a law that permits the USG to exchange personnel for a limited time with non-profit organizations.
33. Crawford joined CISE at the beginning of November 2002.
34. Timeline of Significant Events. Charles Babbage Institute.
35. This section is based primarily on our recent interview with Deborah Crawford (interviewed by William Aspray, June 20, 2017. Charles Babbage Institute). Also, based on personal recollections.
36. <https://www.nsf.gov/cise/sci/reports/atkins.pdf>; last accessed 25 September 2018.
37. Congress appropriates funds down to the program level and provides oversight to prevent re-appropriation during the budget year.
38. “Task force on CISE strategic issues,” draft of internal document to all CISE personnel. January 13, 2003. Charles Babbage Institute.
39. Moore’s Law is the archetypical expression of this.
40. For some of the defining ideas, see P. J. Denning et al. January 1989. Computing as a discipline, *CACM*, 32(1): 9–23. DOI: [10.1145/63238.63239](https://doi.org/10.1145/63238.63239); last accessed 23 September 2019. See also: National Research Council. 1992. *Computing the Future: A Broader Agenda for Computer Science and Engineering*. National Academy Press.
41. Interestingly, initial support for research on computing included a strong strain of science information and information systems. See Chapter 2 in this volume.
42. “Task force on CISE strategic issues,” January 13, 2003. Charles Babbage Institute.
43. “CISE strategic initiatives task force preliminary findings and recommendations,” draft internal document. March 25, 2003. Charles Babbage Institute.
44. This love-hate relationship with CI still resonates and is still not settled, leading among other things to the responsibility for CI moving in and out of CISE.
45. S. Jackson. September 2013. CISE FY 2004 update, *Computing Research News*, 15(4): 4. Accessed from <http://archive.cra.org/CRN/issues/0304.pdf>; last accessed 23 September 2019. See also: P. Freeman. 2003. NSF/CISE: Looking back, looking forward. *Computing Research News*, 16(3): 4. <http://archive.cra.org/CRN/issues/0403.pdf>; last accessed 23 September 2019.
46. In 1971/1972, as a new faculty member at UC Irvine moving from a position at CMU, I had a local-phone connection to one of the first ARPANET nodes in order to do file transfers to/from other ARPA-supported research sites. Earlier, in 1964/1965, I designed and built a time-sharing operating system for a minicomputer (CDC160A) that was used in production access to a larger machine at the University of Texas at Austin. As a graduate student at CMU beginning in 1965, I had remote hard-wired access to production machines and, shortly afterward, dial-up access using one of the first experimental acoustic couplers.
47. In the mid-1990s I headed a project to completely (re)wire the Georgia Tech campus with fiber and implement a series of production networks that utilized the very latest technology of the time and provided capabilities not then found on any major campus network. The FUTURENET Project (a concept of Ron Hutchins) was part of the technology preparations for the 1996 Olympics, which utilized the Georgia Tech campus as the Olympic Village.

48. <http://conferences.sigcomm.org/sigcomm/2002/adprog.html>; last accessed 23 September 2019.
49. Shenker's speech was not written and is not published anywhere. He has graciously provided a copy of his slides, which can be found at Charles Babbage Institute.
50. https://web.archive.org/web/20170202190223/https://www.merit.edu/wiki/NSFNET_final.pdf; last accessed 23 September 2019.
51. This was a bit of role reversal, since it usually had been, and still is, the case that in many areas DARPA is more focused on application than NSF is.
52. Oral history, Larry Landweber, interviewed by William Aspray, July 21, 2017. Charles Babbage Institute. Oral history, Steve Wolff, interviewed by Rick Adrion, July 20, 2017. Charles Babbage Institute. <https://www.livinginternet.com/doc/merit.edu/index.html>; last accessed 23 September 2019.
53. As dean of a growing college in the 1990s, I was on the front lines of this drain of talented faculty. This was not always a total loss in the long term if the researchers eventually returned to the academy—or were able to fund good research with the fortunes they had quickly made in industry!
54. Rita Colwell, NSF Director at that time, was instrumental in this. Our interview with her speaks to this point.
55. He had been one of the main originators of Theorynet. L. Landweber. 1979. Theory Net: An electronic mail system. *ACM Proceedings of the 1979 Annual Conference*, 158. DOI: [10.1145/800177.810053](https://doi.org/10.1145/800177.810053); last accessed 23 September 2019.
56. <https://www.congress.gov/bill/107th-congress/house-bill/3394>; last accessed 23 September 2019.
57. <https://www.planet-lab.org/history>; last accessed 23 September 2019.
58. <https://www.emulab.net/portal/frontpage.php>; last accessed 23 September 2019.
59. <http://www.winlab.rutgers.edu>; last accessed 23 September 2019.
60. See Chapter 9 in this volume: R. McGeer, M. Berman, C. Elliott, R. Ricci, eds. 2016. *The GENI Book*, Springer International Publishing. M. Berman, J. S. Chase, L. Landweber, A. Nakao, M. Ott, D. Raychaudhuri, and R. Ricci, and I. Seskar. 2014. GENI: A federated testbed for innovative network experiments. *Computer Networks*, 61: 5–23. DOI: [10.1016/j.bjp.2013.12.037](https://doi.org/10.1016/j.bjp.2013.12.037); <http://www.geni.net/about-geni/geni-bibliography/>; last accessed 23 September 2019.
61. "Science of design," vision statement. April 2, 2003. Charles Babbage Institute.
62. "Science of design—Vision and FY03 expectations," memo. June 16, 2003. Charles Babbage Institute.
63. "Science of design," Solicitation NSF 04-0552. <https://www.nsf.gov/pubs/2004/nsf04552/nsf04552.htm>; last accessed 14 February 2019.
64. "Science of design," visual. October 4, 2005. Charles Babbage Institute.
65. I had consulted at the Rand Corp. in 1972 on computer security, and then in 1998 was the program chair for the Sam Nunn Bank of America Policy Forum on Information Security held at Georgia Tech.

66. <https://www.congress.gov/bill/107th-congress/house-bill/3394>; last accessed 23 September 2019.
67. CyberTrust-draft 4-7-2003.doc. Charles Babbage Institute.
68. As DD, I gave a talk to researchers deep inside NSA HQ in 1988/89. (My only other visit inside NSA was the morning of 9/11, when I was Acting Director of the Georgia Tech Information Security Center, for a meeting of directors from other universities that had just convened—and was ended by a mass evacuation of the NSA buildings.)
69. See Chapter 10 in this volume.
70. See Chapter 10 in this volume.
71. See Chapter 11 in this volume.
72. <http://www.igert.org/public/about.html>; last accessed 23 September 2019.
73. https://www.nsf.gov/funding/pgm_summ.jsp?pims_id=5383; last accessed 23 September 2019.
74. Perhaps with NSF financial support.
75. My interest was perhaps foreordained—I have my grandmother’s handwritten geometry notes from 1896 and my mother’s master’s thesis in math from 1950.
76. <https://www.ncwit.org/profile/jan-cuny>; last accessed 23 September 2019.
77. <https://www.nsf.gov/pubs/2005/nsf05562/nsf05562.htm>.
78. <https://www.nsf.gov/cise/bpc/>; last accessed 23 September 2019.
79. <https://cra.org/jan-cuny-acm-award/>; last accessed 23 September 2019.
80. <https://www.ncwit.org>; last accessed 23 September 2019.
81. https://www.nsf.gov/awardsearch/showAward?AWD_ID=0413538&HistoricalAwards=false; last accessed 23 September 2019.
82. <https://www.ncwit.org/ncwit-fact-sheet>; last accessed 23 September 2019.

2007–2016: The Growing Centrality of CISE to NSF

William Aspray

This chapter covers the history of computing at NSF for more recent years, 2007 to 2016. This was a period in which CISE was led successively by Jeannette Wing, Farnam Jahanian, and James Kurose, but also led by Acting ADs for significant periods of time before each AD's term: by Deborah Crawford, Peter Arzberger, and Suzi Iacono, respectively. During this era, CISE was already a well-established directorate supporting a well-established scientific discipline. Nevertheless, the computing field and CISE continued to grow rapidly. This growth was enabled in part by a large budget increase arising from President Obama's stimulus package. There continued to be adjustments in the way NSF handled cyberinfrastructure support organizationally, and how that related to support for "basic" computer science research. This was also a period in which there was growing interaction between CISE and other directorates, other federal agencies, and organizations in other countries—primarily because of the growing recognition of the centrality of computer science to most scientific and engineering fields, and to society at large. The chapter is organized chronologically, with one section about each AD or Acting AD of CISE.

5.1 Deborah Crawford (Acting AD, Early 2007)

When Peter Freeman's term as CISE AD ended in January 2007, Deborah Crawford served as Acting AD for five months. Her background included a bachelor's degree in electronic and electrical engineering from the University of Bradford and a doctorate in information systems engineering from the University of Glasgow. Before coming to NSF, she worked in high-speed optical and optoelectronic systems at AT&T Bell Laboratories; the University of California, Santa Barbara; and the Jet

Propulsion Laboratory. In her 17 years at NSF (1993–2010), she held various senior management positions in the Computing, Engineering, and Education and Human Resources directorates, and worked in the Office of the Director. After leaving NSF, she served as senior vice provost for research at Drexel University, president and executive director of the International Computer Science Institute in Berkeley, California, and as vice president for research at George Mason University.

Crawford had worked as Deputy Director of CISE under Peter Freeman. During the first years of Freeman's time as CISE AD, cyberinfrastructure programs were supported in CISE. However, some computer scientists worried that cyberinfrastructure grants were eating into support for more traditional computer science, while other scientific and engineering professions felt that cyberinfrastructure should not solely be under the control of CISE. NSF Director Arden Bement decided to pull cyberinfrastructure activities out of CISE and in July 2005 created an Office of Cyberinfrastructure in the Director's Office. Crawford had previously worked for NSF Deputy Director Joe Bordogna and had been tasked with figuring out the Foundation's response to an important earlier report of a committee on advanced cyberinfrastructure led by Dan Atkins. When she moved to CISE, Crawford worked closely with Rich Hirsh and Dick Hildebrand, who had led the cyberinfrastructure activities inside CISE. Not surprisingly, the Director tapped her to direct this new Office of Cyberinfrastructure. It was from this position that Crawford moved back to CISE to be Acting AD after Freeman's departure, until Jeanette Wing's arrival as the new CISE AD.

Crawford had worked closely with Freeman, and her term as Acting AD largely involved continuing the programs that she had run jointly with him. Crawford notes that Freeman left the CISE Directorate in good organizational and financial shape (e.g., not too many ongoing financial commitments), thus making it possible for Wing to take more initiative and also serving as a programmatic and organizational basis for the work of future ADs through the Kurose era.¹

5.2 Jeannette Wing (AD, July 2007 to June 2010)

Jeannette Wing's background included bachelor's and master's degrees in electrical engineering and computer science and a doctorate in computer science—all from MIT. She taught computer science at the University of Southern California (1982–1985) and at Carnegie Mellon University (CMU, 1985–2012; while on leave at NSF 2007–2010). She served two terms at CMU as head of the computer science department. She is a leading scholar in the area of formal methods in software. Since her AD appointment at NSF, she has served as corporate vice president of

Microsoft Research (2013–2017) and since 2017 as professor of computer science at Columbia University and as director of its Data Sciences Institute.

Wing noted that when she joined CISE in 2007, DARPA support for computer science had waned (compared to the 1980s) and NSF had taken the lead position in computer science funding. Fully 86% of academic computer science research was funded by NSF—a much greater percentage than that of NSF funding in other science and engineering disciplines. The numbers of graduate students and young faculty were growing rapidly. It was a time to be optimistic about federal support for computer science. *Rising Above the Gathering Storm* (2007) and other reports had called for increased funding for federal agencies carrying out and supporting computer science research, and Congress seemed amenable.²

Wing entered CISE with two goals in mind—the first one largely logistical and strategic, the second one promoting some of her own ideas on collaboration:

One goal was to address some dissatisfaction the computer science community had with NSF. Some concerns were logistical in terms of the way processes and programs were run; some strategic, such as where the money was going. I wanted to help the computer science community in terms of streamlining some of the processes and ensuring that priorities were clear. I changed processes so people would submit fewer proposals, but on their best ideas. I also made sure the core areas of computer science were protected.

The second priority I had was in recognizing the expanse of computer science, and that the field itself needed to start reaching out to other disciplines, and thus other directorates at the National Science Foundation. I worked hard to collaborate and partner on friendly terms with my fellow Assistant Directors and Office Directors.³

Wing had inherited three major initiatives. The first involved the Cyber-enabled Discovery and Innovation (CDI) program. She recognized the importance of the Information Technology Research (ITR) program, and she saw CDI as a new large-funding opportunity that might have a similar impact.⁴ While CDI was to be a Foundation-wide initiative, she believed CISE should drive its intellectual agenda. The goal of this five-year initiative was to apply computational thinking to challenging problems in science and engineering research and education. The program hoped to find ways of drawing knowledge from data, understanding complexity in human and natural systems, and using virtual organizations to cross institutional, geographic, and cultural boundaries.⁵

The second initiative involved the Global Environment for Networking Innovations (GENI),⁶ the effort to build up a large-scale platform for carrying out

networking research. In addition to its inherent scientific value, Wing saw GENI as a way for CISE to break in to Major Research Equipment and Facilities Construction (MREFC) funding, a main source of funding for large capital items such as oceanic research vessels and major telescopes.

The third major initiative was the CCC, the Computing Community Consortium.⁷ In 2006, the Computing Research Association won the CISE solicitation and created CCC later that year. The mission of CCC is to “catalyze the computing research community and enable the pursuit of innovative, high-impact research.”⁸ The idea is to identify and articulate visions of computing research and align them with pressing national needs. CCC does this through workshops, white papers, and various means of communicating with their various stakeholders (government officials, the research community, funding agencies, and the public).

One of the first things that Wing did after arriving at CISE was to get an overview of funding, to see if for some areas it was too high or too low in relation to the level of activity and importance. In particular, she ensured adequate funding for core areas of computer science research such as programming languages and software engineering. Wing also immediately increased funding in cybersecurity. There had already been a program in the Computer and Networking Division, but she expanded it to be directorate-wide and then beyond the directorate into a program entitled Secure and Trustworthy Cyberspace.

The transition from the Bush to the Obama presidency occurred during Wing’s time at NSF. The Obama administration was highly supportive of NSF. Soon after President Obama took office, he floated a stimulus package (American Recovery and Reinvestment Act, 2009) and Congress approved it. It temporarily increased NSF funding by three billion dollars (a 50% budget increase). This increase caused significant organizational challenges, as Wing explains:

So the stimulus package gave NSF a lot of money, and we felt throughout the foundation an obligation to be, of course, very responsible in spending this money, but we had to spend it within those nine months. So, it was kind of crazy at NSF, because we decided that we wouldn’t use any of that money to increase staff in any way. We would give all that money to the PIs, to the research community. What that meant was the staff was doing 50% more work on top of their normal load, and that was very, very stressful for everyone at NSF.⁹

Several agencies banded together to use stimulus funds to enhance broadband in the United States. A CCC white paper called for broadband access for every citizen. Both the politics and the pace with which these funds had to be used proved to be very challenging.

Wing made several additional changes. One was to reorganize within CISE not only to support core areas of computer science, but also to fight against “silo-ing” that was common in computer science. As she explained:

First, I made a clear distinction between the core programs and what I called cross-cutting programs. The core programs were the programs within each of the divisions, algorithmic foundations, computing and networking foundations, and information and intelligence foundations. Then I created cross-cutting programs, each cut across the entire directorate because I saw that computing was too siloed and to solve future computing problems required expertise from people across the field. Networking is just one example, where we wanted to study not just computer networks but also social networks; and we wanted to support theoretical, not just experimental research in networking.¹⁰

Wing also reached out to other directorates to create joint programs. Examples included the Social and Computational Systems (SOCS) program and the program on Computer Science and Economics with the Social, Behavioral, and Economic Sciences (SBE) Directorate; and the Cyber-Physical Systems program with the Engineering Directorate.¹¹ As Wing said:

Autonomous cars. Robots at work, at play, at home. Intelligent, energy-efficient, earthquake-proof buildings. Physical infrastructure monitored and controlled by sensor nets. Embedded medical devices. Unobtrusive assistive technology. What is common to these systems? They have a computational core that interacts with the physical world. These cyber-physical systems are engineered systems that require tight conjoining of and coordination between the computational (discrete) and the physical (continuous). Cyber-physical systems are rapidly penetrating every aspect of our lives, with potential impact on sectors critical to U.S. security and competitiveness, including aerospace, automotive, chemical production, civil infrastructure, energy, finance, healthcare, manufacturing, materials, and transportation.¹²

Wing also reached out beyond NSF to create new cross-cutting initiatives. An example was the Health IT initiative that involved NSF and NIH.¹³ Four other Wing’s initiatives are worth noting:

- A multi-agency effort in robotics, which became a major initiative of the White House: “the National Robotics Initiative (NRI), an interagency program with NASA, NIH, and USDA that intends to develop the next generation of collaborative robots to enhance personal safety, health, and productivity.”¹⁴

- A data-intensive computing initiative, called Cluster Exploratory (CluE), in which NSF partnered with Google and IBM to make software and services running on a large cluster available to academic researchers.¹⁵ A second cluster was made available to academic researchers later, through an NSF partnership with Hewlett Packard, Intel, Yahoo!, and University of Illinois at Urbana-Champaign.¹⁶
- A data science initiative, which CISE intended as a follow-on to the CDI program and which became a White House push on big data. The White House Office of Science and Technology Policy under President Obama was interested in harnessing big data to advance discovery in science and engineering and to improve national security and education. The initiative involved DARPA and the Department of Energy as well as the Foundation.¹⁷
- Limiting submissions to core CISE programs to two proposals per year. This helped to reduce the burden of the NSF staff and also was well received by the community of Principal Investigators (PIs), who liked limiting their submissions to their best two ideas each year.

Wing was appreciative of Deborah Crawford, who served as Deputy AD, as well as Acting Deputy Assistant Director Gracie Narcho, Senior Science Advisor Suzi Iacono, and the CISE Division Directors (Sampath Kannan, Ty Znati, and Haym Hirsh).¹⁸

As CISE AD, Wing chaired the Networking and Information Technology Research and Development Committee (NITRD), which was the coordinating body for federal support of information technology.¹⁹ As chair, Wing felt she could influence the agenda for computer science research in other agencies. Suzi Iacono, who chaired the NITRD big data senior steering group, elaborated on the role of NSF and CISE in the activities of NITRD:

CISE is the central player in all of the NITRD interagency working groups, including the coordinating groups, and senior steering groups, and all that. All the other agencies are mission driven. And so their budgets are completely tied down, often years in advance. They have little freedom, little discretion, little autonomy. And at NSF and CISE we have all the discretion, all the autonomy, and we usually have a little bit of money. If you want to try something out you can take a million from this, or a half a million from there. You can, you know, get the support you need. When you have something to get started everyone else is jealous and wants to come here, and so we're leaders across all the agencies. And across all of the areas of computer science.²⁰

There was coordination across agencies in high-performance computing, in particular with DARPA and the DoE Office of Science. Wing also coordinated with the President’s Science Advisor to increase support for quantum information science. Other areas of interagency cooperation included health informatics, big data, and robotics.

Considering the declining enrollments in computer science (CS) in 2004, Wing originated the term *computational thinking* as shorthand for “the ways in which computer scientists think.” She urged that introductory CS courses, especially those for non-majors, teach computational thinking and not just programming. Computational thinking provided a way of addressing a wide range of problems, many of them far removed from computer science; thus, it was a tool that could help everybody. She used both funding (arranging for joint educational programs of CISE with HER and SBE) and her bully pulpit to promote computational thinking. She even characterized CDI as “computational thinking for all scientists and engineers.” Thus, she promoted the CDI program as a way to think about problems from other science and engineering disciplines—or as Wing argued:

it wasn’t so much about pushing computer science on everyone, it was more about pushing the ways in which we think. That was less threatening and less intimidating to their own fields. It was very important, actually, to use the term “computational thinking” for scientists and engineers as opposed to “computer science” for scientists and engineers. It’s subtle, but it was very important.²¹

5.3 Peter Arzberger (Acting AD, 2010 and 2011)

When Wing departed NSF, Peter Arzberger became the Acting AD for about half a year. He had received his bachelor’s degree in mathematics from the University of Massachusetts Amherst, and then worked for a few years in industry before returning to graduate school at Purdue to get a master’s degree and doctorate—also in mathematics, with a dissertation at the interface of biology and computing. He taught for several years at Rochester Institute for Technology and the University of Wisconsin–Madison before moving to NSF for the first time in 1988, as a program officer in the Mathematical and Physical Sciences (MPS) Directorate’s Probability and Statistics program. Later, he moved within NSF to the Division of Biological Infrastructure within the Biology Directorate. He was one of the founders of NSF’s computational biology program and was also one of the leaders in the cross-foundational activities in high-performance computing. He left NSF in 1995 to work at the San Diego Supercomputer Center and the University of California, San Diego. He served as the Executive Director of both the San Diego Supercomputer Center

and the National Partnership for Advanced Computational Infrastructure (NPACI). He returned to NSF in 2009 for a two-year position as Division Director of Biological Infrastructure. A year into that job, he became Acting AD of CISE for about half a year in 2010 and 2011. He returned to UCSD in 2013, and came back to NSF for a third time from 2013 to 2017. During this time, he worked in the Office of the Director and served as an interface to the National Science Board; then he moved to CISE, where he was a senior advisor and then the Acting Division Director for the Computer and Network Systems Division. After NSF, Arzberger became the Director of Life Science Initiatives at UCSD and Director of the National Biomedical Computation Resource.

It was in August 2010 that Arzberger received a call from Tom Peterson, who was the ED for Engineering and Acting Deputy Director of NSF, asking him to serve as Acting AD of CISE. Arzberger had some concern, given that he was not a computer scientist; but Peterson responded that he had sufficient understanding of computer science to complement his deep executive experience. Arzberger served for about half a year, until Farnam Jahanian arrived to serve as the CISE AD.

Because he knew that this was to be a short-term position, Arzberger took a narrow, focused view of his responsibilities as Acting AD. His focus was on certain education issues, laying groundwork so that Jahanian could enter into the job smoothly, and preparing and defending the CISE budget. Arzberger was concerned that many others within NSF were like him in not knowing much about what CISE did. He was surprised by the large number of partnerships CISE had with other parts of the Foundation. (At the time, CISE had a partnership with every directorate except Biology and the Geosciences.) Another major task for Arzberger was to devote significant time to communication: in particular, keeping ties in good standing with the other directorates. But it was also important, he believed, to have lots of communication with the CISE program officers, who he believed needed reassurance and a sense of continuity for several reasons: both Jeannette Wing and Deborah Crawford had left NSF shortly before he arrived, the three division directors were all brand new (two months was the longest tenure), and the appointment of permanent deputy division directors had not yet occurred. As Arzberger said, he “did a lot of walking the hallways.”²² To bring himself up to speed on CISE, he relied—as others before and since have done—on Acting Deputy AD Gracie Narcho and Senior Advisor Suzi Iacono. The new Division Directors (Sampath Kannan, Ty Znati, and Howard Wactlar) provided him with a helpful education in computer science.

Whereas Biology (and many of the other directorates) sold its program by talking about fundamental scientific problems it could address, in CISE the programs were

sold—effectively, Arzberger found—on their potential societal impact. In this way, Arzberger believed, CISE more resembled Engineering.

The one programmatic change Arzberger pointed to in his time as Acting CISE AD was to give greater emphasis to how CISE fit internationally. Like NSF more generally, CISE had never articulated clearly how it fit into the worldwide scientific community. While Arzberger was in Biology, he began conversations with NSF's international office on this issue. When he came to CISE, he updated a strategic plan for international activities for CISE that had first been worked on by Suzi Iacono half a dozen years earlier. The goal was to target scientific activities being done elsewhere in the world that might benefit CISE and the U.S. computer science community. For example, Japan had developed networking testbeds that it would be good for U.S. researchers to have access to. Other examples where international collaboration may be useful to U.S. researchers, Arzberger believed, included the cultural character of algorithms for facial recognition and the cultural dimensions of smart connected communities.²³

Arzberger remained at NSF for about three months after Jahanian arrived before returning to UCSD. This helped Jahanian get up to speed in his new job and eased the transition.

5.4 Farnam Jahanian (AD, March 2011 to July 2014)

Farnam Jahanian began his appointment as CISE AD on March 1, 2011. Born in Iran, he had lived there until he emigrated to the United States after completing high school. He received his undergraduate degree in computer science at the University of Texas at San Antonio, and his master's and doctorate in computer science from the University of Texas at Austin. His research interests are in distributed computing, network security, and network protocols and architectures. Early in his career, he worked for four years at the IBM T.J. Watson Research Center. He was on the faculty at the University of Michigan from 1993 to 2014, where he held a named professorship and served (1997–2000) as director of the Software Systems Laboratory and (2007–2011) as chair of the Computer Science and Engineering Department. He was also the co-founder of the Internet security firm Arbor Networks and was its CEO from 2001 to 2010. After leaving NSF in 2014, he served in rapid succession as vice president for research, provost, and president of Carnegie Mellon University.

Prior to coming to NSF as CISE AD, Jahanian had held both small and large grants from NSF and had served both as a reviewer and a panelist, not only for CISE and Engineering but also evaluating proposals for large-scale center

activities. Susan Graham and Fred Schneider from the CISE AD selection committee encouraged him to consider the position. As Jahanian explained:

[I]t was a tremendous opportunity not only to represent the community and push the agenda of the computer science and engineering community, but also a tremendous opportunity for establishing stronger connection, deeper connection, with all other disciplines, especially at the time where complication and data-intensive approaches [had] become so critical to scientific inquiry in just about every discipline. It was clear to me that it was a tremendous opportunity to put CISE in the middle of a number of conversations that impacted scientific inquiry just about in every other discipline.²⁴

When he arrived at NSF, Jahanian already knew three of the other ADs: Tim Killeen in Geosciences, Ed Seidel in Math and Physical Sciences, and Myron Gutman in Social, Behavioral, and Economic Sciences. It was also a period of time when the “[Obama] White House . . . was so supportive of investment in education and research.”²⁵ In particular, Jahanian worked with Tom Kalil from the Office of Science and Technology Policy, as well as Aneesh Chopra and later Todd Park, who each served successively as the White House Chief Technology Officer, to advance the President’s agenda in science and technology. Suzi Iacono, Gracie Narcho, and the Division Directors helped Jahanian get up to speed within CISE. He listened closely to the needs and interests of the computer science community, in part through the Computing Community Consortium (CCC), Computing Research Association, and the National Research Council’s Computer Science and Telecommunications Board. He also had a supportive NSF Director in Subra Suresh, who provided additional funds to CISE on a regular basis.²⁶

Jahanian was very supportive of the CCC workshops on topics of interest to CISE.²⁷ CCC had received its initial funding from NSF prior to Jahanian’s arrival. However, the CCC funding received a second grant during his term and subsequently a third award. Susanne Hambrusch, Division Director of Computing and Communication Foundations (CCF), indicates the function and importance of CCC to the computer science community and to CISE:

Farnam encouraged everyone to create new things. He was very open to ideas and trying out things. He was supportive of the CCC (Computing Community Consortium) running workshops on topics that NSF was interested in. These workshops provided great insight. . . .

The CCC was supposed to be for the community. It’s not for NSF. The CCC runs workshops and researchers can make proposals for visioning exercises. Based on these activities, reports are written. The CCC organizes events for

junior faculty, like the one giving researchers experience on how science policy is formed and what goes on. CCC does many things and they do them well.

I wished the CCC would have been better known to the community. NSF staff attends many of their events because they really are interesting. Before there was the Smart Health program, there was a CCC workshop on health. I think Eric Horowitz was one of the organizers. There were really good workshops and some of the topics turned into solicitations.²⁸

In addition to various internal programs, Jahanian worked closely with Office of Science and Technology Policy (OSTP) on three major initiatives.²⁹ *A Roadmap for US Robotics: From Internet to Robotics*, a CCC report written by academic and industry computer scientists in 2009, led to a national research agenda for robotics.³⁰ President Obama announced the National Robotics Initiative in June 2011. The National Institutes of Health, NASA, the Department of Defense, the Department of Agriculture, and other federal agencies also joined this initiative.³¹ The idea, as Jahanian describes it, was “a nationally coordinated program across multiple government agencies to develop the next generation of robotics and to advance capabilities and usability of such systems’ artifacts and to encourage existing and new communities to focus on core robots and new innovative applications.”³²

The second major initiative was the Federal Big Data Research and Development Initiative.³³ It arose from the meeting of a cross-agency senior steering group that OSTP had formed. NIH, the US Geological Survey, DARPA, and the Department of Energy all participated, but it was led by NSF and in particular by CISE. As Jahanian explains:

the focus of this was, of course, on [a] sort of investment framework to support the increasing importance in [the] role of data, not only in scientific exploration, but also in every sector of our economy. It has major thrust areas to it. Foundational research to develop new techniques and technologies to drive knowledge from data; new cyberinfrastructure to manage, curate, and serve data to research communities; new approaches for education; workforce development; and also, new types of interdisciplinary collaborations, grant challenges, and competitions.³⁴

The third of these initiatives was the U.S. Ignite program.³⁵ It was a joint initiative of OSTP and CISE. The goal was to build public-private partnerships to create “next generation application services that leverage advanced [high speed gigabit and wireless] networking.”³⁶ More than \$100 million was invested in this program.

A fourth cross-agency initiative, which CISE participated in but did not lead, was BRAIN (Brain Research through Advancing Innovative Neurotechnologies).

It was launched in 2013 by NSF, NIH, and DARPA. CISE participated in the initiatives funding the Computational Neuroscience program. This collaborative research was carried out in collaboration with Spain, France, Germany, Israel, and Japan.³⁷

One new multidimensional initiative, the Secure and Trustworthy Cyberspace (SaTC) program, involved not only every division of CISE but also four other directorates (Engineering, MPS, SBE, and EHR). NSF, and CISE in particular, had supported various cybersecurity initiatives in the past but not any that was this large, this far-ranging, or this costly (\$160 million in 2016 for the program across NSF). The initiative was led by Keith Marzullo in CISE's Computer and Network Systems Division.³⁸ As Jahanian explained, "the aim was to support fundamental scientific advances and technologies to protect cyber systems from malicious behavior, while preserving privacy and promoting usability. It had many components to it."³⁹ Other directorates investigated vulnerabilities in automotive systems and medical devices. Marzullo, director of the NITRD National Coordination Office, provided additional detail:

[W]e brought in Cyberinfrastructure, which at this point . . . was a separate office, because it's important to understand the infrastructure aspects of Cybersecurity as well as the need to protect our supercomputing capacity. We brought in Math and Physical Sciences, because there's a whole aspect to quantum computing and the deep math associated with that, and we brought in the social, behavior[al], and economic scientists, because if you were to look at the strategic plan that was published in 2010 on Cybersecurity by the Office of Science and Technology Policy, you'll see they called out for emphasizing economic incentives; and that meant we needed to bring in the Directorate of Social, Behavior[al], and Economic Sciences. And that was a partnership that grew. And so Secure and Trustworthy Cyberspace (SaTC), I think, was innovative in that we spent a lot of time, especially working with SBE (Social, Behavioral, & Economic Sciences), to forge new partnerships between computer scientists and social scientists, to try to move the needle on cybersecurity.⁴⁰

There were also new initiatives within CISE at this time. A new program provided funding for U.S. researchers working with researchers in Israel. It was the first instance in which NSF did not lead the reviewing process, inasmuch as Israel had high standards and NSF's legal office agreed to accept the Israeli reviews.⁴¹ There was an NSF/Intel Joint Partnership to fund science and technology centers.⁴² CISE created CyberSEES (Cyber-Innovation for Sustainability Science and Engineering) as part of the Foundation-wide SEES (Science, Engineering, and Education for Sustainability) initiative. As Jahanian explained at the time, CyberSEES:

focuses on the central role that computational and data-enabled approaches play in understanding and achieving sustainability. CyberSEES addresses the national priority of sustainability, an urgent and important area to ensure human needs are met equitably without harm to the environment or sacrificing the ability of future generations to meet their own needs.⁴³

The program was the result of a CISE 2011 workshop on the Role of Information Sciences and Engineering in Sustainability⁴⁴ and a National Research Council report entitled *Computing Research for Sustainability*.⁴⁵

XPS (Exploring Parallelism and Scalability) was created to enhance foundational research on parallelism when much of the high-performance computing research was application-focused. As Jahanian explained:

[the program] aims to address a central challenge created by the end of the exponential growth in microprocessor performance (a.k.a. Moore's law). While transistor density continues to scale, power dissipation levels that led processor performance leveled out. Our ability to achieve predictable performance improvements through traditional processor technologies has significant challenges. To avoid a crisis and to continue improving performance, we need a new era of computing driven by novel, groundbreaking research in all areas impacting parallel performance and scalability.⁴⁶

The XPS program was stimulated by a CCC study entitled *21st Century Computer Architecture*⁴⁷ and a National Research Council report entitled *The Future of Computing Performance: Game Over or Next Level?*⁴⁸

NSF created a new program to help protect the career ladder for scientists. With low reward rates in both the CAREER awards program and the core funding programs, it was hard for promising young researchers to win their first grant. The CRII program was targeted at junior faculty and had lower competition.

Looking back on FY 2014, Jahanian provided a snapshot of CISE: budget over \$850 million, 7,500 research proposals received, 1,500 research proposals awarded, 8,000 senior researchers supported, and 7,000 graduate and undergraduate students supported.

Jahanian made some organizational changes in CISE. Half of the CISE program officers and all of its division directors were rotators, and the importance of continuity and organizational memory were quite clear because, as mentioned earlier, both Jeannette Wing and Deborah Crawford had recently left NSF. As they had for Peter Arzberger, Suzi Iacono and Gracie Narcho stepped in to help, but it was clear that something structural needed to be changed. Jahanian introduced new permanent positions of deputy division directors, who could provide continuity to the

division directors (more likely to be rotators).⁴⁹ Jahanian also continued to make the pitch to the computer science community that much of the valuable work accomplished by NSF was done by people who volunteered to be a part of it, and he urged colleagues to serve on panels and come to Washington as rotators.

Also during Jahanian's watch, NSF Director Subra Suresh decided to move the Office of Cyberinfrastructure back into CISE, where it became a division.⁵⁰ Jahanian was surprised by but supportive of this change: "I think that brought a level of cohesiveness to computing research and infrastructure investment that . . . , in the long run, will benefit the country and benefit the science community."⁵¹ While some members of the computer science community might believe it inappropriate to have within CISE an infrastructure program that primarily supports the other science and engineering fields more than computer science, Jahanian disagreed. He pointed out that CISE was already supporting research in many technical areas (machine learning and data science, for example) that were driving transformations in other science and engineering disciplines, and that these initiatives were simultaneously changing and enriching the agenda of computer science research. Computer scientists were already collaborating with biologists on research projects that advanced both disciplines. So why should it be such a problem to have the cyberinfrastructure activities within CISE? Moreover, Jahanian believed that the Foundation was going to spend a lot of money on cyberinfrastructure, whether it was located in CISE or elsewhere in the organization, so why shouldn't CISE shape these investments? In fact, he noted, the big data initiative undertaken by NSF, and of interest to many computer scientists, was enabled by the cyberinfrastructure program.⁵² As he wrote at the time:

The goal is to more tightly couple foundational research in computing, communication, and information with advanced cyberinfrastructure; engage domain scientists to develop and deploy advanced cyberinfrastructure; use cyberinfrastructure to empower and enable knowledge environments and distributed collaboration; and address long-term sustainability of advanced cyberinfrastructure through cross-foundational and cross-institutional partnerships.⁵³

Jahanian also gave a new priority to communication with CISE's many stakeholders. He hired a communications director, the first person employed in such a role in any NSF directorate, and soon other directorates followed this hiring practice.⁵⁴

Jahanian strongly supported work on diversity and inclusion, which was being led by Jan Cuny.⁵⁵ He believed her work cut across all of CISE's divisions, and so he moved her office to be next to his in the CISE AD's office suite. He encouraged

her to build ties with Education and Human Resources. Her work not only helped to build national capacity for the training of computer scientists, in high demand, but also to build computing education for other scientific disciplines:

we started seeing there was a huge interest from almost every other discipline in having some level of computer science education, whether you're a biologist or an engineer or a social scientist. Everybody needs to have an understanding of computational approaches and data intensive approaches. So, again, with Jan's leadership, we developed programs with the EHR directorate to support computer science education for the broader community of the scientists and engineers out there.⁵⁶

The term of an outsider serving as AD is limited by law to four years, and Jahanian wanted the end of his term to coincide with the academic calendar. So, at the end of three and a half years, he left the Foundation and returned to academic life. He reflected:

I learned a lot in the process. I went into it thinking very differently about what government does and what agencies do and I came away with deep appreciation for public servants who serve the government, especially during tough years, where there was a lot of attack on agencies and on federal government in terms of the narrative about the effectiveness of the government itself. I came away with deep appreciation and gratitude for the work that these federal employees do. It's remarkable. I met so many incredibly smart people who are in the government, [contribute] day in, day out, contribute to this country, especially to the research and education mission of the country, to the science enterprise.⁵⁷

He was particularly proud of the fact that, during his term, employee satisfaction rose in the CISE Directorate to be higher than any other NSF directorate and almost any unit within the federal government.

5.5 Suzi Iacono (Acting AD, late 2014)

When Jahanian left the Foundation, Suzi Iacono became the Acting AD of CISE for five months. A social informatics scholar, she had earned her bachelor's and master's degree in social ecology from the University of California, Irvine, and her Ph.D. in information systems from the University of Arizona. Early in her career, she taught in the management school at Boston University, was a visiting scholar at the MIT Sloan School of Management, and was a research associate in the Public Policy Research Office at the University of California, Irvine. She came to NSF in 1998, first as a program officer in CISE's Division of Information and Intelligent

Systems (IIS). Over time, she held various appointments within CISE, including Deputy Assistant Director, and Acting Division Director in both IIS and Computer and Network Systems. She continues to work at NSF, and is currently the head of the Office of Integrative Activities.

Iacono's knowledge of the Foundation and how to get things done had served well in helping both Wing and Jahanian to get up to speed, and was again useful when Kurose arrived as AD. Iacono was a tireless contributor to and leader in NITRD committees, and in that role had made connections across NSF and the government that were very useful for Wing, Jahanian, and Kurose. More generally, her connections across the Foundation were particularly useful as she managed CISE's ongoing cross-directorate initiatives such as those in cyber-physical systems, trustworthy cyberspace, robotics, and big data.

5.6 James Kurose (AD, January 2015–)

The next CISE AD was James Kurose. His background included an undergraduate degree in physics from Wesleyan University in 1978 and a Ph.D. in computer science from Columbia University in 1984. A researcher in the area of computer networks, he joined the University of Massachusetts Amherst the year he graduated from Columbia, and moved up the ranks to Distinguished Professor. He has served as department chair and dean. He also spent a year at IBM Research (1990–1991) and at INRIA and EURECOM in Sophia Antipolis, France (1997–1998). He has won various awards, including the IEEE's Taylor Booth Award and the ACM's SIGCOM Lifetime Achievement Award. He joined NSF as the Assistant Director of CISE in 2015.⁵⁸

Kurose has had a long-standing relationship with NSF. He served on the CISE Advisory Committee when Jahanian was AD. He received almost continuous funding from NSF from the time he graduated from Columbia until he joined NSF. He was the co-PI on one of the large Engineering Research Center grants for Collaborative, Adaptive Sensing of the Atmosphere. He had run several workshops on behalf of NSF in the areas of network research testbeds, undergraduate computing education, and persuasive computing and communications collaborations with India. He had also received funding from other agencies, including DARPA and ONR. From these various activities and his role on the Computing Research Association board, he believed he had a good sense of the state of computing research at the time he joined NSF.

One of the important issues that Kurose faced as AD was the management of the cyberinfrastructure program, which included high-performance computing, net-

works, software, data, and people.⁵⁹ Over the years, it had moved in and out of CISE several times. When Kurose arrived, cyberinfrastructure was back in CISE, and some within the cyberinfrastructure community worried that CISE would use cyberinfrastructure resources for more traditional computer science projects. NSF Director France Cordova initiated a review. The results, available in 2016, indicated that cyberinfrastructure had been well cared for in CISE. To enhance access to decision-makers for the head of the cyberinfrastructure program, it was turned into an Office of Advanced Cyberinfrastructure and its head was invited to the weekly Senior Management Round Table with all of the Assistant Directors to carry out overall planning for the Foundation. Kurose felt very fortunate that he had been dean at University of Massachusetts Amherst when his school, several other universities, and the state of Massachusetts joined together to create the Massachusetts Green High-Performance Computing Center. Through this involvement, Kurose had already learned about many of the issues surrounding high-performance computing and already knew some of the NSF people involved with cyberinfrastructure.

Kurose, like Peter Freeman before him with GENI, has worked to convince the National Science Board to broaden the Major Research Equipment and Facilities Construction (MREFC) program. The National Science Board (NSB) has modified the rules about MREFC, so that a project as small as \$70 million dollars would qualify for these funds—rather than the previous rule that a project had to be at least 10% of the directorate’s budget, so approximately \$100 million for CISE. Nevertheless, no MREFC funds have yet been awarded to cyberinfrastructure (as Kurose requested) or any other CISE-related activities. Both Kurose and MPS AD Fleming Crim continue to push for these changes.

When Kurose arrived at NSF, he had some personal priorities. He wanted to build K-12 computing education to enhance undergraduate computer education. In the networking area, he wanted to move beyond GENI and backbone issues to examine Platforms for Advanced Wireless Research (PAWR). He also wanted to build more partnerships with industry. He understood, however, that “if there was just no interest in the community [for] doing that and the program directors didn’t want to do that, I wasn’t going to force them to do that.”⁶⁰

The initiative with industry was driven by the observation that U.S. federal investment in research was “pretty flat” at 3% of GDP. So Kurose was particularly interested in building long-term research relationships between NSF and industry. CISE engaged with VMWare, the Semiconductor Research Corporation (SRC), and (on at least four projects) Intel. Jahanian had already begun conversations with Intel that led to these joint initiatives, and Kurose continued them. There had already been a long history of working with SRC.

Kurose had been advised by Jahanian to be particularly careful with four activities: rolling out a budget, testifying before Congress, speaking to the National Science Board, and answering inquiries from the Inspector General. It was a trial by fire for Kurose; in his first six weeks, he had done the first three. He has since had some limited contact with the Inspector General, and he has regular contact with the NSB because of the importance of cyberinfrastructure issues to the entire national science and engineering enterprise.

While at NSF, Kurose came to believe that computer science is special among the scientific and engineering disciplines for its centrality to the scientific and engineering enterprise as well as having its own intrinsic value. He found that while there are “amazing questions” and “intellectual challeng[es]” about computing research, there is also a national interest issue here. CISE and the Engineering Directorate both have a particular role to play in building the national economy.

In his FY 2016 budget request, Kurose emphasized a number of cross-cutting investments, while still supporting the core of computer science research as well as advanced cyberinfrastructure. These cross-cutting activities are listed in Table 5.1.

There were also requests for two new cross-cutting initiatives. Innovations at the Nexus of Food, Energy and Water Systems (INFEWS) supports research on the safety and security of food, energy, and water resources. INCLUDES, or Inclusion across the Nation of Communities of Learners that have been Underrepresented for Diversity in Engineering and Science, aims “to develop a scalable, national initiative to increase the preparation, participation, advancement, and potential contributions of those who have been traditionally underserved and/or underrepresented in the STEM enterprise.”⁶¹

Kurose praises his program officers and division directors. He points to their dedication and responsiveness to the research community. Many of the good ideas and good programs bubble up from these professional staff members. Budget issues and workload issues for the professional staff are a major concern since the numbers of computer science students have *tripled* over the past decade, and faculty members educating them are seeking expanded funding from CISE.⁶² The biggest difficulty in attracting high-quality people to serve as program officers and division directors, Kurose believes, is the difficulty of coming to work in Washington, DC, which disrupts the daily life of a professor’s home and work—especially for people who are located far from Washington.

One of Kurose’s strategies has been to get Office of Management and Budget approval for major new initiatives, which will result in “passbacks” with new funding dedicated to each initiative. He pointed to the numerous major funding

Table 5.1 Ongoing cross-cutting investments in the FY 2016 CISE budget

Program Name	Partners
Secure and Trustworthy Cyberspace	Education and Human Resources, Engineering, Mathematical and Physical Sciences, and the Social, Behavioral, and Economic Directorates
Cyber-Physical Systems	Engineering Directorate, the Department of Homeland Security, the Department of Transportation, the National Aeronautics and Space Administration, and the National Institutes of Health
National Robotics Initiative	Engineering, Education and Human Resources, and the Social, Behavioral, and Economic Directorates as well as with the DARPA, NASA, and USDA
Critical Techniques and Technologies for Advancing Big Data Science and Engineering	all the NSF directorates
Smart and Connected Health	Engineering and Social, Behavioral, and Economic Directorates, as well as with NIH

initiatives that were supported during the Obama Administration: the National Big Data Research Initiative,⁶³ the National Robotics Initiative, the Advanced Wireless Research Initiative, CS for All, the Smart Cities Initiative,⁶⁴ the National Strategic Computing Initiative, and the Brain Research through Advancing Innovative Neurotechnologies Initiative. Kurose himself started the Smart and Connected Communities (S&CC) program (building on the Smart Cities initiative) and the Smart and Autonomous Systems (S&AS) program.

Kurose used various mechanisms to solicit community feedback and suggestions about new research directions and new programs. CISE awards funding for 50 workshops each year. For example, the Smart Connected Communities initiative was a result of one of these workshops. Additionally, CCC holds five to ten workshops each year, including an influential one on privacy during Kurose's first two years as CISE AD. There are also various National Academy studies, many of them conducted by the National Research Council's Computer Science and Telecommunications Board—often funded by CISE. A 2017 study on IT and automation, for example, was led by computer scientist Tom Mitchell from Carnegie Mellon and economist Erik Brynjolfsson from MIT.⁶⁵ Feedback also comes through Kurose's

own travel. He visits one or two universities every month. There is also advice from the 20-member CISE Advisory Committee. For example, this group has provided strong input about CISE's role in the NSF-wide big data initiative.

Kurose acknowledged a 3-inch-thick briefing book that senior CISE staff had prepared for him, his reliance on Deputy AD Erwin Gianchandani, and the support he received in particular from Suzi Iacono and Peter Arzberger in getting up to speed and being effective in his job as AD.⁶⁶ Iacono taught Kurose the rigorous regulations for a senior government manager.

With the growing size of the computing research community, there is pressure on NSF programs and very low success rates for applications—sometimes only 5%. Kurose lauds the CRII program established by Jahanian.⁶⁷ Kurose believes that programs with 5% success rates are not sustainable, and he has eliminated programs with low success rates and instead directed faculty members to compete for grants in these areas through general research solicitations.

5.7 Conclusions

We are chary to draw strong conclusions about such recent events as are covered here because it is hard to get historical perspective. During this era, CISE was already a well-established directorate, but it continued to grow as the place of computing in American society continued to grow, the numbers of students wanting to study computing skyrocketed, and the computer science faculty who would teach these students and contribute to basic and applied computing research also grew. The centrality of computing to the wide swath of science and engineering disciplines was becoming increasingly clear, and this fact was reflected in the large increase of programmatic partnerships of CISE with other NSF directorates, other federal agencies, industry, and institutions in other countries.

Some of the notable programmatic efforts that were initiated or continued during this period were Global Environment for Network Innovations (GENI), Cyber-enabled Discovery & Innovation (CDI), various cybersecurity initiatives, Broadening Participation in Computing and the new INCLUDES initiative, and programs involving partnership with other directorates (e.g., Secure and Trustworthy Cyberspace), across agencies (e.g., the National Robotics Initiative), and with industry (e.g., the Cluster Exploratory). There were also efforts to improve the career pathway for young scholars (CRII), the process for applying for funding (limits on number of proposals that could be submitted), and the means for the computer science community to bring good ideas to NSF to be shaped into fundable programs (CCC).

Notes

1. Oral history, Deborah Crawford, interviewed by William Aspray, June 20, 2017. Charles Babbage Institute. Chapter 4 in this volume describes Peter Freeman's initiatives while he was AD at the start date for this narrative.
2. National Academy of Sciences, National Academy of Engineering, and Institute of Medicine. 2007. *Rising Above the Gathering Storm: Energizing and Employing America for a Brighter Economic Future*. Washington, DC: The National Academies Press. DOI: [10.17226/11463](https://doi.org/10.17226/11463).
3. Oral history, Jeannette Wing, interviewed by Rick Adrion, March 14, 2017. Charles Babbage Institute.
4. Peter Freeman wrote the first draft of this program while he was at CISE as AD.
5. The program announcement for CDI can be found at <https://www.nsf.gov/pubs/2011/nsf11502/nsf11502.htm>.
6. See Chapter 9 for more information about GENI; also see: P. A. Freeman. November 2005. NSF/CISE plans GENI initiative. *Computing Research News*, 17(5), https://cra.org/crn/2005/11/nsf_cise_plans_geni_initiative/; P. A. Freeman. March 2006. GENI and your research. *Computing Research News*, 18(2), https://cra.org/crn/2006/03/geni_and_your_research/; P. A. Freeman. November 2006. Moving forward strategically. *Computing Research News*, 18(5), <http://archive.cra.org/CRN/articles/nov06/freeman.html>.
7. This initiative was first developed by Peter Freeman and Deborah Crawford. In our 2017 interview, Marzullo discussed the importance of the CCC to the NSF program staff as they crafted new programs. Oral history, Keith Marzullo, interviewed by William Aspray, July 20, 2017. Charles Babbage Institute.
8. About the CCC, <https://cra.org/ccc/about/>.
9. Wing interview, 2017, *op. cit.*
10. Wing interview, 2017, *op. cit.*
11. There was increasing collaboration during the Wing years between CISE and HER on broadening participation issues. The interview with Jan Cuny (2017), who led the broadening participation efforts within CISE since 2004, provides an overview of these efforts. Oral history, Jan Cuny, interviewed by Peter Freeman, March 22, 2017. Charles Babbage Institute.
12. J. M. Wing. 2009. Cyber-physical systems. *Computing Research News*, 21(1). https://cra.org/cm/2009/01/cyber-physical_systems/.
13. Oral history, Howard Wactlar, interviewed by Peter Freeman, March 15, 2017. Charles Babbage Institute. Here Wactlar briefly discusses both the health and robotics initiatives.
14. F. Jahanian. March 2012. Highlights of the CISE fiscal year 2013 budget request. *Computing Research News*, 24(2). https://cra.org/crn/2012/03/highlights_of_the_cise_fiscal_year_2013_budget_request/.
15. J. M. Wing. March 2008. Data-intensive computing. *Computing Research News*, 20(2). https://cra.org/crn/2008/03/data-intensive_computing/.
16. CRA. November 2008. *CISE Bytes*. 20(5).

17. See G. Gross. March 2012. White House launches big data R&D push. *Computerworld*, vol. 30. https://www.computerworld.com.au/article/420083/white_house_launches_big_data_r_d_push/.
18. Division Director Keith Marzullo, who served in Jahanian’s era as AD, also emphasized the importance of Iacono and Narcho—as well as long-time program officer Anita LaSalle—in providing continuity for CISE and mentoring new staff (Marzullo interview, 2017). For additional background, see: Oral history, Gracie Narcho, interviewed by Peter Freeman, March 15, 2017. Charles Babbage Institute.
19. Starting during the George W. Bush Administration, the CISE AD served as chair and the NCO director served as co-chair.
20. Oral history, Suzi Iacono, interviewed by Peter Freeman, March 7, 2017. Charles Babbage Institute. George Strawn, who was working for NITRD during the Wing era, also briefly discusses NITRD’s and NSF’s role in the Next Generation Internet initiative. (See Oral history, George Strawn, interviewed by William Aspray, July 19, 2017. Charles Babbage Institute.)
21. Wing interview, 2017, *op. cit.*
22. Oral history, Peter Arzberger, interviewed by William Aspray, July 21, 2017. Charles Babbage Institute.
23. Arzberger interview, 2017, *op. cit.*
24. Oral history, Farnam Jahanian, interviewed by William Aspray, July 31, 2017. Charles Babbage Institute.
25. Jahanian interview, 2017, *op. cit.*
26. Oral history, Susanne Hambrusch, interviewed by Rick Adrion, September 20, 2017. Charles Babbage Institute.
27. Hambrusch interview, 2017, *op. cit.*
28. Hambrusch interview, 2017, *op. cit.*
29. One of the advantages that NSF had over some of the other federal agencies, such as Transportation or NASA, was NSF’s much closer connections to the academics who serve as principal investigators. See Marzullo interview, 2017, *op. cit.*, for a discussion of this point.
30. <http://www.us-robotics.us/reports/CCC%20Report.pdf>. This was a follow-on to the work of the congressional robotics caucus, which had been formed in 2007. (See A. K. Noor, *Robotics: Present State, Future Trends*, ASME, <https://www.asme.org/engineering-topics/articles/robotics/robotics-present-state-future-trends>.)
31. See news release 11-129. June 24, 2011. “NSF Leads Interagency Collaboration to Develop Advanced Robotics,” https://www.nsf.gov/news/news_summ.jsp?cntn_id=119911; and “National Robotics Initiative 2.0: Ubiquitous Collaborative Robots,” program solicitation, https://www.nsf.gov/funding/pgm_summ.jsp?pims_id=503641&org=CISE.
32. Jahanian interview, 2017, *op. cit.*
33. See news release 12-060. March 29, 2012. “NSF Leads Federal Efforts in Big Data,” https://www.nsf.gov/news/news_summ.jsp?cntn_id=123607&org=NSF&from=news. Also see the program solicitation, “Critical Techniques, Technologies and Methodologies for Ad-

- vancing Foundations and Applications of Big Data Sciences and Engineering,” https://www.nsf.gov/funding/pgm_summ.jsp?pims_id=504767.
34. Jahanian interview, 2017, *op. cit.*
 35. See news release 12-112. June 13, 2012. “NSF Leadership in Discovery and Innovation Sparks White House U.S. Ignite Initiative,” https://www.nsf.gov/news/news_summ.jsp?cntn_id=124472; as well as the Dear Colleague Letter: “U.S. Ignite: The Next Steps,” https://www.nsf.gov/publications/pub_summ.jsp?ods_key=nsf12085.
 36. Jahanian interview, 2017, *op. cit.*
 37. See the program announcement at https://www.nsf.gov/funding/pgm_summ.jsp?pims_id=5147. Also see: F. Jahanian. September 2013. CISE looks ahead to 2014. *Computing Research News*, 25(8), https://cra.org/crm/2013/09/cise_update_the_need-to-know_for_fy_2014/.
 38. See Marzullo interview, 2017, *op. cit.*, for comments on both the Trustworthy Computing program and another cross-cutting program in Cyber-Physical Systems, which involved, for example, work with mechanical engineering and systems engineering. While Marzullo and other NSF staff saw the importance of the traditional core programs, “we saw a growing importance, a growing interest in the cross-cutting programs as compared to the core programs . . . There was a general feeling that we should be trying to do more investments that brought in more than [the] one domain that was represented by our division” (Marzullo 2017).
 39. Jahanian interview, 2017, *op. cit.*
 40. Marzullo interview, 2017, *op. cit.*
 41. Other international programs included WiFiUS (Wireless with Finland and the U.S.) and a partnership with Japan on networking. These were examples of NSF Director Suresh’s SAVI (Science Across Virtual Institutes) program. See Marzullo interview, 2017, *op. cit.*
 42. Another collaboration with industry during this period was with the Semiconductor Research Corporation, a consortium of semiconductor companies. CCC ran a workshop on hardware cybersecurity at which NSF and the semiconductor companies worked together to identify problems that were of importance to both academic and industrial researchers (Marzullo 2017).
 43. F. Jahanian. November 2012. Highlighting opportunities for the CISE community. *Computing Research News*, 24(5). https://cra.org/crm/2012/11/highlighting_opportunities_for_the_cise_community.
 44. See <http://archive2.cra.org/ccc/visioning/visioning-activities/sees-it>.
 45. National Research Council. 2012. *Computing Research for Sustainability*. Washington, DC: The National Academies Press.
 46. F. Jahanian, November 2012, *op. cit.*
 47. Published May 25, 2012. <https://cra.org/ccc/wp-content/uploads/sites/2/2015/05/21stcenturyarchitecturewhitepaper.pdf>.
 48. National Research Council. 2011. *The Future of Computing Performance: Game Over or Next Level?* Washington, DC: The National Academies Press.

49. Howard Wactlar, the Division Director for CISE's IIS Division, described his first Deputy Division Director, Debbie Lockhart, an experienced permanent employee who had previously worked in the MPS Directorate: "She was a tremendous help in taking care of all of these administrative things and keeping me doing the right thing at the right time when it was due. She gave me the freedom to go try these new initiatives and spend as much time as I needed outside" (Wactlar interview, 2017, *op. cit.*).
50. F. Jahanian. March 2013. CISE welcomes the Division of Advanced Cyberinfrastructure. *Computing Research News*, 25(3).
51. Jahanian interview, 2017, *op. cit.*; also see Hambrusch interview, 2017, *op. cit.*
52. Jahanian interview, 2017, *op. cit.*
53. F. Jahanian. March 2013, *op. cit.* The documentation of this decision can be found in reports of the NSF-wide Advisory Committee for Cyberinfrastructure: April 2011, <http://www.nsf.gov/od/oci/taskforces/>; Cyberinfrastructure for 21st Century Science and Engineering Advanced Computing Infrastructure Vision and Strategic Plan, February 2012, <http://www.nsf.gov/pubs/2012/nsf12051/nsf12051.pdf>; and Revolutionizing Science and Engineering Through Cyberinfrastructure: Report of the National Science Foundation Blue-Ribbon Advisory Panel on Cyberinfrastructure, January 2003, <http://www.nsf.gov/od/oci/reports/atkins.pdf>.
54. On Jahanian establishing deputy division director positions and hiring a communication person, see Hambrusch interview, 2017, *op. cit.*
55. For more background, see Cuny interview, 2017, *op. cit.*
56. Jahanian interview, 2017, *op. cit.*; also see Hambrusch interview, 2017, *op. cit.*, and Marzullo interview, 2017, *op. cit.*
57. Jahanian interview, 2017, *op. cit.*
58. P. Harsha. October 2014. UMass professor and CRA board member Kurose selected to run NSF CISE. *Computing Research News*, 26(9). <https://cra.org/govaffairs/blog/2014/09/umass-amherst-prof-and-cra-board-member-kurose-selected-to-run-nsf-cise/>.
59. For his early remarks about cyberinfrastructure, see: J. Kurose. January 2015. An early greeting from CISE. *Computing Research News*, 27(1). https://cra.org/crn/2015/01/an_early_greeting_from_cise/. Kurose pointed to National Research Council, *Future Directions for NSF Advanced Computing Infrastructure to Support U.S. Science and Engineering in 2017–2020: Interim Report*. Washington, DC: The National Academies Press, 2014.
60. Oral history, James Kurose, interviewed by William Aspray, April 1, 2017. Charles Babbage Institute. Also see E. Gianchandani and G. Jochum. May 2015. NSF's CISE pushing beyond today's Internet. *Computing Research News*, 27(5). https://cra.org/crn/2015/05/nsfs_cise_pushing_beyond_todays_internet/.
61. J. Kurose. 2015. The FY 2016 budget request for NSF computer and information science and engineering. *Computing Research News*, 27(2), https://cra.org/crn/2015/02/the_fy2016_budget_request_for_nsf_computer_and_information_science_and/. On the following two years' requests, see: J. Kurose. March 2016. Highlights of the President's FY 2017 budget request for CISE. *Computing Research News*, 28(3), <https://cra.org/crn/2016/03/highlights-presidents-fy2017-budget-request-cise/>; J. Kurose. June 2017. The President's FY 2018

budget request for CISE. *Computing Research News*, 29(6), <https://cra.org/crn/2017/06/presidents-fy2018-budget-request-cise/>.

62. Kurose interview, 2017, *op. cit.*, also points out that deans, when seeing these growing enrollments and familiar with the cyclical boom and bust cycles in computing enrollments, are reluctant to make long-term commitments to tenure-track lines and thus are more likely to hire instructors. However, Kurose believes that the situation is different this time—in part because of the increasing centrality of computing in other disciplines—and the growth will be sustained.
63. See C. Baru. March 2015. NSF and the national big data initiative. *Computing Research News*, 27(3), https://cra.org/crn/2015/03/nsf_and_the_national_big_data_initiative/.
64. See J. Kurose. October 2015. NSF/CISE plays leadership role in new federal smart cities initiative. *Computing Research News*, 27(9), <https://cra.org/crn/2015/10/nsfcise-plays-leadership-role-in-new-federal-smart-cities-initiative/>.
65. National Academies of Sciences, Engineering, and Medicine. 2017. *Information Technology and the U.S. Workforce: Where Are We and Where Do We Go from Here?* Washington, DC: The National Academies. DOI: [10.17226/24649](https://doi.org/10.17226/24649).
66. Kurose interview, 2017, *op. cit.*
67. Kurose interview, 2017, *op. cit.*



PART

**SELECTED
SUBJECT
STUDIES**

Pre-CISE Computing Facilities and Education Programs¹

William Aspray

This chapter discusses computing facilities and education programs in the era prior to the formation of CISE in 1986. As we will see in Chapter 7, NSF was supporting research in computer science, especially through the Math and Physical Science (MPS) Directorate. However, NSF's major early contribution to computing was in the provision of computing facilities (for both research and education) and education programs of several sorts. These computer facilities and computer education programs were important in the early years of modern computing. In the end, however, NSF did not have the resources to meet the national demand for either computing facilities or computer education.

6.1 Facilities Program

This section describes the computer facilities programs conducted by the National Science Foundation in the years prior to the formation of the Computer Science and Engineering Directorate from 1959 to 1986. From 1950 to 1967, the facilities program was arguably the most important contribution the Foundation made to the computing field—more important than its direct contributions to computing education or computing research.

Already by the mid 1950s, the computer was a versatile tool used in the physical and biological sciences, engineering, and the social sciences. Scientists and engineers used computers to find exact and approximate solutions to problems; model complex structures; organize, analyze, and present data; and control laboratory equipment. In the educational realm, computers delivered instruction in various

subject areas; and computers enabled teachers to introduce students to scientific subjects before the students were able to handle the subject's full complexity, e.g., through the use of models or statistical packages. Universities used computers for administrative purposes, like any other business or organization. Unfortunately, computers were then high-capital items, not readily within the financial means of most colleges and universities. The Foundation became the principal federal agency providing computers to institutes of higher learning.

By 1971, NSF provided 233 computing center facilities grants to institutions in all 50 states. Many types of institutions benefited, including community colleges, private and public colleges, Ivy Leagues, women's colleges, historically black colleges and universities, liberal arts colleges, and research and technology universities.²

6.1.1 University Computing Centers

In 1953, the Foundation surveyed the status of applied mathematics in the United States. It concluded that the computer was causing "an unprecedented mathematization; not only of fundamental scientific research in the physical and biological sciences but also in the management of our industrial and social systems."³ The computer resembled other large-scale facilities, such as nuclear reactors and wind tunnels, both accelerating research and causing researchers to cross traditional scientific and engineering boundaries.⁴

At the time, computers were scarce and expensive. In 1954, there were only 20 computers in the United States, and only four commercial models were available for rent or purchase.⁵ Universities could not afford them.⁶ In response, in May 1955, the National Science Board decided the Foundation should help provide computing facilities to the nation's colleges and universities, with five small grants made the next year to support university computation centers and research in numerical analysis.⁷ A formal program for computing equipment was established in 1959, although various programs around the Foundation received proposals to fund computing equipment before that. For example, the Mathematical, Physical, and Engineering Sciences Division (MPES) received 19 facilities proposals in 1958, most of them for computing equipment. MPES was able to support three of them for a total of \$200,000.

During the second half of the 1950s, computing spread rapidly across university campuses, and by 1959, 150 colleges and universities were conducting some kind of computing activity. Computers were used in research involving linear programming, game theory, automata theory, artificial intelligence, adaptive mechanisms, psychometrics, neural psychology, learning machines, information theory, coding theory, statistics, cybernetics, and a wide range of modeling techniques.⁸

The largest donor of computing facilities to colleges and universities in the United States during the 1950s was not any federal agency, but IBM. By 1959, IBM had donated small Model 650 computers to more than 50 schools and had provided larger computers (Models 704 and 709) to several universities. While IBM had a strong charitable sensibility, it was also good business practice to help the universities to train IBM's future workforce and the workforce of IBM's customers. However, one academic, Louis Fein of Stanford University, expressed concerns about the IBM program: machines were sometimes awarded without a strong computer curriculum in place, many schools assigned unqualified instructors just to obtain a free computer, and there was little consideration given to computing's theoretical foundations. Other computer manufacturers—including Burroughs, Sperry Rand, Bendix, and Royal McBee—also had university donation programs similar to that of IBM, but smaller.⁹

The growth in campus computing grew unabatedly in the 1960s. Academic computing facilities grew from 100 in 1961, to 300 in 1963, to 700 in 1965, to over 2,000 (at 160 institutions) in 1969.¹⁰ Between 1959 and 1971, the Foundation awarded 414 computing facilities grants—about equally divided between first computer acquisitions and equipment improvements to established computing centers. In 1959, the first official year of the computing facilities program, the Foundation awarded \$1.5 million in matching grants, spread across five institutions, which each received between \$100,000 and \$500,000. Preference was given to universities that were able to provide substantial funding themselves. The Mathematics budget and some general funds were used to provide almost a million additional dollars to 13 other schools for installation, rental, and operating costs of computing centers.¹¹

As early as the second year of the computing facilities program, 1960, the Foundation recognized its funding could not keep up with the growth in academic computing.¹² It considered—but dismissed for the time being—funding regional centers. With IBM support, MIT had been providing computing facilities to a number of colleges across New England since 1957.¹³ However, the Foundation staff decided that the regional center idea was deficient by both not providing adequate access to researchers and not providing hands-on instructional experience. Nevertheless, the Foundation revisited and implemented a regional centers program only a few years later.¹⁴

After setting aside the idea of regional centers, NSF supported computing facilities for individual colleges as best it could.¹⁵ MPES anticipated the total need for 1962 to be \$10 million, and a program was established in 1962 within Mathematical Sciences to handle these computing facilities requests. As more colleges and universities opened computing centers, and as the computing centers at large universities installed mainframe computers, the projections of national need rose

dramatically, ranging from \$30 million to \$300 million.¹⁶ Despite this demand, the budget increased much more slowly. In 1961, the Foundation awarded 6 grants totaling \$1.6 million from the facilities budget and another 20 small grants totaling \$796,000 from other funds.¹⁷ Between 1961 and 1967, the computing facilities budget grew from approximately \$2.4 million to \$11.3 million. 1967 represented the end of the era, with computing facilities grants dropping steeply thereafter. In 1967, the Foundation awarded 214 grants, mostly awards between \$20,000 and \$200,000, to expand existing computing facilities. There were a few larger grants: Case Western Reserve, Cornell, North Carolina-Chapel Hill, Pittsburgh, Princeton, Purdue, Washington-Seattle, and Yale each received between \$500,000 and \$700,000, while California-Berkeley, Texas Christian, and Wisconsin-Madison each received between \$1,000,000 and \$1,700,000. With the creation in 1967 of the Office of Computing Activity and the refocus on computer education over facilities, the Foundation budget for computing facilities dropped to \$6.5 million in 1968.

6.1.2 Instructional Use of Computers

In the mid-1960s, universities were rapidly increasing their instructional use of computers. Jerome Wiesner, the Dean of Science at MIT, estimated that instructional use was growing at twice the rate of research use of computers.¹⁸ Between 1964 and 1968, the average cost for operating an academic computing center doubled. Moreover, the needs of research users and instructional users of computers were different: researchers frequently needed high-speed processing, large memories, specialized input-output equipment, and data converters; while instructional use required facilities that could handle large numbers of small programs with rapid turnaround time.¹⁹ By 1968, nearly all universities, one-third of four-year colleges, one-fourth of junior colleges, and one-fifth of high schools could offer students access.²⁰

Although the Foundation staff had set aside the idea of regional computing centers, the Office of Computing Activities (OCA) advisory committee returned to the idea almost immediately. One reason was simply to provide economies of scale given the great need and limited budgets. Another reason was to advance OCA's educational mission, as "a way of distributing not just computing power, but intellectual power within regions," and "as a way of transferring knowledge from the university all the way down to the secondary school and as a support for educational innovation." A third reason was technological—to uncover system and hardware problems with different network topologies.²¹ A fourth reason was political: "Also it sounded good . . . regional computing. You know, that appeals to every Congressman."²²

To test the feasibility of regional computing centers, OCA in 1968 funded 10 pilot projects involving 8 universities, 82 colleges, and 23 secondary schools located across 25 states. The Foundation provided \$4 million, representing two-thirds of the total cost. These pilot projects sought to evaluate regional networks comparing service delivery costs, computer programming language choices, curriculum development, and dissemination methods as well as identifying institutional barriers to cooperation.²³

OCA was happy with the pilot program, and Milton Rose, the head of OCA, hoped to create an additional 10 regional centers each year.²⁴ The Foundation did fund 30 regional centers between 1968 and 1973.²⁵ The centers trained faculty members to use computer services in their science teaching and provided remote access through terminals and telephone lines to schools without computers. The program made large computers, extensive program libraries, experienced computer service staffs, and inexpensive computing time available to a broad community of academic users.²⁶ Nearly 350 institutions participated in the regional centers program, including 26 universities, 240 colleges, 40 junior colleges, and 40 high schools.²⁷

Three of the regional networks were particularly successful.²⁸ One was the Triangle Universities Computation Center, formed in 1965 with grants to the University of North Carolina at Chapel Hill, North Carolina State University, and Duke University. The three universities formed a jointly owned, independent corporation to provide computing services from a single computing center in the Research Triangle research park. The follow-up grant, as part of the regional centers program, supported the North Carolina Educational Computing Service, which provided computing services to smaller public and private schools throughout the state.²⁹ A second successful regional network was the MERIT network, which linked the computing centers at the University of Michigan in Ann Arbor, Michigan State University in East Lansing, and Wayne State University in Detroit. The network later became one of the managing agencies of the National Science Foundation Network (NSFNET). A third major success was the New England Regional Computing Center (NERComp), originally formed through an agreement between IBM and MIT, and later funded through the Foundation's regional centers program.³⁰

While regional centers provided a model, the program had some limitations. Small institutions were frequently unable to pay for the services of a regional computing center once the generous subsidy from the Foundation ended. In some cases, even during the period of Foundation support, some aspiring computing centers were more interested in developing powerful and novel services than in supporting the mundane needs of users at smaller institutions.³¹

A couple of years after the regional centers program was initiated, OCA proposed a plan to create centers specializing in the computational needs of particular disciplines, such as theoretical chemistry, or particular national social problems, such as environmental pollution, transportation, and communication.³² The Foundation developed this plan partly in response to the wishes of the Nixon Administration to apply science to problems of national need. These plans were never implemented.

The clearest response to this desire of the Nixon Administration was the short-lived Research Applied to National Needs (RANN) program. There were computing-related grants given in only one year—1971—under the RANN program: work on gallium arsenide charge-coupled devices, fault detection in digital sequential circuits, computer applications in a national earthquake information center, and simulations of interconnected power systems.

When John Pasta arrived as OCA head in 1970, he prepared an annual budget that included \$10 million for specialized computing centers and another \$11 million for the regional centers program. The federal Office of Management and Budget (OMB) was opposed to all such facilities grants. Instead, it wanted research grants to include budget lines for facility use because it believed the government could then pay only for actual expenses of specific individual research projects and not otherwise subsidize the academic computing centers.³³ Foundation Director William McElroy sided with the OMB and in 1970 wrote to university and college presidents announcing the end of the Institutional Computing Services grants.³⁴ No new grants were awarded to regional computing centers, and none of the specialized computing centers were ever funded.³⁵ This ended the Foundation's support of large computing machines to colleges and universities.³⁶

6.2 The 1970s and the 1980s

After the Institutional Computing Services grants were terminated, there was significantly less Foundation support provided for computing facilities. Individual computer science research grants sometimes included modest budget lines to access computing facilities. In 1977, the Foundation pooled some of the funds it was expending on small equipment acquisitions in order to fund computers for computer science departments, under a program known as the Research Equipment program.³⁷ This program was affordable because of the arrival on the market of powerful and inexpensive minicomputers, generally VAX computers built by Digital Equipment Corporation. The rationale was that providing one minicomputer

to the entire department saved money compared to funding computing for several individual research projects. Most of the schools that received Coordinated Experimental Research grants—one of the most influential of the Foundation’s computing programs of the 1980s—started with equipment grants from the Research Equipment program.³⁸ This program was continued into the 1990s.³⁹

The computing facilities program was the Foundation’s most important contribution to computer science and computationally driven science in the 1950s and 1960s. However, the Office of Computing Activities, when created in 1967, shifted focus to computer education, largely at the expense of computing facilities grants. A facilities program that could truly support the instructional as well as the research computing needs of U.S. colleges and universities was made unsustainable by the increasing need for instructional computing on American campuses in the 1960s. Efforts by the OMB also created a barrier to a computing facilities program because the Nixon Administration wanted to control what it supported rather than providing a general subsidy to academic computing centers. With the emergence of minicomputers, the Foundation once again could provide computing facilities, but at the department rather than the university level.

6.3 Support for Computer Education, 1950–1986⁴⁰

The legislation establishing the Foundation charged it with improving scientific education as well as supporting scientific research. In the computing field, the Foundation carried out its educational mission in three ways: *computer science education* (fellowships, traineeships, faculty sabbatical programs, teacher training, and curriculum development in both computer science and computer engineering); *computers in science education* (course improvement grants to introduce computing into the science classroom); and *computer-aided instruction* (the use of the computer in all classrooms, not just science classrooms).⁴¹

Throughout most of the 1950s, computer education was addressed only incidentally through other programs, such as fellowships offered through the mathematics program. After Sputnik and the passage of the National Defense Education Act in 1958, the Foundation bolstered all of its programs in science education, including those in computing. A major increase in the Foundation’s computer education activities occurred in the 1960s as a result of the Pierce Report and President Lyndon Johnson’s Great Society education initiatives.⁴² NSF established the Office of Computing Activities in 1967 in part to carry out a substantial program in computer education. This led to a sharp increase first in computer-aided instruction and the

use of computers in science education, and later in improvements to computer science education. In 1972, the education program shifted from OCA to the Education Directorate, after which the emphasis of computer education within the computing program was on computer science curriculum development and “manpower” issues. In the early 1980s, with the coming of the Reagan Administration, federal support for education (including computer education) came under assault and was scaled back across the Foundation.

6.3.1 Educational Support in the 1950s

In the Foundation’s early years, the staff moved cautiously in science education not only because of a small budget and small staff, but also because of a strong public view that education was a local rather than a federal concern.⁴³ It was also unclear whether the computer would ever become a major tool for the delivery of education. Initial efforts were focused on providing graduate fellowships and supporting faculty enrichment through sabbaticals and conferences; these activities did not impose national standards on colleges, universities, and public schools. Moreover, these types of support were closely aligned with the goal of the Foundation to support scientific research. In the 1950s it was widely believed that effective use of the computer as a scientific research instrument required advanced mathematical training, and not surprisingly the Foundation’s initial efforts in computer education were targeted at increasing the number of mathematicians and their knowledge of computers.

Leon W. Cohen, the program director for Mathematical Sciences, gave the earliest known talk on the shortage of computing personnel at a 1954 conference NSF funded at Wayne University in Detroit.⁴⁴ He encouraged sabbatical support for mathematicians to spend a year at a large university computing facility such as UCLA or the University of Illinois. He also advocated graduate fellowships and postdocs to mathematicians and thought university computing centers should offer research seminars on computing topics such as numerical analysis and coding. The Mathematical Sciences division had been providing summer institutes for college mathematics teachers since 1953, and Cohen suggested that they be expanded to include computing topics.

In 1955 the Foundation created an *ad hoc* Advisory Panel on University Computing Facilities, chaired by the eminent mathematician John von Neumann. The committee recommended “a limited program to provide computing equipment and partial support for appropriate staff in order to carry out research and training in high-speed computation.”⁴⁵ The following year, the Foundation noted that “at present only a fraction of the number of mathematicians needed for computer

Table 6.1 NSF fellowships and traineeships in computer science, 1965–1971⁴⁸

1965	33
1966	27
1967	24
1968	54
1969	58
1970	76
1971	70

work are being graduated at the various levels.”⁴⁶ The same annual report noted that “scientists in other fields, also, must be trained in methods of applying computer techniques to their own problems” as a means to develop adequate staffing for industrial and defense needs.

Following the Sputnik crisis of 1957 and the passage of the National Defense Education Act the following year, the Foundation ramped up all aspects of science education. This led directly to an increase in fellowships and traineeships, curriculum development grants, and teacher training institutes in the computing field. Computer science fellowships were awarded by Mathematics in this era, and only beginning in 1965 were the fellowship awards in computer science listed separately from those in mathematics. (See Table 6.1.) Between 1965 and 1974, awards going to computer science increased from 10 to 20 percent of the total mathematics program fellowship awards. In 1974, the Foundation temporarily discontinued most of its fellowships and all of its traineeships, later explaining that the national shortage of “scientific research manpower” was reduced and that there was a new “need for a range of scientific and technical competencies well beyond those possessed by individuals whose academic preparation is primarily for pursuits of careers in basic research.”⁴⁷

In the mid-1960s, two Foundation-wide programs occasionally supported educational computing: the Institutes for Science, Mathematics and Engineering Teachers program and the Course Content Improvement program. The first computing summer schools were held at the University of Oklahoma in 1959 and 1960. Between 1964 and 1968, five or six computing summer schools were held each year—typically targeted at high school, junior college, or four-year college teachers. Curriculum development grants through the Course Content Improvement program followed a similar pattern. The first two curriculum development grants

in computing were awarded in 1959, and throughout the 1960s between two and ten course content improvement awards were made each year in computing. By the time the program was terminated in 1970, the Foundation had supported 73 course improvement grants in computing. Projects ranged from the development of curricula for teaching computer principles, to using computers as an instructional tool in both scientific and non-scientific disciplines. Perhaps the most significant of these awards was made to John Kemeny and Thomas Kurtz to develop the Beginner's All Purpose Symbolic Instruction Code (BASIC) programming language and make it an integral part of the Dartmouth education for all undergraduates.⁴⁹

6.3.2 Computer Education in the Great Society

Education was a major element of President Lyndon Johnson's Great Society program. In 1965 the President's Science Advisory Committee convened a Panel on Computers in Higher Education, chaired by John Pierce of Bell Laboratories.⁵⁰ The panel reported that 35% of undergraduates could benefit from access to adequate computing facilities, but fewer than 5% had this access—and then only at a few “favored schools.” The committee recommended that the federal government bear much of the estimated \$400 million price tag to provide adequate computing facilities to universities, and that federal agencies provide short courses to train faculty to teach computer science and also provide support for research and education in computer science. The committee also proposed that the Foundation and the Office of Education establish a study group on computers in high school. President Johnson's message to Congress on health and education in 1967 gave force to these recommendations.⁵¹ Mainly in response to this presidential message, the Foundation created the Office of Computing Activities (OCA) in July 1967.⁵² In his 1967 message to Congress on health and education, President Johnson called for NSF to work with the Department of Education to develop the potential of computers in education.

When Milton Rose first convened his OCA advisory committee in August 1967, there was no simple consensus about what the office should be doing.⁵³ So OCA decided to support varied educational experiments.⁵⁴ Numerous efforts across the country brought computers into higher education in the United States. By 1968, 90% of all U.S. colleges and universities enrolling 2,500 or more students offered some instructional use of computers, mostly programming courses.⁵⁵ This made it difficult for OCA to settle on an appropriate mission. OCA's difficulties were compounded by the Vietnam War and the drastic cuts in domestic spending made

to fund the war effort. For example, OCA received less than 20% of its proposed budget of \$72 million for the 1970 fiscal year.⁵⁶

Beginning in 1968, OCA made awards supporting computer science curriculum development, application of computer techniques in teaching science, and the development of computer-assisted instruction (CAI) projects. The CAI projects were the most numerous—about 25 each year through the 1970s—awarded at first by OCA and, beginning in 1974, by the Science Education Division. In the curriculum development area, OCA supported both summer teaching institutes and course development grants. About half of the summer programs were targeted at secondary and vocational teachers. Funding was in short supply, so that the number of annual awards fell off rapidly from a dozen in 1968 to between zero and five per year in the 1970s. Awards in this area received a new boost in 1978, when a new program was created in the Science Education Directorate—making more than 50 awards in each of its first two years—driven by the opportunity to use the newly created microcomputers in undergraduate instruction. One of the projects funded was John Hamblen’s well-known survey of computing activities in higher education.

The computer-aided instruction area was the most controversial for OCA.⁵⁷ The Johnson and Nixon Administrations were supportive of CAI as a means to reduce the high costs of labor in education; and RCA, Westinghouse, and other companies saw CAI as a promising business opportunity. But OCA decided its funds were better spent on academic-based experimental CAI systems rather than on development and implementation projects. Critics of CAI argued that it was shallow and did not result in real learning. But the Foundation persisted in its support.

Between 1958 and 1980, OCA provided 30 grants to support the research of Patrick Suppes in the Institute for Mathematical Studies in the Social Sciences at Stanford University, widely regarded as the leading research center in CAI.⁵⁸ Suppes had built CAI courses to teach subjects ranging from elementary logic to Mandarin Chinese. His group worked with six-year-olds in his Stanford laboratory, underprivileged children in Kentucky and Mississippi, and members of a pueblo in New Mexico. Under this program, OCA also funded Seymour Papert’s Logo project at MIT. OCA bowed to political pressure—deviating from its decision to support only research—when it supported two large CAI system development projects in the 1970s: the Program Logic for Automatic Teaching Operations (PLATO) project at the University of Illinois (a mainframe computer connected to remote terminals to deliver educational material, later commercialized by the Control Data Corporation) and Time-shared Interactive Computer Controlled Information

Television (TICCIT) at the MITRE Corporation (a minicomputer connected to television technology, with course development support from Brigham Young University).⁵⁹ While some of the OCA projects were well executed, the overall impact of the program was limited:

What we couldn't really break through . . . was the economic barrier . . . how to get this stuff out in sufficient quantities to really make an impact on an institution as large as the U.S. school system We underestimated tremendously the resistance of the educational establishment to change and the amount of institutional change that that technology would be forcing in order to be really useful.⁶⁰

At the same time, the Foundation was supporting the use of computers for science education. In 1967, the OCA staff organized a conference on computer use in the teaching of statistics, physics, chemistry, and mathematics at the Science Teaching Center of the University of Maryland.⁶¹ From 1970 through 1978, the Foundation provided support to an annual Conference on Computers in Undergraduate Education. At these meetings, people reported on their actual classroom use of computers across all academic disciplines. These annual conferences were reorganized in 1979 as the new National Educational Computing Conference. Arthur Melmed, Andrew Molnar, and Frederick Weingarten from NSF were actively involved as organizers and participants in these conferences, which disseminated practical knowledge about the uses of computers in education.

6.3.3 A Division of Labor

In 1974 the Foundation moved the education programs out of OCA and into the Education Directorate, and OCA was reorganized as the Division of Computer Research (DCR) with the physicist John Pasta as head. Thereafter, DCR's emphasis was on curricular development and "manpower" issues, although it continued to support a few computer-based education projects.⁶² These few grants awarded in the middle 1970s by DCR supported computing in mathematics education, instrumentation for education in data capture and analysis, and work on interactive computing in laboratory instruction. There was also an educational component associated with the 25 regional computing centers that the Foundation funded.⁶³ Federal budgets for all aspects of computers in education were much less generous after 1972 than they had been in the golden years of 1967 to 1972. Indeed, from 1972 until 1989, support for education was generally lean at the Foundation, increasing only after the end of the Reagan Administration.

The Education Directorate worked closely with the computer division on regional computer centers.⁶⁴ During the 1970s, the Education Directorate supported a consortium called CONDUIT, led by Gerard Weeg of the University of Iowa and Thomas Kurtz of Dartmouth, to make the regional centers more effective. The five regional centers participating in CONDUIT were centered at Oregon State, North Carolina, Dartmouth, Iowa, and Texas at Austin. The networks cooperated in the exchange of materials, translated these materials into BASIC and FORTRAN, prepared documentation, tested the materials in the classroom, and sponsored workshops to promote their use. CONDUIT was regarded as a success.⁶⁵

In the mid-1970s, the Education Directorate supported a project in which MITRE Corporation wired homes in Reston, Virginia, with two-way interactive televisions and computers; this system delivered the drill-and-practice materials in mathematics that Patrick Suppes had developed at Stanford. Later, a similar project was funded in Buffalo, New York, using cable services to deliver instruction to handicapped children.

The availability of inexpensive microcomputers stimulated new activity in the Education Directorate. A new Undergraduate Instructional Development program in 1978 supported microcomputer applications to teach science, awarding more than 50 grants that year and a similar number the following year. For example, the University of Utah enhanced the use of computer graphics in engineering education, and the Wicat Corporation developed an intelligent videodisc system for teaching developmental biology. That same year, the Education Directorate sponsored a major conference on the application of computer technology to science education.⁶⁶ The next year, a second conference was also convened by the Education Directorate. The two major European manufacturers of videodisc systems, Philips and Thompson, were invited and subsequently designed an interactive learning system incorporating a small computer and optical discs. Other Foundation-funded projects resulting directly from this conference included a project at Brigham Young University to experiment with an interactive video system to teach about DNA, and a project at the University of Utah to test a different interactive system in the teaching of physics and engineering.⁶⁷

The Foundation had played an important role in the early study of computer applications to education. By the mid-1980s this had become an identifiable professional subfield of study, with more than 20 journals devoted to it.⁶⁸ However, in 1983, at the direction of the Reagan Administration, the Foundation cancelled all of its programs in science education. The lone program officer managing the completion of already awarded grants was Andrew Molnar, who oversaw a portfolio of 500 grants, covering all levels of education. He, together with another program

officer, Dorothy Derringer, managed to carry out a small computer education program in their spare time, funded by industry.⁶⁹

There was one last effort in computer-based education in the 1980s. Despite the political push to reduce the role of the federal government in education during the Reagan years, there were five major studies originating from the science policy community on computer-based education published during the first half of the 1980s.⁷⁰ As a result, just as the new computer science directorate (CISE) was being formed, a new program was created in the Application of Advanced Technologies, supporting both research and development in computer-based education.⁷¹ Projects funded by this program included learning systems for basic algebra, problem-solving in geometry, fundamental mathematical concepts for grade-school students, and algorithm discovery for undergraduate computer science students. Perhaps the best-known project supported by this program was again work by Patrick Suppes—a system developed at Stanford to teach college-level calculus to seventh and eighth grade students.

6.3.4 Building the National Computer Science Community

Before turning to an examination of the Foundation's role in developing a curriculum for computing, this section presents the national context for the development of computer science and the role of the universities in training computer scientists. As we will see, the Foundation's role was primarily to support the efforts of professional societies and individual universities in these curricular efforts, as well as to play a coordinating role for the emerging computer science community. Those topics are discussed in detail in the next section.

Some of the earliest computers were built on university campuses, and it was those universities that taught the first computer courses: Harvard, MIT, and the University of Pennsylvania were already teaching computing courses in the late 1940s, and by the early 1950s they were joined by UCLA and Berkeley. Computing courses were originally taught only at the graduate level because the equipment was too expensive and too scarce for masses of undergraduates. MIT, which had begun graduate instruction in computing in 1947, offered its first undergraduate computing course in 1953—probably the first undergraduate computing course in the United States.

Early computing courses included mathematical topics, especially numerical analysis, and electronic engineering topics such as switching circuits. As more universities began to offer computing instruction in the 1950s, the instructional programs tended to fall into one of two categories. Some universities established computing centers, and these centers offered practical instruction in how to use

computing equipment. Examples included the Wayne University Computation Laboratory, the University of Michigan’s Willow Run Research Center, Georgia Tech’s Rich Electronic Computing Center, the University of Illinois Computation Laboratory, and Purdue’s Computation Laboratory. At other universities, however, either electrical engineering or mathematics departments offered introductory courses in logical design, programming, or applications—often even before the university had acquired its own computing equipment. By the mid-1950s, other academic units, such as business schools and agronomy departments, were beginning to offer computing courses. A 1954 survey by the Institute of Radio Engineers of 68 universities conducting some activity in digital or analog computing found 29 were offering at least one computing course—and of these, 9 were offering three or more courses.⁷²

Already by 1954 there were discussions about the directions for computing instruction. Howard Aiken, the director of Harvard’s Computation Laboratory and an early leader in computer education, argued for a broad education, of both “sociology and computing devices.”⁷³ F. Joachim Weyl, Director of the Office of Naval Research (ONR) Mathematical Sciences Division, argued that broad computer instruction was urgent:

an unprecedented mathematization; not only of fundamental scientific research in the physical and biological sciences but also in the management of our industrial and social systems. This is about to assign to mathematics an entirely new part in our civilization with far-reaching implications on what should be taught, how it should be taught and to whom.⁷⁴

Throughout the 1950s and the early 1960s, the Foundation made no concerted effort to develop computer science and computer engineering curricula. As the 1960s went on, however, the Foundation took an increasingly active role—though primarily a supporting role—in developing a computer science curriculum. The late 1950s and the decade of the 1960s was an important time for computing in the United States. As price-performance characteristics of computing systems became increasingly more favorable, the demand for both computers and computer professionals grew. Labor statistics for computing in the 1960s varied widely (Bureau of Labor Statistics, American Federation of Information Processing Societies (AFIPS), the Pierce Report, and various industrial projections)—because of different definitions regarding whom to include as a computer professional and from taking statistical portraits at slightly different times; but they all showed steady and significant annual increases in demand for computing personnel in the 1960s and first half of the 1970s. For example, AFIPS claimed there were 10,000 systems analysts and 40,000 programmers working in the United States in 1960, and 60,000 systems

analysts and 60,000 programmers by 1965. The need in each of these professions was projected to exceed 200,000 workers by 1970.⁷⁵

Computer labor was changing. By the 1960s, military and industrial computer projects of great size and complexity, such as the SAGE air defense system and the operating system for the 360 family of IBM computers, demonstrated an acute need for large numbers of skilled programmers.⁷⁶ There was increasing concern about writing reliable software on time and on budget, and a famous 1968 UNESCO conference coined the term “software crisis.”⁷⁷ The primary expense of a major computing project was steadily shifting away from hardware costs and increasingly to personnel costs.⁷⁸

In response to this rapid advancement in computing and the rise in demand for computing workers, universities around the United States began to establish computer science departments—first at the graduate level, then at the undergraduate level. Purdue established the first computer science department in the United States in 1962. Others soon followed. Of the computer science departments existing in the United States in 1988, over 60% were founded between 1962 and 1972.⁷⁹ By the end of the 1960s, about half of the students studying computer science and computer engineering were at the bachelor’s level, half at the master’s level—with very few Ph.D.s.⁸⁰

There were early efforts at national standardization. The two largest computing societies active in the United States, the ACM and the AIEE (later the IEEE) Computer Society, both of which had been formed in the late 1940s, were active in these curriculum standardization efforts in the 1960s and 1970s. The ACM formed a permanent curriculum committee in 1964 and finalized a computer science curriculum in 1968. The Foundation supported the work of this committee throughout the 1960s and 1970s.⁸¹ Between 1967 and 1972, the COSINE committee (Committee on Computer Science in Electrical Engineering Education) of the National Academy of Engineering developed a curriculum for computer engineering. In 1977, both the ACM and the IEEE Computer Society (founded 1970) proposed curricular revisions—with some overlap but reflecting the difference in viewpoint of the engineers in IEEE and the scientists in ACM. The Foundation was represented by Bruce Barnes on both committees.⁸² These curricula developed by the two professional societies and the National Academy were widely adopted by American colleges and universities.⁸³

6.3.5 The Foundation’s Role in Meeting National Needs for Computer Scientists

The Office of Computing Activity helped individual computer science departments and programs to establish effective curricula and, more generally, to build up insti-

tutional size and strength. The need for computer scientists first became apparent in the mid-1960s with the increasing demand for computer professionals across all sectors of American society.

Alan Perlis, at that time chair of the Carnegie Mellon computer science program and a member of the OCA advisory committee, identified 11 strong graduate programs in computer science (Berkeley, Carnegie Mellon, Harvard, Illinois, MIT, Michigan, NYU, Penn, Purdue, Stanford, and Wisconsin) as of 1967, and he projected a total of 81 programs of varying quality. As of 1967, approximately 200 computer science faculty graduated 40 new Ph.D.s; but Perlis forecasted the need for 400 faculty to meet teaching needs. In academic year 1964–1965, 4,300 undergraduates and 1,300 graduate students were enrolled nationally in computer science degree programs (including data processing programs), and the numbers were rising rapidly—the numbers quadrupled only two years later.⁸⁴

OCA might have increased the number of fellowships and traineeships to help create the next generation of computer science professors. However, a glut of scientists in other scientific disciplines led to NSF-wide changes harmful to computer science. In 1969, the advisory committee for Mathematics and Physical Science noted a significant surplus of American scientists compared to research funding available; some graduate-trained scientists were leaving science entirely.⁸⁵ The Nixon Administration's Office of Management and Budget pressured the Foundation to terminate its student traineeship program in 1971. This change exacerbated the already difficult situation for computer science. OCA responded to the 1969 MPS advisory report with an appeal for special dispensation for computer scientists, but no new fellowships or traineeships in computer science resulted. One effect was that leading computer science departments, such as Cornell, Purdue, and Stanford, were only able to fund (and thus admit) approximately 25% of their applicants, whereas the overall admission rates in graduate programs at these universities in other disciplines was approximately 60%.

In 1967, OCA began a trial program to support computer science programs. In its first year, OCA provided funds to help Johns Hopkins, Ohio State, and NYU strengthen their graduate programs, and for Colgate to create an undergraduate program. The following year, OCA provided \$1 million to strengthen graduate programs at Berkeley, Purdue, Rhode Island, University of Southern California, SUNY Stony Brook, and Washington University in St. Louis. After these two trial years, the OCA advisory committee recommended continuing and increasing funding to build up national educational capacity in computer science. Unfortunately, during the Nixon Administration years, funding was well short of what OCA needed.

As of 1975, 62 computer science departments and 53 other types of departments (mathematics, engineering, information science, statistics, etc.) were awarding doctorates with a computing emphasis. Just over 2,000 computer scientists were working in these programs—1,500 of them in computer science departments. While undergraduate computing enrollments continued to rise throughout the 1970s, computer science doctoral production peaked for the decade in 1976—at 244.⁸⁶ The flattening of Ph.D. production probably reflected insufficient faculty advisors and insufficient funding. Another reason may have been the perception of life as a faculty member: In a survey of people who left the university for industry, the Foundation found that people made this change primarily because of heavy teaching loads and job insecurity.⁸⁷ All engineering fields were subject to this academic flight, but it occurred in computer science and computer engineering at twice the rate as for the rest of engineering. There were also concerns because tenure committees did not understand the nature of computer science research, which often had high cost and time requirements and yielded relatively few publications.⁸⁸

One long-term solution to the computing “manpower” shortage was the wide-ranging Coordinated Experimental Research program implemented in the early 1980s (see chapter 2). Nobody believed that it was providing short-term relief to this problem, but that instead it might have an impact over the long term.

6.4 Conclusions

NSF carried out efforts in computer science education (fellowships, traineeships, faculty sabbatical programs, teacher training, and curriculum development in both computer science and computer engineering), computers in science education (course improvement grants to introduce computing into the science classroom), and computer-aided instruction (the use of the computer in all classrooms, not just science classrooms), especially after the passage of the National Defense Education Act in 1958, until 1981, when the Reagan Administration took office. Some of this activity was undertaken by OCA, then later by the Education Directorate. The growing needs for computing professionals during this era were not always in step with the perceived personnel needs in the rest of science and technology, and especially in the early 1970s, this hampered NSF’s abilities to support education for computer scientists.

Modern computing is largely contemporaneous with NSF. Scientists were among the earlier users of high-speed computers, and the Foundation began to support computational science within a year of its formation. During the 1950s, the Foundation’s major contribution to the computing field was its facilities pro-

gram, which provided computers to both large and medium-sized universities for research across a wide range of computational sciences. As it became clear how valuable computers were to both scientific research and education, before the end of the 1950s the demand had outstripped the Foundation's ability to provide individual machines to colleges and universities. NSF tried to provide a partial solution to this demand through its support of regional computer centers that linked research universities to nearby colleges and high schools, through minicomputers for individual computer science departments and computer science laboratories, and eventually to the formation of national computing centers equipped with high-performance computing facilities.

The 1960s was a time when the academy recognized how valuable the computer could be as an instructional tool, not only for teaching computer science—or science and engineering more generally—but also for providing education across a wide range of disciplines. NSF was involved from early times in this effort to demonstrate the educational value of computers. It was, however, beyond the Foundation's budget to provide computers for educational use in colleges, much less in public schools; so the Foundation's main contribution in this area was to help demonstrate the potential of computers to education. One example in particular, the PLATO project at the University of Illinois, showed this potential. The Foundation provided support for experimentation in computerized education in the late 1960s and early 1970s (although the program was cut off abruptly in the early 1980s by the Reagan Administration's belief that education was a state and local matter, not a federal responsibility or prerogative). During the same period, the Foundation provided grants to individual universities and professional societies to develop computer science and computer engineering curricula and implement new college-level courses. From the time of the passage of the NDEA in 1958, the Foundation provided fellowships and traineeships in computer science, which helped to train computer science professors and other people who worked in a spectrum of jobs that involved high-skill computing and computer science.

Notes

1. This section is similar to W. Aspray and B. O. Williams. 1994. Arming American scientists: NSF and the provision of scientific computing facilities for universities, 1950–1973. *IEEE Annals of the History of Computing*, 16(4), 60–74. DOI: [10.1109/85.329758](https://doi.org/10.1109/85.329758). Both this section and that paper were based on an unpublished manuscript: W. Aspray, B. O. Williams, and A. Goldstein. 1992. *Computing as Servant and Science: Impact of the National Science Foundation*, which was an independent study of computing programs of the National Science Foundation through 1980. This section differs from the other paper by staying more

focused on the facilities programs and less on the organizational history of computing at NSF; and by extending the story into the 1980s. This chapter also benefits from additional archival materials and additional historical research on related topics.

2. For a full list of the institutions that received grants, please refer to the Charles Babbage Institute.
3. F. J. Weyl. 1955. Summary of conference discussions and proposals: Panel discussion; Opening remarks. In A. W. Jacobson, ed. *Proceedings of the First Conference on Training Personnel for the Computing Machine Field*. Detroit: Wayne University Press, pp. 84–85. (This book contains the proceedings of the first conference dedicated to the subject of computer education; the conference was held in Detroit at Wayne University in 1954.) In fact, as early as 1953 the Foundation was receiving grant proposals with computing requirements. The Foundation asked the Applied Mathematics Laboratory of the National Bureau of Standards for help in evaluating these proposals, in keeping with an arrangement that had assigned the Bureau the responsibility of advising federal agencies on computing matters.
4. Punched card tabulating equipment had had a similar, but less extensive impact on scientific research when it was first introduced in the 1930s. See G. W. Baehne, ed. 1935. *Practical Applications of the Punched Card Method in Colleges and Universities*. New York: Columbia University Press.
5. The four commercially available computer systems were the Remington Rand UNIVAC 1, the IBM 701, the Engineering Research Associates 1103, and the Burroughs DATATRON—according to M. H. Weik. June 1957. *A Second Survey of Domestic Electronic Digital Computing Systems*. Aberdeen Proving Ground, Maryland: Ordnance Research and Development, Department of the Army. <https://catalog.hathitrust.org/Record/000427964>.
6. National Science Foundation. 1955. *Fifth Annual Report*, pp. 53–55. https://www.nsf.gov/pubs/1955/annualreports/ar_1955.pdf.
7. National Science Foundation. 1965. *Sixth Annual Report*, pp. 57–58. https://www.nsf.gov/pubs/1956/annualreports/ar_1956.pdf.
8. L. Fein. September 1959. The role of the university in computers, data processing, and related fields. *Communications of the ACM*, 2: 7–14. DOI: [10.1145/368424.368427](https://doi.org/10.1145/368424.368427).
9. Fein, 1959, *op. cit.*, pp. 9, 13, 14.
10. V. H. Swoyer. 1980. Computer system changes. In J. W. Hamblen and C. P. Landis, eds. *The Fourth Inventory of Computers in Higher Education: An Interpretative Report*. Boulder, CO: Westview Press, pp. 43–60.
11. Oral history, Arthur Grad (program director for Mathematical Sciences), interviewed by William Aspray, October 19, 1990. Also see National Science Foundation, 1958, *Annual Report*, https://www.nsf.gov/pubs/1958/annualreports/ar_1958.pdf; and *Annual Report of the Division of Mathematical, Physical, and Engineering Sciences*, 1959, NSF Historian's Files. For a lengthy case study of the computing facilities program from the perspective of one university recipient (Duke University), see Aspray and Williams, 1994, *op. cit.*
12. National Science Foundation. 1959. *Annual Report of the Division of Mathematical, Physical, and Engineering Sciences*. NSF Historian's Files, p. 7.

13. T. E. Kurtz. 1974. The NERComp network. In M. Greenberger et al., eds. *Networks for Research and Education*. Cambridge, MA: MIT Press, p. 282.
14. National Science Foundation. 1960. *Annual Report of the Division of Mathematical, Physical, and Engineering Sciences*. NSF Historian's Files, pp. 6–7.
15. The Foundation staff evaluating equipment proposals would often make site visits to the applicant campus—something that was not then common at the Foundation. The site visit teams included not only computer scientists but also biologists, chemists, and physicists. (See Oral history, Milton Rose, interviewed by William Aspray, November 6, 1990. Charles Babbage Institute; and Oral history, Frederick Weingarten, interviewed by William Aspray, September 26, 1990. Charles Babbage Institute. But also see Oral history, Frederick Weingarten, interviewed by Peter Freeman, July 11, 2017.) The program officer responsible for the computing facilities program was John Aufenkamp. Arthur Melmed, Glenn Ingram, and Frederick Weingarten also worked on the facilities program.
16. In 1965 the Foundation convened a Working Group on Computer Needs in Universities and Colleges, with representatives from the Mathematical Sciences Section, the Engineering Division, the Education Division, the Planning Office, and the Comptroller's Office. The group expressed concern about the reliability of the projections concerning need for computers because of the thin evidence on which they were based.
17. National Science Foundation. 1961. *Annual Report of the Division of Mathematical, Physical, and Engineering Sciences*. NSF Historian's Files, p. 3. Occasionally, there were computing facility grants from other parts of the Foundation. For example, in 1961 the Scientific Personnel and Education Division program in Undergraduate Instructional Scientific Equipment made grants of small computers to 20 colleges. The computing staff in MPES were highly critical of these awards because of the loose standards that had been applied in making them. As a result, MPES demanded greater information about a university's computing needs and existing equipment before making an award, and also increased the research requirements on grants for small computers (National Science Foundation. 1962. *Annual Report of the Computers and Computing Science Program*. NSF Historian's Files).
18. U.S. Congress, House of Representatives. June–August 1965. “Government and Science; Review of the National Science Foundation,” Hearings before the Subcommittee on Science, Research and Development of the Committee on Science and Astronautics, 89th Congress, 1st Session (Washington, DC: Government Printing Office), 1: 651, 660, 667–9, 747. Also see 787, where NSF Director Haworth notes the “big money” it would cost to meet this need for academic computing facilities. <https://catalog.hathitrust.org/Record/008467056>.
19. National Science Foundation, Office of the Director. April 11–12, 1968. Subject files—1968, Advisory Committee for Computing Activities, Background Materials and Agenda of the Third Meeting, Records Accession No. 307-74-038, Box 1, Washington Federal Records Center. At the same time as these computer facility grants were being made, the Foundation also supported research (awarding grants in 1957 and 1962) by Philip Morse at MIT on time-sharing, which was regarded as more promising for instructional computing than batch processing. Indeed, time-sharing proved to be a better solution to instructional computing than batch processing, but the cost of time-sharing systems delayed their spread into universities.

20. Answer to question No. 4, attached to Statement of Dr. Milton E. Rose, Head, Office of Computing Activities, before the Subcommittee on Science, Research, and Development of the Committee on Science and Astronautics, U.S. House of Representatives, March 1969. Copy in National Science Foundation, Office of the Director, Subject Files, 1969, Records Accession No. 307-75-052, Box 2, Washington Federal Records Center.
21. Weingarten interview, 1990, *op. cit.*
22. All the quotations in this paragraph are from Rose interview, 1990, *op. cit.*
23. Early draft of briefing notes attached to Statement of Dr. Milton E. Rose, 1969, *op. cit.*
24. The program officer in charge of the regional computing centers was Larry Oliver.
25. Draft briefing notes, Statement of Dr. Milton E. Rose, 1969, *op. cit.*
26. C. Mosmann. 1973. *Academic Computers in Service*. San Francisco: Jossey Bass, pp. 43–44; see also L. H. Williams. 1974. A functioning computer network for higher education in North Carolina. In Greenberger et al., *op. cit.*, pp. 222–232.
27. National Science Foundation. 1972. *22nd Annual Report*, pp. 50–51.
28. Not all of the regional networks were successful. A good example is a plan to create a computing center in Madison, Wisconsin, to serve all campuses of the University of Wisconsin, which was abandoned because of technical and management problems arising from inter-campus rivalries.
29. Mosmann, 1973, *op. cit.*, pp. 43–44; Williams in Greenberger et al., 1974, *op. cit.*, pp. 222–232.
30. T. E. Kurtz, The NERComP Network. In Greenberger et al., 1974, *op. cit.*, pp. 282–287.
31. G. P. Weeg. 1974. Regional star networks as seen by the user and server, 320–337, and N. R. Nielsen, Network computing, pp. 64–73 (both in Greenberger et al., 1974, *op. cit.*).
32. National Science Foundation. 1969. Advisory Committee for Computing Activities. January 16, 1970. “Annual Report 1969” to W. P. McElroy from D. Alpert, chairman, Office of the Director. Subject Files—1970, Records Accession No. 307-75-053, Box 2, Washington Federal Records Center.
33. J. R. Pasta. December 15, 1970. “Introduction and Background,” Directors Program Review: Computing Activities, NSF, 5–6. To the extent that this auditing practice was established in university computing centers, it caused computing centers to retreat from the open access policy mandated by the Foundation and restrict the number of free hours of computing time given.
34. Notice No. 30. July 30, 1970. “Important Notice to Presidents of Universities and Colleges and Directors of Non Profit Institutions,” Office of the Director, Subject Files—1970, Records Accession No. 307-75-053, Box 2, Washington Federal Records Center. An earlier draft of the announcement retained the program for specialized computing centers, but it was excised in the final announcement. See the draft of May 13, 1970, attached to the Agenda for the Seventh Meeting of Advisory Committee for Computing Activities, June 11–12, 1970, in the same record collection.
35. National Science Foundation. June 10, 1970. Memorandum to David E. Ryer, Special Assistant to the Director from John R. Pasta, Head, Office of Computing Activities, Office of the Director, Subject Files—1970, Records Accession No. 307-75-053, Box 2, Washington Federal Records Center. In Pasta’s elope, Kent Curtis and colleagues said that Pasta did not

- agree with the OMB position but accepted the cuts in facilities as inevitable (K. K. Curtis, N. C. Metropolis, W. G. Rosen, Y. Shimamoto, and J. N. Snyder. July 1983. John R. Pasta, 1918–1984. *Annals of the History of Computing* 5: 224–238).
36. There had been hope among the OCA staff for a continuation of the computing facilities program when Guy Stever arrived in 1972 as the new Foundation Director because the Foundation had made a major computer grant to Carnegie Mellon University while he was president there. However, one of his first decisions at the Foundation was to end the facilities program. In fact, Stever was also opposed to the Foundation's involvement in educational activities (Oral history, Andrew Molnar, interviewed by William Aspray, September 25, 1991. Charles Babbage Institute). Weingarten remembers: "Lee Haworth was . . . very interested in forming OCA and would support it. But afterwards, . . . Guy Stever never believed that computer science even existed. He used to resent it. He would say, 'Well, why don't you have car science?' . . . every time you would bring something up to him it was a constant battle . . ." (Weingarten interview, 1990, *op. cit.*). This is interesting inasmuch as Stever was CMU president at a time in the late 1960s when the computing program was building up to become a national powerhouse.
 37. The Research Equipment program bought computers only for computer science departments. Other Foundation divisions were buying computers for departments and laboratories in other sciences. There was no apparent coordination of these computer facilities programs within the Foundation.
 38. See Chapter 2 for a more detailed discussion of the Coordinated Experimental Research program.
 39. The material on the Research Equipment program is taken from three oral histories: Bruce Barnes, interviewed by William Aspray, September 26, 1990; Thomas Keenan, interviewed by William Aspray, September 28, 1990; and Harry Hedges, interviewed by Frederick Nebeker, September 26, 1990. All of these are held by the Charles Babbage Institute.
 40. This section is similar to the hard-to-access paper: W. Aspray and B. O. Williams. 1993. Computing in science and engineering education: The programs of the National Science Foundation. *Electro 93 International*, Conference Record, April 27–29, 1993. Edison, NJ. 2: 234–239. Both were based on Aspray, Williams, and Goldstein, 1992. This section benefits from additional archival materials and additional historical research on related topics.
 41. A longer study would place this account into the larger history of engineering education in America. One might consider the works listed here and others by established technology historians such as A. Aker, 2006. *Calculating a Natural World*. MIT Press; A. S. Bix. 2014. *Girls Coming to Tech!* MIT Press; E. Layton. 1986. *The Revolt of the Engineers*. Johns Hopkins University Press; R. Kline. 2000. The paradox of "Engineering Science"—a Cold War debate about education in the U.S. *IEEE Technology and Society Magazine* 19(3): 19–25. DOI: [10.1109/44.868938](https://doi.org/10.1109/44.868938); and T. Reynolds and B. Seely. 1993. Striving for balance: A hundred years of the American Society for Engineering Education. *Journal of Engineering Education*, 82(3): 136–151. DOI: [10.1002/j.2168-9830.1993.tb00092.x](https://doi.org/10.1002/j.2168-9830.1993.tb00092.x).
 42. There are histories of some of the early computer science departments in the United States, e.g., Georgia Tech, Purdue, and Stanford. Each of them in one way or another had a connection to NSF in the creation of its department. For example, Georgia Tech received NSF grants in 1961 and 1962 to hold international conferences on information

science, and in 1964 Georgia Tech received an NSF grant to help it form a department of information and computer science. See, for example, History of GT Computing, <https://www.cc.gatech.edu/node/2650>, last accessed 21 January 2019, which includes links to various articles about computing at Georgia Tech, but see in particular: P. Freeman. July 27, 2015, Origins of the College of Computing, 1945–1990, https://www.cc.gatech.edu/sites/default/files/documents/coc_origins_to_post_0.pdf; R. L. Pyle. 2015. *First in the Field: Breaking Ground in Computer Science at Purdue University*. Purdue University Press; and Department Timeline, Computer Science, Stanford Engineering, <https://cs.stanford.edu/about/department-timeline> (last accessed 21 January 2019).

43. For example, the state-wide debate topic for AY1955–56 in Texas was whether the federal government should provide support for public education (private communication, Peter Freeman, August 2018). See also W. Aspray. July 2000. Was early entry a competitive advantage? US universities that entered computing in the 1940s. *IEEE Annals of the History of Computing*. DOI: [10.1109/85.859525](https://doi.org/10.1109/85.859525).
44. Leon W. Cohen. 1955. Cooperation between the National Science Foundation and educational institutions for mathematical research and education. In A. W. Jacobson, ed. *Proceedings of the First Conference on Training Personnel for the Computing Machine Field*. Detroit: Wayne University Press. Computer Science, Department Timeline, Stanford Engineering, <https://cs.stanford.edu/about/department-timeline>; last accessed 18 January 2019.
45. NSF, *Annual Report*, 1955, pp. 54–55.
46. NSF, *Annual Report*, 1956, p. 58.
47. NSF, *Annual Report*, 1977, p. 87.
48. Table notes: These are primarily graduate fellowships, but they also include a few cooperative graduate fellowships, graduate traineeships, graduate teaching assistantships, postdoctoral fellowships, senior postdoctoral fellowships, senior faculty fellowships, secondary school teacher fellowships, and senior foreign scientific fellowships. The numbers for the years 1967–1971 are probably low because NSF did not break out the field of graduate traineeships in these years. (The total number of traineeships awarded by the Foundation in all science and engineering areas was 2,784 in 1965, 5,656 in 1968, and 3,458 in 1971; but the percentage awarded in computer science was probably very low and unlikely to swamp the numbers reported in this table.)
49. A. R. Molnar. June 21, 1989. “Computers in Education: A Historical Perspective of the Unfinished Task,” unpublished typescript based on comments at the National Educational Computing Conference, Boston; A. W. Luehran and J. M. Nevison. May 31, 1974. “Computer Use Under a Free-Access Policy.” *Science* 184: 957–961. Examples of some of the other curriculum development grants of this time are ones awarded to Case Institute of Technology on automata and nets, the University of Arizona on molecular biology with computer-controlled displays, SUNY Buffalo on computer systems, and Cornell for building a computer system for introductory instruction.
50. President’s Science Advisory Committee, Panel on Computers in Higher Education, John R. Pierce, chairman, “Computers in Higher Education,” Washington, DC: Government Printing Office, 1967.

51. "Section C. Computers, Administrative History of NSF during the Lyndon Baines Johnson administration," p. 116. Draft copy in the NSF Historian's files.
52. For more on the collaboration between the Foundation and the Office of Education, the involvement of the White House and Defense Department, and the role of state governments in computer education in response to President Johnson's call, see chapter 8 of Aspray, Williams, and Goldstein, 1992, *op. cit.*
53. OCA Advisory Committee. August 9–10, 1967. "Minutes of the First Meeting." Records Accession No. 307-75-051, Box 1, Washington Federal Records Center. Arthur Melmed ran the education program in OCA.
54. National Science Foundation. Office of the Director, Subject Files—1968, "Background Materials and Agenda of the Third Meeting, 11–12 April, 1968," Records Accession No. 307-74-038, Box 1, Washington Federal Records Center.
55. U.S. Congress, House Subcommittee on Science and Astronautics. March 1969. Statement by Dr. Milton E. Rose attached to agenda, "Fifth Meeting of Advisory Committee for Computing Activities, May 22, 1969."
56. National Science Foundation. 1970. *Twentieth Annual Report*, p. 80. https://www.nsf.gov/pubs/1970/annualreports/ar_1970.pdf.
57. W. McElroy. December 1, 1969. To F. E. Westheimer, President's Science Advisory Committee, Director's Note Files in NSF Historian's Files. On the history of CAI, see R. P. Niemiec and H. J. Walberg. 1989. From teaching machines to microcomputers. *Journal of Research on Computing in Education*, 21(3): 263–276, <https://eric.ed.gov/?id=EJ404188>. The OCA staff did not like the term "computer-aided instruction" but it was used in the popular press. The staff preferred the name "computer-based education" or later "computer-managed education" (Molnar interview, 1991, *op. cit.*).
58. Dick Atkinson, later the director of NSF, worked in Suppes's lab and had received a grant for computer education work in the late 1960s from the Foundation's Office of Education (Molnar interview, 1991, *op. cit.*).
59. For an in-depth study of OCA's activities in CAI, and its supporters and critics, see chapter 8 of Aspray, Williams, and Goldstein, 1992, *op. cit.*
60. Weingarten interview, 1990, *op. cit.* Also see Weingarten interview, 2017, *op. cit.*
61. "Computers in Undergraduate Education: Mathematics, Physics, Statistics and Chemistry," Proceedings of a conference sponsored by the National Science Foundation and conducted at the Science Teaching Center of the University of Maryland, College Park, December 8–9, 1967, distributed by J. David Lockard, Director, Science Teaching Center, University of Maryland.
62. Molnar interview, 1991, *op. cit.*
63. On the educational activities in OCA in the last few years before many computer education programs were moved to the Education Directorate, see the following NSF *Annual Reports*: Nineteenth (1969), pp. 92–93; Twentieth (1970), pp. 82–83; Twenty-first (1971), 50–51; Twenty-second (1972), pp. 50–51; and Twenty-third (1973), p. 89.
64. Molnar interview, 1991, *op. cit.*

65. See, for example: L. T. Parker, Jr. and J. P. Denk. Spring 1974. A network model for delivering computer power and curriculum enhancement for higher education: The North Carolina Educational Computing Service. *EDUCOM Bulletin*, 9: 24–30.
66. “Technology in Science Education—The Next Ten Years—Perspectives and Recommendations,” Science Education Directorate, NSF, 1979. For an analysis of the positions taken by computing researchers and educators at this conference, see chapter 8 of Aspray, Williams, and Goldstein, 1992, *op. cit.*
67. Molnar interview, 1991, *op. cit.*
68. I. Rubincam. Summer 1987. Frequently cited authors in the literature on computer applications to education. *Journal of Computer-Based Instruction*, 14(3): 23–28.
69. Molnar interview, 1991, *op. cit.*
70. These five studies were: National Science Foundation, *The 5-Year Outlook on Science and Technology*, 1981; Office of Technology Assessment, Congress of the United States, *Informational Technology and Its Impact on American Education*, Washington, DC: Government Printing Office, 1982; National Science Board Commission on Precollege Education in Mathematics, Science, and Technology, *Educating Americans for the 21st Century: A Report to the American People and the National Science Board*, Washington, DC: Government Printing Office, 1983; Committee on Science and Engineering and Public Policy, National Academy of Sciences, *Information Technology in Precollege Education*, Washington, DC: Government Printing Office, 1984; and Office of Science and Technology Policy, Executive Office of the President, *Biennial Science and Technology Report to the Congress, 1983–1984*, Washington, DC: Government Printing Office, 1986.
71. A. R. Molnar. 1986. “Intelligent Tutors and Knowledge-Based Systems in Education,” unpublished manuscript presented at the Conference on Applications of Artificial Intelligence and Expert Systems, sponsored by the Learning Technology Institute, Warrenton, Virginia, October 29–31, 1986.
72. H. D. Huskey. 1955. Status of university educational programs relative to high speed computation, pp. 22–25. In A. W. Jacobson, 1955, *op. cit.*
73. H. Aiken. “What Is a Computer?,” pp. 85–87 in Jacobson, 1955, *op. cit.*
74. F. J. Weyl. “Summary of Conference Discussion and Proposals: Panel Discussion; Opening Remarks,” pp. 84–85 in Jacobson, 1955, *op. cit.*
75. These AFIPS numbers were reported in the Pierce Report, p. 57. For a more detailed account of computer labor statistics of the 1960s and 1970s, see chapter 8 in Aspray, Williams, and Goldstein, 1992, *op. cit.*
76. See, for example: F. P. Brooks, Jr. 1975. *The Mythical Man-Month*. New York: Addison-Wesley. Brooks was the leader of the software development for IBM System 360.
77. P. Naur and B. Randell, eds. 1969. “Software Engineering,” NATO, *Report on a Conference Sponsored by the NATO Science Committee* (Garmisch, Germany), October 7–11, 1968.
78. Statement by Dr. Milton E. Rose to U.S. Congress, House Subcommittee on Science, Research, and Development of the Committee on Science and Astronautics, March 1969, attached to the Agenda, “Fifth Meeting of Advisory Committee for Computing Activities, May 22, 1969,” H-35.

79. L. S. Peters and H. Etkowitz. November 1–19, 1988. “The Institutionalization of Academic Computer Science,” unpublished paper presented at The Study of Science and Technology in the 1990s, joint conference of the Society for Social Studies of Science and the European Association for the Study of Science and Technology, Amsterdam.
80. By 1988, student enrollments in computer programs in the United States were overwhelmingly at the undergraduate level (86%), with only 13% at the master’s level and 1% at the doctoral level (Peters and Etkowitz, 1988, *op. cit.*).
81. Keenan interview, 1990, *op. cit.*
82. For a more detailed discussion of these model curricula, see chapter 8 in Aspray, Williams, and Goldstein, 1992, *op. cit.* See Barnes interview, 1990, *op. cit.* for his experiences at the Foundation.
83. Also of importance to the national efforts to develop a standard computer science curriculum was the Computer Science Board, a group of Midwestern university computer science department chairs who began to get together to talk about curriculum. They evolved into the Computing Research Association and were the group that organized the biennial computer science conference in Snowbird, Utah, beginning in 1974. The Foundation was represented at these early Snowbird meetings by Kent Curtis. (See Hedges interview, 1990, *op. cit.*; W. Aspray. March 2003. CRA: 30 Years of Service to the Computing Research Community of North America. *Computing Research News*, 15(2): 4, 20.)
84. Statement by Dr. Milton E. Rose, 1969, p. H-9.
85. NSF, *Report of the Meeting of the Advisory Committee for Mathematics and Physical Science*, March 7–8, 1969, Office of the Director, Subject Files 1969, Records Accession No. 307-75-052, Box 2, Washington Federal Records Center.
86. K. Curtis. May 20, 1975. “University and Industry Research,” Computer Science, Directors Program Review, NSF Historian’s Files. Also see “Agenda, Mathematical and Physical Sciences Division Director’s Retreat, 15–16 November 1982,” appendix, “Degrees Awarded in Computer Science,” Office of the Director Subject Files, File MPS, Records of the National Science Foundation, Accession No. 307-87-219, Box 3, Washington Federal Records Center.
87. B. H. Barnes. October 1981. “Computer Science Mobility,” Computer Science Program Report, National Science Foundation. 6(5): 5–6.
88. K. Curtis. 1984. Computer manpower—Is there a crisis? In R. F. Cotallessa, ed. *Identifying Research Areas in the Computer Industry to 1995*. Park Ridge, NJ: Noyes, pp. 12–53.



Pre-CISE Computing Research

William Aspray

It is quite well known that DARPA (the Defense Advanced Research Projects Agency), since its creation by President Eisenhower in 1957, has supported fundamental research in networking, artificial intelligence, graphics, and other areas. However, NSF has also been an important supporter of computing research in the United States: its funding has, in fact, consistently covered an even wider range of topics, with the funding spread to a larger number of universities, compared to that of DARPA. To illustrate the importance of this support, we begin this chapter with four quotations:

The Foundation has had a very important impact on the field of numerical analysis over the past forty years—probably more than any other government agency (and perhaps more than that of all others combined).¹

It is also my impression that the United States would never have had a significant lead in this area [computer graphics] had it not been for the courage and financial support of the National Science Foundation.²

NSF has played a crucial role in the development of AR [automated reasoning] and AI [artificial intelligence] during these two decades. The funding has not been as generous as that from DARPA but the freedom extended to the researcher more than makes up for the difference.³

Generally, the NSF has made an important contribution to the development of computer graphics technology over the past twenty-five years. Those early demonstrations of the potential of computer graphics for business and the scientific community were instrumental in advancing the state of knowledge. The NSF deserves great credit for providing the opportunity for basic computer graphics development.⁴

The creation of the Office of Computer Activities (OCA) in 1967 formalized this decade-long role the Foundation had established as a patron of computer research. In the late 1950s and early 1960s, mathematicians and electrical engineers had come to the Foundation with proposals to study various aspects of computers. By 1966, the Foundation had awarded 155 grants, through the mathematics and engineering directorates, for computer science research—totaling over \$6,000,000. OCA greatly increased the support of computer science research. Over the next two years alone, 101 new grants were awarded, amounting to almost \$6,000,000.

The Foundation's computing research program fully kicked off in the late 1960s. The Mansfield Amendment to the Military Authorization Act, introduced in 1969, directed the Defense Department to divest itself of research that did not have a direct connection to specific military functions. The DoD halted more than \$300 million of research, and much of that research was redirected to the NSF. As a result, NSF funding for computing research almost doubled in 1972.⁵ Annual funding passed \$10 million, while the number of grants increased from 97 to 167. Another large infusion of funds occurred in 1976, with a 60% increase in both the number of grants and the total annual funding for research. After scaling back in 1977, growth in funding for the remainder of the decade barely kept pace with inflation. All told, over the 26 years between 1955 and 1980, the Foundation provided \$150 million to hundreds of researchers at more than 230 institutions to work on 2,700 different grants, with an average of \$57,000 per grant.⁶

The Foundation supported scientists working in every area of computer science. Software and architecture research garnered the most funds. Computer theory was the third most heavily supported area, even though theory grants were somewhat smaller. Artificial intelligence and numerical analysis were next in total funding. Graphics, databases, circuits and components, theoretical computer engineering, communications, and robotics received less support.

The Foundation's support for computer research differed from that of other funding agencies. Although the Foundation supported technological areas, such as computer architecture, circuits and components, artificial intelligence, and robotics, its grants were intended to develop scientific research rather than promote technological development. In particular, the Foundation funded a substantial amount of theoretical research of an abstract and mathematical nature.

Although the Foundation staff might shape a research agenda for funding, they did not generally set one. They relied instead on the scientific community to set the agenda, both through the proposals individual scientists submitted and the reviews by the scientific community. Because the program's direction was driven by the scientific curiosities of the applicants, it is more appropriate to consider examples

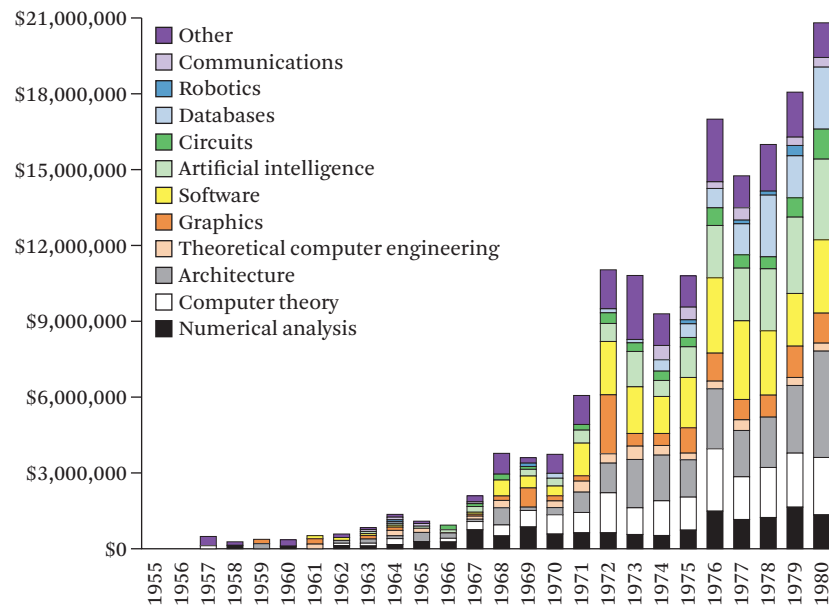


Figure 7.1 NSF funding of computer science research, 1955–1980.

of researchers sponsored by the Foundation than to analyze the formal programs established by Foundation staff. In the seven case studies that begin in the next section of this chapter, we present examples of the research done by some of the most respected NSF-funded computer scientists. Many other illustrious names could be added. These case studies are intended to give a flavor of the breadth of topics, the character and significance of the work, and the nature of the relationship between the Foundation and its grantees.

An overview of NSF funding of computing research prior to the creation of CISE is given in Figures 7.1 and 7.2. The next seven figures, placed in the sections that follow, will provide a chronological picture of funding in seven distinct research areas: circuits and components (Figure 7.3); computer architecture (Figure 7.4); software (Figure 7.5); numerical analysis (Figure 7.6); theoretical computer engineering—that is, control, network, systems, and information theories (Figure 7.7); artificial intelligence (Figure 7.8); and graphics (Figure 7.9).⁷

The Foundation provided a favorable environment for computer research and for the establishment of a science of computing. More than any other agency, the Foundation made it possible for faculty members at schools of practically every type and quality to learn about the latest advances in computing and contribute to them




Figure 7.2 NSF funding percentages by category of research, 1955–1980.

by carrying out individual research projects. The Foundation’s facilities grants, fellowships, traineeships, young investigator awards, teacher training institutes, and curriculum development grants all supported this effort, but the main instrument was the individual research grant.

One of the great strengths of the individual research grants program was that it provided academic research freedom. Grants were not awarded on the basis of the mission needs of some government agency, but on the basis of what was good science. Thus, the Foundation encouraged researchers to develop their own research agendas. The scientific community, working in tandem with program officers, created a system in which there was relatively impartial, merit-based reviewing of proposals, allowing the most promising scientific ideas to get studied. To a degree greater than at any other federal agency, the Foundation’s research program was one of and for the academic scientific community.

The Foundation took risks in pursuing good science. (We discuss below all the examples mentioned in the next several paragraphs.) Zadeh's work on fuzzy logic, Muroga's efforts at automating logic design, Hellman's cryptanalysis, and graphics research by Csuri and Greenberg were all supported by the Foundation even though there was no promise of short-term payoff. The Foundation's support of Bledsoe's analogy approach to automated theorem proving and Kuck's research on parallel systems illustrate how it would sometimes support projects that tried out speculative new approaches to research problems.

The Foundation staff understood that the research community was supporting promising researchers as much as promising research projects. The Foundation did its part in supporting this philosophy by being flexible about what research was conducted under its grants, sometimes even allowing principal investigators to change the research plan in the middle of the grant if initial results or other developments in the field warranted this. We see this in the case study of Martin Hellman below, for example.

The computer research funded by the Foundation had a strongly theoretical orientation. The best example of this was the Foundation's strong support of computer theory. This is also true, for example, of the work of Bledsoe in artificial intelligence and of Zadeh and Hellman in computer engineering theory, but this is not surprising because these fields were largely mathematical and theoretical. The Foundation's preference for theory even showed through in hardware-driven areas of computer science, such as computer architecture (e.g., Kuck's case study) and circuit design (e.g., Muroga's). When projects, such as Raphael's robotics project at Stanford Research Institute, had both hardware and theoretical aspects, the Foundation largely supported the theoretical components, leaving industry, the military, and other funders to support the hardware side.

It is difficult to assess the Foundation's overall impact on computer research, given that its support was spread across hundreds of researchers at almost as many different institutions.⁸ The Foundation nevertheless clearly had an important impact. From the early 1960s until the late 1980s, the Foundation almost single-handedly supported the field of theoretical computer science. The Foundation supported research by leading practitioners, such as Manuel Blum, Arthur Burks, Juris Hartmanis, and Richard Karp, that advanced the theory field. From an examination of the grants lists, it is clear that the Foundation actively supported every major area of computer science since the mid-1970s and has been centrally involved in many, if not all, of the major advances in computer science since 1955. Among the topics the Foundation actively supported are, for example, semantics of programming languages, structured programming, software engineering,

computational complexity, numerical analysis packages, control theory, information theory, and parallel computer design and algorithms.

The Foundation has had a less considerable, but not insignificant role in computer science applications. It should be remembered that the Foundation was expressly prohibited from funding development work and that the vast majority of its support was directed to the academic rather than the industrial sector. The Research Applied to National Needs (RANN) program, which was intended to advance research applications, was short-lived (terminated in 1977); and grants in computer science were awarded under its aegis for only one year. Some important computer applications have nevertheless resulted from Foundation funding. Wayne Cowell's NATS project was of tremendous importance to scientists in other research disciplines. Martin Hellman's research, although based in theory, had applications to cryptanalysis of great practical significance to both government and industry. The Foundation played a critical role in the transition of the Internet from a military to a public system.

What is abundantly clear is that faculty and students at almost every research university and many other two- and four-year colleges in America were able to carry out computer science research because of the Foundation's support. This support both strengthened the educational programs and added to our research knowledge. Both fortified America's industrial base.

7.1 Saburo Muroga and Computer Circuit Research⁹

We begin our analysis of research areas with the most hardware-oriented one, circuitry and components. The Foundation issued only 109 grants in this area between 1960 and 1980. Only three or four grants were active in any year, typically totaling under \$75,000. This support falls into three categories: materials used in computer hardware, computer-aided design of circuits, and miscellaneous/other. Almost 60% of the \$5,700,000 in funding for circuitry falls into the miscellaneous category. This includes efforts to support research on circuits for alternative computer architectures such as optical computers and fault testing.

Early funding in circuits was largely for either circuit devices as elements of computer hardware or "thin films," one of the promising storage technologies. Funding rose dramatically in 1972 due to a sudden and sharp interest in optical computing and superconducting circuit elements. The Foundation sponsored research on Josephson junctions, a superconducting switching element believed in the late 1960s to have great promise. These two new areas remained active research

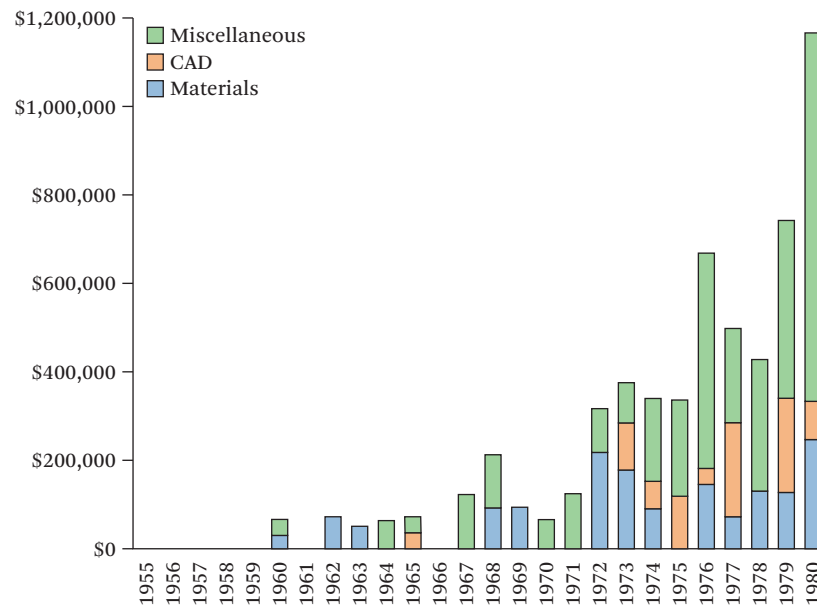


Figure 7.3 Annual funding for circuit and component research, 1955–1980.

fields throughout the 1970s, although support to them did not increase. In fact, by 1990 neither had proved itself in practice, and research in Josephson junctions was largely abandoned.

Overall funding for circuit research grew through the 1970s. However, most of the money was spent on research into several kinds of circuits used in computer hardware, including register transfers, shift registers, and multi-value sequential switching circuits. During that period the Foundation began to support research on computer-aided design systems (CAD) for large-scale integration and very-large-scale integration computer chips. This support began in 1973 and increased slowly.

Support for circuits and components was a minor part of the Foundation's computer research program. This can be attributed partly to the theory orientation within the Foundation and partly to the fact that component design was seen largely as industry's province. Nevertheless, as the work of Saburo Muroga (described below) indicates, the Foundation made a few notable contributions in this area. Most of the work on long-term, high-payoff materials (e.g., thin films) or switching devices (e.g., optical switching or Josephson junctions) did not have a short-term payoff, unlike the work on circuit design.

The value of computers rests in their ability to perform a staggering number of calculations with unerring precision. Computers gain capacity at the expense of extraordinary complexity in their design. In arranging the vast number of individual components to operate in perfect synchrony, computer designers patiently considered and reconsidered multitudes of subtle variations and evaluated each one for its effectiveness. The task was repetitive and tedious, precisely the sort of job for which computers were created. It was fitting, then, that an early application for computers was the automation of computer design.

Great promise lay in automating the design of computer logic circuitry. These circuits, which controlled and executed the calculations done by a computer, comprised numerous discrete elements, called *logic gates*. Engineers designed an intricate network of these gates for each computer function and others to coordinate the computer's operation. The logic circuitry design determined the computer's speed and efficiency. By automating the design process, engineers expected to reduce costs and improve design productivity. These goals appeared attainable because computers could easily manipulate the Boolean expressions that engineers used to represent networks of logic gates, and the high speed of computers could evaluate a vast number of working solutions to a particular design problem, screening for one that was optimal.

Many criteria existed for assessing the optimality of automatically designed logic circuitry, including the size of a circuit, its speed, the presence of redundant elements, and issues of manufacturing. A logic network was traditionally considered minimal, and therefore optimal, if it employed the smallest number of logic gates. Challenging themselves to design the simplest circuits to accomplish essential operations within computers, engineers hoped to apply computers to minimize the number of logic gates. The rewards for success at this challenge would be quicker, more reliable designs for logical networks.

The task proved difficult, however. A number of methods for creating minimal logic networks existed in the early 1960s. Conventional switching theory offered some promise, but the networks that switching theory techniques produced were often hopelessly complex designs that demanded performance characteristics of the components that exceeded the specifications of any known physical device. Physical implementation of the minimal networks developed by switching theory techniques were thus frequently impossible. Alternative techniques existed, but they were either restricted in the complexity of the networks they could design or operated too slowly to be practical. The "exhaustion" method worked by checking

every conceivable network within a range of parameters. It was used in the early 1960s but only for small networks. An automated system for designing logic networks that was clearly superior to manual design remained elusive. In the 1960s and 1970s, the emergence of very-large-scale integrated circuit (VLSI) technology diminished the incentive to design intricate control circuitry because many functions could be accomplished by using integrated circuit packages, each containing a standard logic network.

One of the few sites where research on logic networks continued was at the University of Illinois at Urbana-Champaign, where the case study of Saburo Muroga begins. After leaving his native Japan in 1960, Muroga had worked for the IBM research labs in Yorktown Heights, New York; in 1964, he joined a group at the University of Illinois working on the design of the ILLIAC IV Supercomputer. His assignment was the threshold logic gates. His background studying integer linear programming gave him unusual insight into the problem. While working on the ILLIAC IV, he began to express networks of threshold logic gates in terms of the mathematical statements of that field. He designed a minimal network by setting up a system of linear programming inequalities that represented a network of gates and the constraints on that network.

By increasing the number of gates in his network, Muroga made the linear programming problem that corresponded to it easier to solve. The mathematical problem remained unsolved until an adequate number of gates were added. As long as the linear programming problem was unsolved, there was an insufficient number of gates in the circuit. When the problem was solved, it meant that enough gates had been added. Using this correspondence, Muroga was able to assemble networks that performed as they were intended to by adding gates, one at a time, until the linear programming problem was solved. This ensured that the finished network used the minimum number of gates. The key innovation was the application of linear programming; once Muroga described networks of gates outside the traditional framework of Boolean algebra, he avoided the limitations of classical switching theory.

Muroga announced his first results in 1965, and his research group tested many different linear programming algorithms. One researcher in the group, T. K. Liu, implemented one of these methods in a program called ILLIP (Illinois Integer Programming) in 1968. The program proved highly successful after Liu modified it to discount minimal network solutions that were identical except for an inconsequential permutation of gates. In an elementary test with only eight gates and no more than three connections per gate, ILLIP found a minimal network in less than two minutes running on an IBM 360/75 computer, whereas it would have taken

a designer of the previous generation several thousand hours using an IBM 7090 and the exhaustion method.

Despite Muroga's progress, no one showed much interest in automated logic design prior to 1970. In that year, however, Muroga convinced the Foundation that his techniques could be important in the design of computer circuits, particularly those of parallel adders. His \$44,700 award in 1970 began a history of agency support to his research group that continued into the CISE era.

Although Muroga's work in the early 1970s was a continuation of earlier pursuits, he soon started off in a new direction. He had already expanded his repertoire of useful integer programming algorithms to include branch-and-bound methods, and his group developed a program to derive minimal networks using them. As the group studied the networks produced by the branch-and-bound methods, they became aware that non-minimal networks produced in the iterations of their procedure could sometimes be turned into minimal ones by applying well-known transformations. By studying the networks that could not be transformed in this way, they developed more flexible transformations and a powerful and entirely new technique for designing minimal logic networks known as the *transduction method*.

This was a major advance over integer programming techniques because it was highly efficient in the analysis of large networks. It was based upon a heuristic technique that analyzed any network generated by conventional methods, identified redundant nodes, and eliminated them. The group published results about transduction during the mid-1970s and wrote several automatic logic synthesizing programs based on them. While transduction was originally conceived as a method for minimizing networks containing only NOR gates, over time its use extended to more complex logic networks.

In the 1980s, automated logic design became more popular. Changes in design philosophy and chip technology provided a strong economic incentive to use automated design methods. Computer engineers turned increasingly to application-specific integrated circuits (ASICs). These semi-custom chips required specially designed logic networks, but their production runs were generally not large enough to justify an extensive design effort. ASICs were notoriously demanding to design. As special-purpose, non-programmable chips, ASICs' value rested primarily in their speed. ASICs had to be fast, and this requirement necessitated careful, streamlined design. Furthermore, improvements in fabrication techniques increased the density of gates to thousands on a single chip. This high number of gates increased the design complexity above a human scale. With the introduction of complementary metal-oxide-semiconductor (CMOS) technology, gates themselves became more varied and complex. The design process also had to be completed in only a few

weeks in order to remain competitive in the ASIC market. In this environment, Muroga's research became ever more important.

After 1980, other research groups created logic design systems, such as the Yorktown Silicon Compiler, SOCRATES, MIS, BOLD, and LSS, which supplemented the systems produced by Muroga's group. Industry was cautious, however, about adopting automated design, partly because of poor performance in earlier years. For example, a logic synthesizing system designed at IBM in the late 1960s by Theodore D. Friedman and Sih-Chin Yang called ALERT produced networks that used approximately 160% more gates than those designed by hand. With steady improvement in computer-aided design of logic networks, however, the method began to be routinely and extensively employed in industry. Muroga's work had only modest impact on commercial design; the primary audience for his work was the academic community. But it was through his persistence and the Foundation's willingness to fund him that these techniques were developed to the point where they could be turned over to industry.

7.2 Walter Karplus, David Kuck, and Computer Architecture Research¹⁰

The Foundation issued its largest computer science research grants in the area of computer architecture. The 282 grants awarded for architecture research between 1955 and 1980 were funded at an average level of \$78,921, 38% higher than the average computer science grant. The total support of \$22,260,000 for computer architecture during this period was almost 12% more than the amount expended on the 411 grants that supported research in computer theory.

Computer architecture grants from the Foundation were divided among four categories: traditional digital architectures, alternative digital architectures, hybrid and analog architectures, and miscellaneous research. Support for the first of these, traditional digital architectures, includes proposals to study systems architectures not described as distributed systems or networks, as well as non-specified aspects of digital architecture. This was research on systems of, and components in, machines built with classical von Neumann architectures.

Alternative digital architecture research includes computers that employed multiple processors or non-electronic technologies. The Foundation awarded its first grants here in the late 1960s, and by 1973 this line of funding was a firmly established portion of the Foundation's architecture research program. During the 1970s, the emphasis of the research changed. In 1973, when support for these architectures constituted a third of all computer architecture research, less than

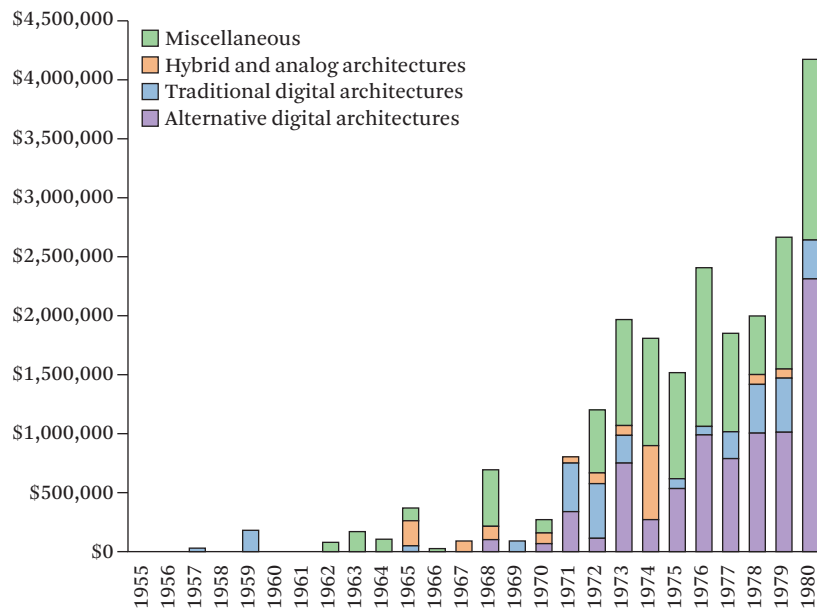


Figure 7.4 NSF funding for computer architecture research, 1955–1980.

half was applied to multiple processor machines, including parallel, vector, and array processors, and distributed systems. The plurality of alternative digital architecture research funds went to fault-tolerant systems, with approximately 10% for investigation of optical computers. By 1979, over three quarters of alternative digital architecture research was for multi-processing machines. The large increase in funding in 1980 reinforced that trend, with 20 new grants for multiple processor research.

After steady support for the study of analog and hybrid computers between 1965 and 1974, projects on these architectures waned. The Foundation awarded only one grant in the field between 1975 and 1980.

Because of the prominence of digital technology today, people often forget that engineers and scientists used analog computing equipment for almost a century before digital computers became available. For years, many applications remained the province of analog or hybrid analog-digital machines. In the 1960s, the Foundation supported Walter Karplus, a pioneer in the investigation of both analog and hybrid analog-digital machines. Karplus, an electrical engineer teaching com-

puter science at the University of California at Los Angeles, simulated large physical systems, such as air pollution and water resources, with complex mathematical models. The partial differential equations in his models were frequently too computationally demanding to resolve with existing methods, so he developed computer techniques for them. This research led to advances in both hardware and software.

Karplus published his first paper on analog computing in 1955, shortly before earning his Ph.D. in engineering from UCLA. At about that time, Ramo-Wooldridge Corporation and Convair Astronautics independently launched the first large-scale efforts at joining digital computers with analog ones. Trouble communicating signals between the analog and digital segments hampered early work, but by the mid-1960s, nearly every large company that needed to simulate physical systems, primarily aerospace industry firms, were using hybrid computers.¹¹

Karplus began his Foundation-sponsored research in hybrid computers in 1962 as a natural extension of his interest in modeling physical systems. Although his background was in analog computing, he focused on an approach to hybrid computers that emphasized the digital segment of the machine. His investigations used the class of hybrid computers known as Discrete-Space-Discrete-Time (DSDT) machines. These computers left no quantity, neither time nor space, as continuous, smoothly changing variables. Instead, each number was quantized, changed from an analog to a digital variable. DSDT machines were distinguished from the other major families of hybrid computers, Continuous-Space-Discrete-Time (CSDT), Discrete-Space-Continuous-Time (DSCT), and Continuous-Space-Continuous-Time (CSCT), which all left either space or time as a continuous variable.

A DSDT machine computed problems primarily in digital form. Small portions, however, were transferred to an analog computer and computed as analog subroutines. At these points, the system converted all the key variables from digital values into analog signals—that is, output voltages and currents, which were fed into a network of electronic components. The network relaxed, that is, reached a steady state representing an approximate answer to that portion of the problem, almost immediately. A *digitalizer* (more frequently called a *digitizer* today) reconverted the output voltages of the circuit network to digital values; and these were re-input to the digital computer, where more computations were done, if necessary. Using this hybrid method, a particular differential equation, which by purely numerical methods might take an impractical length of time to solve or might accumulate unacceptable round-off errors, could be solved quickly and with less danger of error.

The hybrid approach took advantage of the particular strengths of digital and analog computers. Digital machines were general-purpose and could be

programmed to solve almost any equation, but the solution might take an impractically long time to calculate or introduce unacceptable numerical errors in the process. Analog machines were fast and relatively inexpensive, but usually applicable to only one kind of problem. Karplus attained the generality necessary to make his work economical by using the analog segment only as a subroutine in a largely digital program. Because similar laws governed a large number of different physical systems, a small number of DSDT analog networks could be applied to many problems. With an analog network that computed general parabolic differential equations, for instance, he might solve problems in heat transfer, the diffusion of pollutants into the atmosphere, and many other applications.

Advances in digital computer technology tended to crowd out hybrid computers in the late 1960s. Karplus predicted that his digitally oriented DSDT approach would constitute a “third generation” of hybrid computers, keeping them on the forefront of research, but soon even he turned to purely digital techniques. His last Foundation grant to study linkage between analog and digital systems was awarded in 1971, to study the feasibility of direct brain-computer communications.

Karplus’s labor at hybrid architecture had been an effort to solve complex mathematical models of physical systems. With hybrid computers fading as an effective approach, he began researching other means. A Foundation grant in 1970 supported him to create a computer language to ease the creation of mathematical models. The result, the Partial Differential Equation Language, found its way into a number of national laboratories, universities, and private industry.

Karplus was adept at finding research areas amenable to his computer and mathematical techniques. He canvassed a wide range of engineers, who directed his attention to interesting problems. Through contacts with the aerospace industry, he worked on flight simulation and training controls for Apollo space vehicles. He worked with civil engineers on underground water reservoirs, electrical engineers on power distribution systems, and nuclear engineers on nuclear reaction simulators and thermohydraulic transients (important for their effect on the temperature and pressure of cooling water in nuclear power plants). The same types of equations appeared in each of these applications; his general methods for modeling were applicable to all of them.

After 20 years of nearly continuous support, Karplus ceased applying for Foundation support in 1982. By that time, he had published 70 papers and 9 books; his modeling techniques were sufficiently refined that commercial interests were eager to use them. His research contracts in the 1980s were with companies including IBM, Hughes Aircraft, TRW, Teledyne Controls, Universal Computing, Doelz Networks, and Aerojet General Corporation. He also received support from NASA and

the State of California. Karplus stands as an example of the practical engineering consequences resulting from Foundation support to computer scientists.

The introduction of parallelism in computing promised greatly amplified computational power, but also posed many challenges. These included the design of parallel machines, and also the preparation of algorithms and software necessary to take advantage of a computer's parallel processing capabilities. Many computer scientists have investigated some aspect of hardware, software, and algorithms to develop effective parallel systems; David Kuck is one of the few people who took an integrated approach. In his research, he considered all three aspects and how their interaction reinforced the performance improvements they each promised independently. The Foundation was a major sponsor of this research.

Kuck received his Ph.D. in electrical engineering from MIT, where he worked on the architecture for the time-sharing Project MAC. In 1965, he joined the faculty of the University of Illinois at Urbana, where he worked on the ILLIAC IV supercomputer. He learned how difficult the "massively parallel" ILLIAC IV was to program while writing compilers and applications for it. The supercomputer's architecture was not suited to the programming procedures he had learned. He believed that it would be easier to design a new architecture than to effect sweeping changes in programming methodology. He consequently analyzed existing programs in terms of the constructs programmers actually used to define criteria for new computer architectures. Cognizant of the potential advantages of parallel architecture, Kuck resolved to improve the match between it and its applications. Through the process of transforming sequential algorithms into ones that were as parallel as possible, he developed an understanding of the architectural requirements for parallel processors.

Kuck found little precedent for this research when he began in 1969. E. C. Russell and Gerry Estrin at UCLA had devised a graphing technique to explore program operations, but Kuck gained only limited guidance from their efforts. He adopted their global perspective as he drew charts called *dependence graphs*, which indicated the relation of one module of a program to another—that is, how the performance of one module depended on the performance of another. Complementing this global analysis, Kuck added tools for gaining a local perspective—that is, for understanding the function of the individual elements of a program. He collected FORTRAN programs from advanced scientific users at the University of Illinois and the national energy laboratories and analyzed them with his new techniques. At a very

early stage in his research, he applied for Foundation support. In 1971, the Foundation supported his proposal for Research on Computer Hardware and Software Organization.

By 1972, Kuck and two of his students, Steve Chen and Yoichi Muraoka, had completed a program analyzer and run analyses of more than 20 programs. The resulting paper attracted attention among specialists in computer architecture and scientific computation because Kuck presented the first experimental evidence that conventional programs could be automatically translated into parallel ones. Despite this success, he began to lose faith in the possibility of inferring architectures from the analyses of programs. His analysis revealed considerable diversity in the structures and techniques of the programs he examined, and he concluded that any particular architecture derived from such analyses would be so heavily oriented toward one special purpose as to be almost unprogrammable for other uses. He continued to work on his analyzer, later named Paraphrase, as a compiler of parallel FORTRAN programs, but he refocused his architectural efforts on developing building blocks for parallelism. This goal he pursued with a Foundation grant in 1973 on general-purpose parallel memories and parallel interconnection networks.

Because a parallel machine has many different processes running simultaneously, close coordination at high speeds represents a critical design issue. The memories and the networks connecting them to the processors must be able to supply the right piece of data to each processor in each and every memory cycle—a pace measured in microseconds. Kuck set out to design parallel memories and parallel networks that worked together. Just as a telephone company used a crossbar switch to enable any telephone subscriber on a network to connect to any other, Kuck needed a network that allowed ready communication between the different processors of a parallel computer. He and Muraoka figured out how to decompose crossbar switches into numerous smaller switches. He discussed his idea with scientists at Bell Laboratories, but they were not interested. The telephone company's speed requirements for automatic switching apparatus were vastly different.

One of Kuck's former students, Duncan Lawrie, at the time the head of the Computer Science Department at the University of Illinois, observed that while equal ability to connect between each pair of nodes was valuable in some networks, it was a wasted resource in others. Many networks, including telephone switchboards, utilized only a few of their possible connections, but those that were used were used frequently. Lawrie described a class of networks, called *omega networks*, in which connections could be made at speeds required by parallel computers. Paraphrase was an important tool in this development. Lawrie simulated the execution of a program on a parallel computer and, using Paraphrase, kept statistics on the ex-

change of data among the different processors and memories. Statistics generated by the Paraphrase analysis of many different programs helped Lawrie to describe the connection patterns characterizing omega networks.

The development of omega networks opened the door for parallel computers with a far greater number of processors. Parallel machines designed in the 1960s using crossbar switches were limited to four or eight processors. With an omega network (also called a *shuffle exchange network*), Kuck envisioned computers with hundreds or thousands of processors. In the following years, some computers, such as IBM's RP3 and the Alliant FX/8 computer system, were built using the shuffle exchange network principle. The research Kuck carried out with his students also informed the design of other machines. These included the Cray X-MP and Y-MP supercomputers (designed by Kuck's student Steve Chen), the Burroughs Scientific Processor Project computer, and machines of the three major Japanese supercomputer manufacturers: Fujitsu, Hitachi, and NEC.

The shuffle exchange network was not the only approach to parallelism. There was ongoing debate over the relative merits of a shuffle network versus the crossbar, mesh, and hypercube systems. Kuck continued to champion the omega network developed by Lawrie. He implemented it in the CEDAR system, constructed at the University of Illinois Center for Supercomputing Research. He used CEDAR as a test bed for other architectural innovations. CEDAR was built from a small number of Alliant computers—each one a parallel machine using a shuffle exchange network—linked together in a cluster with a common memory. Kuck linked several of these clusters together and they, too, shared another common memory. The linking process could be iterated through many levels. This concept has been explored by IBM and several Japanese computer companies.

The Foundation supported Kuck's work on memory hierarchy management, such as CEDAR, from the 1970s. Through the use of Paraphrase, Kuck discovered ways of managing the paging of a parallel program in virtual memory. *Paging*, already well established in non-parallel machines, improved a computer's performance by moving standard-sized units of a program between the computer's fastest memory and its slower, lower-level memory. This technique freed large blocks of the computer's faster memory. It was exploited to great advantage by Kuck's colleagues working on parallel routines for commonly used linear algebra operations. This collection, called BLAS3 (Basic Linear Algebra Subroutines), was so successful that Kuck himself applied the memory management techniques to workstation architectures he was developing. Using ideas generated by BLAS3, along with other architecture innovations, including RISC (Reduced Instruction Set Computer) processors and memory caching, Kuck's company created hardware in the 1980s that

made one-processor workstations run up to 30% faster. This was another example where Kuck's research on parallel software led to architectural improvements.

Kuck's company, an independent commercial venture, was not the sole outlet for his research. Kuck, Lawrie, and Ahmed Sameh, a long-time collaborator with Kuck on mathematical software, organized and operated the Illinois Center for Supercomputing Research in the early 1980s with funding from the Foundation, the Department of Energy, and later the Air Force and DARPA. The Center carried out the research on the CEDAR system. By supporting Kuck's work, the Foundation fostered substantial improvements to parallel computing, supercomputer and workstation design, and advanced scientific computation.

7.3 **Mary Shaw, Barbara Liskov, and Software Research**¹²

The Foundation devoted more funds to research in software than to any other area of computer science. Between 1954 and 1980, it granted over \$24,000,000 to software research, more than 15% of all computer science research funding. The grant size was relatively large; 331 grants averaged \$73,200 per grant. Numerical analysis and computer theory were the only fields that received more grants, and larger grants were awarded only for computer architecture and robotics.

The Foundation's support for software research falls into five categories. One involves the development of languages, compilers, and operating systems. A second deals with the development of applications software—that is, programs that facilitate the work of some application, such as economics, mathematics, or education. The third area, software design, examines the general processes of software creation. This research was directed toward methods for efficiently producing high-quality software including software engineering, software verification, and reliability. Questions about implementation, including software portability and related issues, constitute the fourth category; and the fifth category is miscellaneous software research.

Languages, compilers, and operating systems comprised the software research area receiving the most support from the Foundation between 1961 and 1980. Awards in this area constituted the bulk of funding in 1968, when annual support for software research first exceeded \$500,000. In that year, software applications were funded for the first time. Applications funding was at first focused on computer-aided instruction, but the emphasis on education disappeared when the education initiatives were moved out of the Office of Computing Activities. By the late 1970s, approximately half of the funding in this area was dedicated to research on applications to assist scientists in mathematical disciplines, such as econom-

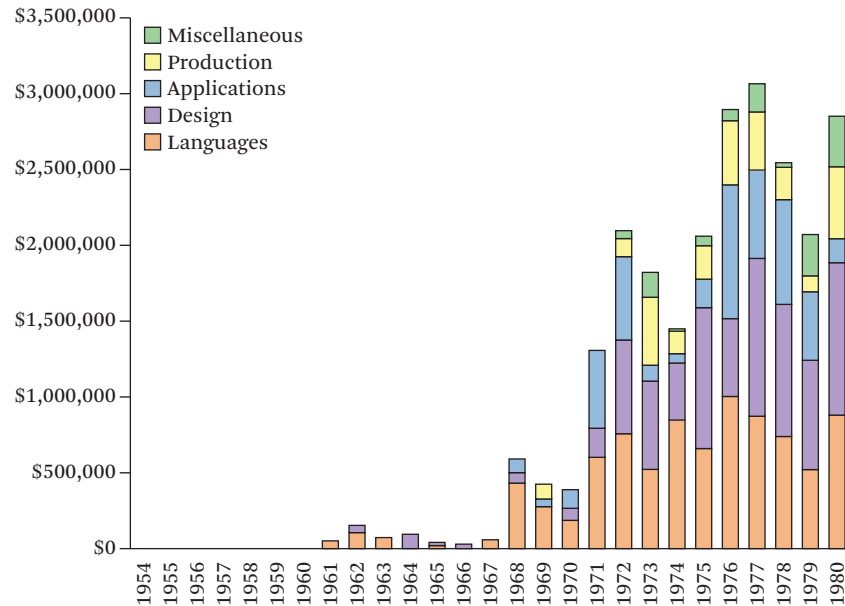


Figure 7.5 Annual funding for software research, 1954–1980.

ics. The rest was committed to the study of applications computer scientists used in their own work.

Some of this work overlapped heavily with issues considered in software design, an area that became of critical importance to computer professionals during the 1970s. Foundation funding for software design expanded markedly when questions about design techniques started to attract the academic and commercial software communities. The Foundation first funded software design research at a comparatively significant level in 1972. From then on, throughout the 1970s, funding for this area of software research predominated over all other areas.

In the late 1960s, computer professionals became acutely aware that expensive computer hardware was not being fully utilized. One problem was that high-quality software could not be acquired in a timely or economical way. Although many programming languages existed, most software was still being written in the aging standards: COBOL, FORTRAN, and LISP. The sentiment grew that these languages interfered with efficient programming. Concern about this so-called “software crisis” prompted the NATO Science Committee to convene a conference in 1968 to

investigate. Participants introduced the notion of software engineering; and papers by Edsger Dijkstra, Niklaus Wirth, and others elaborated this concept. From these papers emerged the idea of structured programming, an approach to programming computers that would check the tendency of programs to develop along excessively complex, convoluted lines. With Foundation support, Wirth derived the language Pascal from ALGOL 68 as one of the earliest ventures into implementing structured programming practice. This research brought structured programming to center stage.

Another group working along similar lines included Mary Shaw and William Wulf at Carnegie Mellon University, who collaborated with Ralph London at the Information Sciences Institute at the University of Southern California. They noticed that the most popular programming languages did not contain features that encouraged programmers to write with a structured approach. They began work in 1974 on a programming language called Alphard as a practical tool for constructing high-quality software. The team wanted Alphard to bring together recent developments in programming methodology and program verification, particularly the concept of abstraction. Contending that even moderately difficult programs are too complex for the human mind to grasp, they seized on abstraction as a way for the programmer to retain cognitive control of the program. By designing Alphard to allow easy and useful abstractions of intricate tasks and structures, Shaw and her colleagues hoped to facilitate straightforward conceptualizations of computer programs.

The most fundamental abstraction technique used in Alphard was the user-defined abstract data type. In a program module called a *form*, the programmer completely specified all the properties of a particular type of data in terms of pre-existing types. The specification described a data structure that was a representation of the abstract construction in the mind of the programmer. Programmers were already familiar with a few data abstractions such as stacks and arrays. These common data types were pre-defined in the familiar languages. Programmers needed only to declare that a variable was one of the pre-defined types, and the language would automatically know how to process instructions involving that variable. The language prescribed all possible manipulations of the variable.

Shaw's group was convinced that dependence on pre-defined data types sometimes prevented appropriate data abstraction and hampered clear, efficient programming. They pointed out that FORTRAN, which offered only rectangular arrays, created problems for programmers who wanted to use triangular matrices, a common form in scientific computation. In order to be more economical with memory space, programmers often packed two triangular matrices into the same square ar-

ray, rotating the orientation of the matrices to achieve the most compact storage. This practice confounded subscribing conventions and made locating any particular matrix element a process requiring an extra calculation, thereby increasing the possibility for error. Shaw's team hoped Alphard would bypass the need for any special routines that translated between the way the user thought about the data and the way the computer represented it.

Alphard was devoid of all but two pre-defined data types: *rawstorage*, a set of contiguous, addressable, untyped storage locations, and *boolean*, variables with only two values, representing "true" and "false." Alphard incorporated powerful tools to define all other types from these. An Alphard program started with forms declaring all the data abstractions. The language separated the specification of the data abstractions, where all the qualities of the data type and the legitimate operations that could be performed on variables of that type were declared, from the section describing their implementation by the computer. Any specification might have several possible implementations. The programmer chose the implementation considered most appropriate to the computer system on which the program was running. A publication in 1978 described the language informally and emphasized the importance of the difference between specification and implementation.

Shaw's group also confronted the problem of providing assurance that implementations faithfully represented specifications made in Alphard. They found their solution in C. A. R. Hoare's 1972 paper, "Proof of the Correctness of Data Representations."¹³ Hoare introduced an instruction function to prove that any given specification and implementation were consistent with each other. This allowed Shaw's group to conceive of a language that separated the two elements and left their elaboration in the hands of a programmer. In papers describing Alphard, the designers wrote sample forms for common data structures such as stacks, queues, and binary trees, as well as for several more unusual ones. The verifications for these forms were each a mathematical proof ranging in length from a few lines to several pages.

For many applications, conventional data types were perfectly adequate for user needs. Shaw and her colleagues included, with the Alphard compiler, verified forms describing many of the most common data types, such as integers and vectors. Alphard programs had, conceptually at least, a standard prelude that declared the common data types. More exotic data types were in libraries of forms. Once a form had been validated, it could be used with confidence in any application demanding that particular data type. To increase efficiency, the Alphard language allowed developers to reuse code as extensively as possible.

Alphard proponents contended that abstraction would produce better programs. The designers knew that, to convince others of their approach, they must consider the completeness of specification and verification, quality of programs, life-cycle costs, and reusability of programs. They believed that the ability to implement the language on different hardware systems was also crucial. To test the prospect, they implemented in Alphard a piece of software that had been developed in a conventional software environment and compared the results. They chose a text editing program of medium length and recast it as an Alphard program. Although the initial results they reported in December 1978 were encouraging, at that point they had only completed a specification in prose.

By 1979, Shaw and her colleagues believed they had demonstrated the value of data abstraction in the production of superior code. They disseminated their results through publications and extensive contact with other groups studying similar issues. For example, Alphard cross-pollinated influence with Barbara Liskov, who was working at MIT on a data abstraction language called CLU. Alphard also reached the developers of other data abstraction languages such as Euclid, Gypsy, Mesa, Concurrent Pascal, Modula, and ADA. The concepts underlying Alphard had become so widespread by 1979 that Shaw and her group no longer felt that the language was urgently needed. Other languages, influenced by the work on Alphard and sporting comparable qualities, had appeared. Instead of completing the tedious details of Alphard, Shaw and her colleagues turned to new projects.

The impact of software engineering is plainly evident in the research career of Barbara Liskov, a computer scientist at MIT. She received Foundation support during the last half of the 1970s to develop a programming language called CLU, which, like Shaw's Alphard, promoted good design practices through the use of abstraction. Liskov used her CLU experience in the 1980s to explore software issues raised by new architectures. The Foundation supported both of these stages of her research.

The ambitions that motivated Liskov to work on CLU were similar to those of Shaw's group at Carnegie Mellon. Liskov hoped that the abstractions that CLU encouraged would result in more effective programming: faster production of software that was both of higher quality and easier to maintain. She emphasized three sorts of abstraction: data, procedural, and control. Data abstractions were achieved through program items called *clusters*. In clusters, Liskov carefully segregated the description of a data abstraction from its implementation, as Shaw did in Alphard

forms. Procedural abstractions were ways of thinking about the operations the program performs, and control abstractions influenced the sequence of those procedures. Liskov's control abstraction supplemented the *if*, *while*, and *for* commands familiar from other languages.

CLU differed from Alphard in several ways. CLU employed semantics similar to the programming language LISP, while the semantics of Alphard were closer to those of Algol. CLU had no particular formal specification mechanism. In CLU, new data types were defined algebraically—that is, axioms defined an algebra that expressed the data structure. On the other hand, Alphard expressed new types in terms of already existing types. The similarities between the two languages are more important than their differences, however. Research on CLU, like the simultaneous work on Alphard, solidified the role of data abstraction as a practical approach to computer programming.

Liskov never intended CLU to be a commercial programming language. When she completed it in 1979, she did not seek funding to develop a production-level compiler, nor did she attempt to find an industrial partner to market CLU. She distributed implementations for the DEC 20, DEC Vax, and Motorola 68000-based Unix systems to universities for a nominal fee. Over 200 sites acquired the language, but CLU received its most serious use at MIT. Liskov herself used CLU to begin research on software for distributed systems, a new type of computer configuration.

Distributed systems feature any number of computers—not necessarily of the same make—operating simultaneously at geographically distinct locations. Much was known about the architecture of distributed systems, but little about how to program such machines. Reliability was critical. In a distributed system, any one of the processors working on an element of a program might fail to accomplish its particular assignment. Liskov recognized that the capacity of the system to continue operation despite the failure of any single component constituted an essential feature of distributed system software. Along with these reliability issues, Liskov concentrated her development efforts on understanding how to program the concurrent activities of multiple computers so that they operate efficiently.

CLU was useful not only because it encouraged good programming, but also because it was an object-oriented language. This meant that programmers conceived of programs in CLU as operating on objects, such as databases and files. Major applications of distributed systems, such as airline reservations networks, deal with just such objects. When Liskov resolved in the mid-1980s to develop a new language dedicated to run on distributed systems, she decided to retain this characteristic of CLU. She began work on a language called Argus, addressing the question of robustness in the face of component failure. In designing and implementing Argus,

she adopted an idea from database management software—the atomic transaction, which she generalized for use with distributed systems software. In Argus, programs were constructed out of a number of atomic transactions, which were computational tasks that either succeeded completely or failed completely (as if they had never run at all).

Liskov had an early version of Argus running as early as 1986, and within two years she was writing fairly sophisticated distributed programs with relative ease. She used Argus to write several applications, including an electronic mail system, an editor that permitted many users at different locations to work on the same document, and a game played simultaneously by remote users. The approach using atomic transactions influenced other researchers and stirred debate but was widely accepted. Al Spector, a researcher at Carnegie Mellon, drew inspiration from Argus in his work on Camelot, software intended to support lower-level distributed computing. Argus also influenced Avalon at Carnegie Mellon and EDIE at the University of Washington. Argus's lack of portability, however, prevented it from being implemented outside of MIT and thus limited its influence.

Liskov's work on Project Mercury concerned heterogeneous computing in a distributed network. Answering a need introduced by the increasing use of computer networks, which involve computers of many different types and powers, Project Mercury investigated how to create a program when components of it were written in different languages and implemented on different machines. Project Thor, started in 1990, concerned object-oriented databases on a heterogeneous computer network. The goal was to facilitate the harmonious sharing of data across a distributed system. By the second year of Project Thor, Liskov designed implementation strategies for her databases using Argus. Network-oriented databases demanded a great deal of data replication to ensure that stored information was locally available. Thus, Liskov began work on a file replication system. That work proved so fruitful it was spun off into its own project. Liskov hoped to achieve implementation of the system, called Project HARP, by the end of 1991.

Much of Liskov's research was supported by block grants from DARPA to the MIT computer science laboratory, but she also received NSF support. She found that Foundation support provided independence from larger laboratory and military objectives to pursue basic science. She also valued the Foundation's peer review process, which informed her of the research community's regard for her work. Foundation funding was important in enabling her to keep her work at the leading edge of software research.

7.4 Wayne Cowell, David Young, and Numerical Analysis Research¹⁴

Numerical analysis antedates electronic computing. There is a centuries-long tradition of applying new computational tools to mathematical problems that defy analytic solution, including by such luminaries as Isaac Newton and Carl Friedrich Gauss. Numerical analysis stagnated, however, after Gauss's contributions in the nineteenth century. The introduction of the high-speed digital computer in the 1940s changed the situation. Computers were first used to speed up the calculations demanded by numerical methods that had been developed prior to computers, but computers were soon being used on numerical problems never attempted by hand or with desk calculators. The computer transformed the field of numerical analysis.

Although the Foundation's facility grants helped equip schools with digital computers for science, it also directly supported research in numerical analysis. Between 1955 and 1980, the Foundation awarded 412 grants, totaling nearly \$14,000,000, for such work. The 412 grants numbered more than for any other category of computer science research, yet the total funding for numerical analysis was less than 9% of all Foundation funding for research in computer science. The average Foundation grant for numerical analysis amounted to \$33,500. This was the smallest grant size out of all categories of computer science, reflecting the mathematical rather than experimental character of numerical analysis. Expenses for numerical analysis research were low; the Foundation awarded only 16 grants of more than \$100,000 between 1955 and 1980.

The Foundation's numerical analysis research divided into nine sub-categories. The largest of these, general numerical analysis, is a catch-all that includes all grants with titles that obscure the precise topic of the research. It is most likely that these grants, numbering over 100, were for the study of differential equations, non-linear mathematics, probability, or one of the other topic-specific sub-categories. Support for differential equations far outweighs support for non-linear mathematics, but, in fact, many grants for differential equation research considered non-linear cases, and much of the non-linear mathematics research looked at differential equations. It is best not to compare these two areas, rather to note how support for the two compares to areas such as discrete mathematics, optimization, and approximation.

By the late 1960s, there was research activity in mathematical software, as 59 grants by the Foundation indicate. Funding for mathematical software, \$130,000 in 1968, declined in the following years but was revived twice in the 1970s by

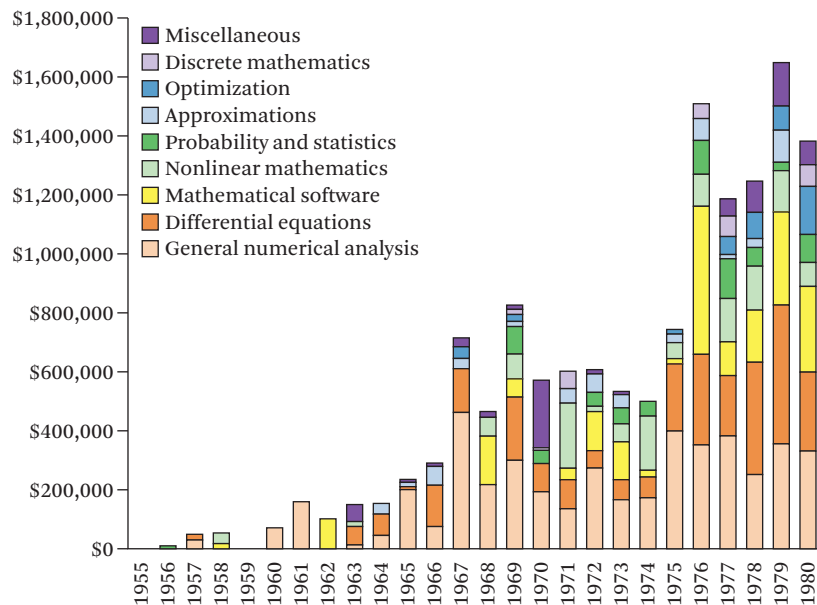


Figure 7.6 NSF funding for numerical analysis research, 1955–1980.

the large-scale PACK projects organized under Wayne Cowell of Argonne National Laboratories. This is the subject of our next case study.

After its introduction in the 1940s, the computer rapidly replaced desk calculators and punched-card tabulating systems in scientific computation. John von Neumann noted that the computer had a different “internal economy” from previous calculating equipment. For the computer, arithmetic was inexpensive but storage was precious. With earlier equipment, mass storage was available in the form of punched cards and other technologies but arithmetic operations were limited. The new economy of computing machinery completely altered the research agenda for numerical analysis. The computer made it possible for the first time to employ techniques that demanded many computations, such as approximation methods for non-linear partial differential equations and probabilistic methods such as Monte Carlo. Mathematicians and other computationally oriented scientists in the late 1940s and throughout the 1950s explored numerical analysis with vigor.

There were problems, however. Although researchers developed new numerical methods to exploit the capabilities of computers, these methods were not always widely available to researchers. The proliferation of incompatible computing machines and machine-specific languages limited the utility of any given piece of mathematical software, and mathematicians had little incentive to package their numerical methods into software for general distribution. It was only in the mid-1960s, as computers became more common in research environments, and as computer systems and languages became more standardized, that an opportunity opened to implement these new routines for general use. The ACM Special Interest Committee on Numerical Mathematics, founded in 1966 by Joseph Traub, who later became the chairman of the Carnegie Mellon computer science program, stressed the need for high-quality machine implementations of numerical algorithms. In 1969, John Rice of Purdue University coined the term *mathematical software* and called for computer programs that implement widely applicable mathematical procedures. Several independent efforts at compiling libraries of numerical routines were launched. In 1971, British mathematicians James Hardy Wilkinson and Christian Reinsch pulled some of this early work together into a handbook entitled *Linear Algebra*.

Wilkinson and Reinsch's procedures were elegant, efficient software implementations of well-known numerical methods for solving linear systems. Derived from work published in the German journal *Numerische Mathematik* during the 1960s, they were written in Algol, a language few American scientists used, and so the programs were of limited appeal in the United States. Virginia Klema, a researcher at Argonne National Laboratory, started translating some of them into FORTRAN, the most popular programming language among American scientists. This work excited Wayne Cowell, another Argonne scientist. Already in 1970, he had conceived of a team effort to employ mathematicians from Argonne, the University of Texas, and Stanford University to produce mathematical software, and he arranged for support from both the Foundation and the Atomic Energy Commission. His team began working on high-quality, certifiably correct FORTRAN versions of the Wilkinson-Reinsch algorithms. In 1971, the Foundation issued awards to Cowell, as well as to Y. Ikebe at Stanford and Cleve B. Moler at the University of Texas, to produce, test, and distribute these mathematical programs.

It was an unusual partnership. As a national laboratory, Argonne was funded by the Atomic Energy Commission and did not customarily rely on other government agencies for support. Since the other two participating institutions were universities, however, the researchers saw the Foundation as an appropriate source of funding. Cowell secured the Foundation's support when he pointed out that,

although the project would be centered at Argonne, much of the work would be done at Texas and Stanford.

The project was called NATS, reflecting the partnership between the NSF, Argonne, Texas, and Stanford. Cowell's group prepared FORTRAN versions of the key Wilkinson-Reinsch routines before the end of 1971. The versions went out to 20 university, industrial, and government laboratories for testing on a wide range of computer systems. Argonne collected the returns and integrated the suggestions into a finished package, called EISPACK, released in May 1972. There were versions of each piece of software for each of seven popular scientific computers: IBM 360 and 370, CDC 6600 and 7600, Univac 1108, Honeywell 635, PDP-10. The software efficiently calculated eigenvalues or eigenvectors for a variety of matrix types. The software drivers called between one and five subroutines, each performing a single operation on a matrix. An advanced user could summon these subroutines in complex and innovative ways to solve broad classes of problems.

EISPACK was a great success. Users of all levels of expertise employed its routines. Distributed through the Argonne Code Center (later, the National Energy Software Center) and the independent corporation IMSL, EISPACK was distributed to over 1,000 sites. Argonne scientists disseminated the routines further by inserting them as components in other programs, which were shared with scientific users. The routines in EISPACK spread quickly. EISPACK benefited from its exceptional portability; the code ran on a wide variety of computers. The portability and the overall quality made EISPACK popular in scientific computing circles.

The first release of the package in 1972 was only the beginning. The NATS team needed to prepare documentation for the routines. The group also began to extend the routines. In 1974, the Foundation made awards to Cowell, Edward Ng at the Jet Propulsion Lab, and Henry Thatcher at the University of Kentucky, and small grants to a number of test sites for a second version of EISPACK. The project was dubbed NATS II, but the acronym was stripped of its old meaning and reinterpreted to stand for the National Activity to Test Software. The new name stressed the Foundation's particular interests. Both NATS and NATS II produced mathematical software packages, but the Foundation conceived of the projects as efforts to study the process, including production, certification, and dissemination, by which that software was created.

NATS II funding was spent also to continue work on another project that the NATS group undertook. This was a package called FUNPACK, a collection of numerical routines for more specialized mathematical operations, such as Bessel functions, complete elliptical integrals, exponential integrals, Dawson's integral, and the psi function. Since these functions were less commonly used than the eigenoperations covered by EISPACK, FUNPACK's intended consumer base was much

smaller. The NATS group released FUNPACK in 1973 to a small but enthusiastic community of users. Versions for IBM, Univac, and CDC computers were made, but transportability was not given as much prominence. Taking the position that high-performance function programs were by their nature highly machine-dependent, the FUNPACK group, under the leadership of W. J. Cody, placed little emphasis on machine independence and transportability. Indeed, the group celebrated the deliberate machine specificity of its routines as being useful to their performance.

The FUNPACK group explored issues involved in generating a software package. With the NATS II grant, the Foundation stipulated that the research should focus on the methodology of production, certification, distribution, and maintenance of mathematical software. The authors of FUNPACK felt strongly that “NATS [(and NATS II)] was not funded to produce software, it was funded to study how to produce software.” They recognized that FUNPACK’s impact was not as significant as EISPACK, but they believed the effort was just as valuable. Others at Argonne, who were looking for new mathematical software to result, did not share this view.

Among the detractors of FUNPACK were the authors of LINPACK, a 1976 Foundation-supported effort similar in purpose to EISPACK. LINPACK was a collection of routines to analyze and solve linear equations and linear least-square problems. Like the EISPACK routines, these were originally written by Wilkinson in Algol, and the project translated them into FORTRAN. In 1976, when the NATS II grant expired, the Foundation committed its support to the LINPACK proposal with Cowell as the principal investigator. Other funding went to J. R. Bunch at the University of California at San Diego, G. W. Stewart at the University of Maryland at College Park, Cleve Moler at the University of New Mexico, and a few testers. The project leader, Jack Dongarra at Argonne, explained that although “LINPACK was not funded as a development project . . . it is safe to say that the authors were more interested in development of the package than in software research.”

The PACK projects, highly esteemed by the computational science community, stimulated considerable research on both mathematical software and software methodology. In the late 1970s and 1980s, mathematicians developed numerous software packages bearing names such as ELLPACK, FISHPACK, ITPACK, MINPACK, PDEPACK, QUADPACK, ROSEPACK, SPARSPACK, and LAPACK. Each one incorporated routines for a particular kind of mathematical problem. These packages increased the availability of mathematical software and solidified the professional legitimacy of numerical analysts working as programmers. Insights about software production gained in the PACK projects were also applied in the creation of TOOLPACK. It was a collection of programming tools designed to help developers create high-quality scientific software. It included routines to analyze FORTRAN programs, detect semantic errors, reformat them, and carry out other tasks that the

authors of EISPACK, LINPACK, and other mathematical software previously had to program anew each time.

Cowell worked on TOOLPACK in the late 1970s at the suggestion of Web Miller of the University of California at Santa Barbara. The Foundation and the Department of Energy provided his initial funding, but the Foundation dropped out once the project was established. DOE continued support of TOOLPACK alone. The Foundation, however, continued to support related research, such as Dongarra's work in the 1980s on LAPACK.

David Young of the University of Texas at Austin was one of the several numerical analysts who received a small grant as a tester for the LINPACK project and later developed a PACK himself. When he assumed LINPACK testing responsibilities in 1976, he already had a long history of Foundation-supported research. His research concentrated on iterative methods for solving a particular class of linear systems. The matrices associated with these systems were "sparse,"—that is, they had few elements with other than zero value. These matrices were closely associated with elliptical partial differential equations, which commonly occurred in diffusion, steady-state heat flow, and fluid dynamics problems. Young's research contributed, for example, to describing the flow of oil in water and for designing sections of nuclear reactors.

Young began his work on iterative methods for the numerical solution of linear systems in the late 1940s as a graduate student of Stanley Frankel at Harvard. In 1950, they developed an influential method known as successive overrelaxation (SOR). Young worked at refining the method during the 1950s and determined key parameters necessary for its application in particular cases. By 1963, when he received his first grant from the Foundation, SOR had become the preferred method for addressing linear systems. The Foundation supported him to improve computational aspects of SOR.

Engineers and scientists in other disciplines frequently called on Young to assist them with the solution of difficult numerical problems. His research in applied mathematics equipped him with an arsenal of useful iterative techniques for linear systems. Realizing that his techniques had wide scientific applicability, Young decided, at the suggestion of Harvard mathematician Garrett Birkhoff, to make them available in a software package.

A package of iterative routines was problematic, however, because of the special attention that must be given to selecting appropriate operating parameters. Young

had painstakingly worked to develop techniques for determining the omega parameter used in SOR; it was difficult to imagine a generalized algorithm that would reliably produce accurate solutions iteratively. He and his collaborators engaged in years of intensive research. He issued results of his work in a package as early as 1978. In the early 1980s, he released ITPACK 2C to the scientific community. The package included Jacobi, SOR, Symmetric SOR, and RS method algorithms. Young used mathematical techniques, such as Chebyshev and conjugate gradient methods, to accelerate the performance of these algorithms.

ITPACK 2C was eventually consolidated with ELLPACK, another package of routines for solving linear systems of elliptical differential equations. Young then turned to extending the power of his routines. The Foundation was interested in seeing him recast the iterative algorithms in ITPACK to take advantage of more of the sophisticated computer hardware available in the 1980s. He wrote iterative algorithms that ran on vector and parallel machines, not just the VAX computers that supported the original ITPACK. These new architectures enhanced algorithm performance because of their higher calculation speed. He discovered some tradeoffs. The algorithms best suited to parallel computation were not always the ones that converged on a solution most quickly. Nevertheless, using the increased computing power, such as the Cray supercomputers he used for his own research, Young was able to make significant inroads on complex problems.

For most of his career, Young worked with symmetric, positive definite matrices—much of this research funded by NSF. By applying the computational power available to him in the 1980s, however, he approached the more intractable case of non-symmetric matrices. This enabled him, for the first time, to investigate realistic models of physical systems. The extension to non-symmetric methods vastly increased the applicability of his work. His colleagues in the petroleum engineering department at the University of Texas used his methods to study diffusion across the oil-water boundary in an effort to improve oil recovery techniques. His mathematical techniques were applied to describe the vacillations of the boundary between water and oil when water was pumped into an oil reservoir. It had formerly been necessary to introduce unrealistic simplifications to keep the matrices symmetric.

Young continued to improve ITPACK. He released a version for vector computers. He collected iterative methods for non-symmetric systems into the packages ITPACK 3A and ITPACK 3B. He improved the efficiency of his technique to store matrices and thereby reduced the memory requirements for systems running ITPACK. These packages were widely distributed to university, industrial, and national laboratories.

Young's work on packages of useful numerical routines was but one part of his work in numerical analysis. He also worked with emerging new numerical methods. In the closing years of the 1980s, he worked on solving equations describing complex three-dimensional systems that evolve over time. These models were the most realistic he worked with; the equations were, correspondingly, the most computationally demanding. He also studied multigrid methods, a complex numerical approach for solving equations that required few iterations before determining a solution. Multigrid routines posed a special challenge for inclusion in a software package because their complexity demanded that the user have good understanding of their principles. Young's work on multigrid routines was funded principally by the Foundation and the Department of Energy.

7.5 Lofti Zadeh, Martin Hellman, and Theoretical Computer Engineering¹⁵

Some research funded by the Foundation has investigated topics at the boundaries of engineering, mathematics, and computer science. Not surprisingly, these proposals may fall under one of several directorates. Using non-standard language, we call this research *theoretical computer engineering*. It includes research in information and coding theory, systems theory, network theory, control theory, and other related work.

Foundation support for theoretical computer science (e.g., automata and complexity theory) has been modest; the \$5,000,000 distributed over 112 grants up to 1980 is only a little more than 3% of the total that the Foundation awarded for computer research. Only numerical analysis grants were smaller on average. In the early days of Foundation-sponsored research, however, theoretical computer engineering played a more significant role. In 1961, the Foundation committed \$300,000, nearly 60% of its computer research budget, to theoretical computer engineering grants. Most of those funds went for systems theory research. Of all theoretical computer engineering categories, systems theory won the most Foundation support during the 1960s. The majority of these funds went to the University of California at Berkeley, where Lofti Zadeh and Charles Desoer were doing important work.

In the 1970s, as the Foundation's budget for other areas of computer science ballooned, support for theoretical computer engineering increased only slightly. There was a shift in emphasis during this decade from systems theory to work on information and coding theory. As excitement about the latter subject, with its applications to data communications and cryptography, grew during the 1970s, the Foundation increased its support. Between 1975 and 1980, fully 70% of funds

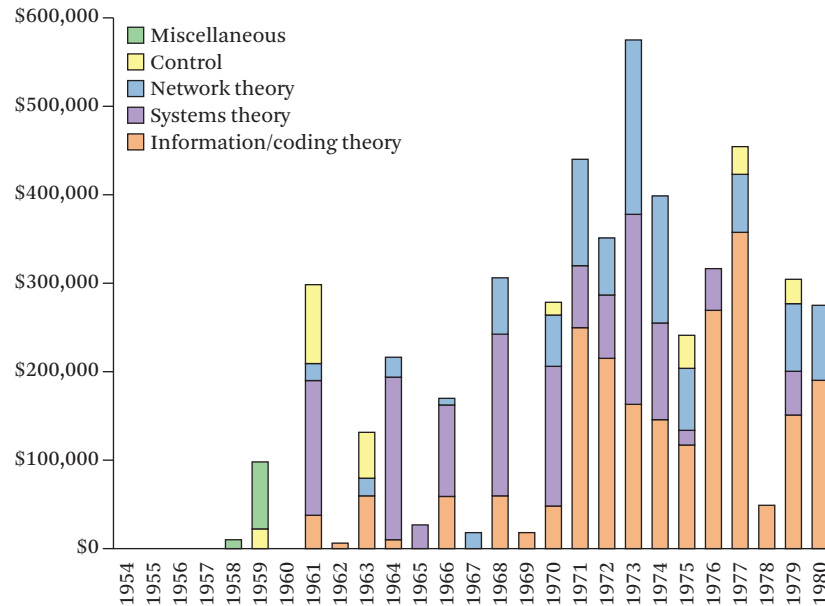


Figure 7.7 Annual funding for theoretical computer engineering research, 1954–1980.

for theoretical computer engineering went to the study of information and coding theory.

The first phase of Foundation support for theoretical computer engineering is best exemplified by the career of Lofti Zadeh. His research began before the Foundation was founded. As an electrical engineering doctoral candidate at Columbia in the late 1940s, he worked on systems analysis, creating mathematical theories to explain the behavior of systems of electrical devices. His 1949 dissertation on time-varying networks introduced the concept of the *time-varying transfer function*. In that year, he joined the faculty at Columbia, continuing his work on systems analysis and related areas.

In 1950, Zadeh and John R. Ragazzini published a paper that generalized a theory of time-series prediction proposed by Norbert Wiener two years earlier. In 1952, they published work on sampled-data systems, proposing a *z*-transformation that made analysis of such systems similar to the analysis of more familiar linear time-invariant systems. In both of these papers, Zadeh explored uncertain events with probabilistic methods—a tool that would figure prominently in his later research.

Zadeh first received Foundation support in 1957; the result was a well-received book, written in 1963 with Charles Desoer, describing the state-space theory of linear systems. The state-space approach became the standard technique for optimizing linear control systems. Zadeh also contributed to the theory of non-linear systems by defining a hierarchy of them. His work facilitated the design of non-linear filters and predictors, which were more effective than their linear counterparts. This work found many military applications involving signal processing, fire control, and trajectory calculations.

Despite his important incremental improvements in systems analysis, Zadeh was pessimistic about applying existing systems analysis approaches to highly complex systems. As early as 1961, he began to think that the richness of human systems stemmed from their use of imprecise judgements and shades of interpretation—an imprecision unknown in the electromechanical systems with which he was working. He was convinced that any adequate approach to analyzing complex systems would involve the flexibility to handle intermediate, indefinite quantities. He believed that models of complex systems should not be circumscribed by rigid quantifiability but that variables able to take imprecisely defined, fuzzy values should be introduced.

Radical suggestions like this might have been derided or ignored by the research community had they originated from a less well-respected scientist. Zadeh, however, was known as an innovative and productive theorist. His papers appeared in respected journals, and his methods were widely circulated. In 1964, he applied to the Foundation for a grant to develop a theory of sets without sharply defined boundaries. His past record of impressive results with Foundation support earned him the benefit of the doubt, and he received a grant to study these so-called *fuzzy sets*.

Zadeh published his first description of fuzzy sets in 1965 and worked exclusively in that area thereafter. During the late 1960s and early 1970s, he refined the concepts of his theory with an eye toward constructing a system that functioned analogously to the way in which humans reason. The “unsharp” boundaries of fuzzy sets left room for uncertainty, a quality that Zadeh was unable to integrate satisfactorily into earlier system models. In a 1968 paper, he integrated probabilistic measures of events into his fuzzy approach. Working with Richard Bellman in 1970, he developed a theoretical underpinning for his study of fuzzy sets.

As his understanding of fuzzy sets deepened, Zadeh discovered applications that he had not anticipated. His original idea was to develop systems that mimicked human reasoning. In a 1972 paper, he proposed using fuzziness to improve control of electromechanical systems. Within a year, he broadly expanded the possibilities of fuzzy control through the introduction of linguistic variables, common words

such as *hot* or *tepid* or *chilly*, which embraced an imprecisely defined range of quantitative values. Linguistic variables simplified control of electromechanical tasks because they made it easier to describe the way in which the task was to be performed. By 1976, Danish engineers were using fuzzy control to govern the operation of cement kilns, which came to dominate that industry.

A number of other industries, including consumer electronics, automobiles, and home appliances, such as washing machines and air conditioners, adopted fuzzy control in the years that followed. Almost all of these applications were made outside of the United States, mostly in Japan. Skepticism on the part of U.S. industry about the applicability of fuzzy control systems mirrored the skepticism among U.S. academic researchers about the theory. From its inception, fuzzy logic (as the field came to be called) was controversial. Many established researchers in systems analysis denigrated the idea. Without wide community support, the Foundation was hard pressed to support additional research in fuzzy logic. An applicant proposing research in the area had a low chance of finding sympathetic reviewers. If the application came from an established researcher, then reputation alone might carry the proposal. But young, inexperienced researchers were not so fortunate. For example, E. Rustini, a well-respected researcher at the Stanford Research Institute who eventually made solid contributions to fuzzy logic, was unable to win Foundation support at the beginning of his career.

Zadeh, on the other hand, continued to enjoy generous support. He received some support from ARO (Army Research Office) and NASA, but his primary source of funding was always the Foundation (principally the Engineering Directorate, but also Information Science and CISE after it was created). During the 1980s, Zadeh prolifically published research results on fuzzy logic. He concentrated on adapting fuzzy logic principles to expert systems and natural language processing. In both of these efforts, he took advantage of the flexibility of fuzzy systems to bestow common-sense knowledge on them. To do so, he created techniques for expert systems to represent imprecise information and combine evidence. In 1986, he proposed a way of representing meaning in natural language systems. With advances of these kinds, he led a steadily growing community of researchers investigating the promise of fuzzy logic. The Japanese practice of incorporating fuzzy logical control into consumer products made academic and industrial computer engineers in the United States reexamine their attitudes toward fuzzy logic.

Critics sometimes ask why the federal government should be supporting theoretical investigations such as those in computer engineering theory. Implicit in

this criticism is the view that these studies do not have relevance to practical national needs. Martin Hellman's Foundation-supported research on applications of information theory to cryptology is an example of how theoretical research can have important practical implications.

In 1948, when Claude Shannon published his seminal paper on information theory, there was a close association between information theory and cryptography because of wartime concern with transmitting secure messages. Work on data compression and on error correction and detection continued by information theorists after the war, but public work on cryptography waned. Most cryptographic research was classified and carried out by government agencies, such as the National Security Agency. Only in the late 1960s did interest in cryptography appear outside these federal organizations.

One pioneer in the newly opening field was Martin Hellman. Through his contributions, he legitimated cryptography as an academic research subject. After obtaining his Ph.D. from Stanford in 1968, he spent a year at IBM's research center in Yorktown Heights, New York, just as encryption research was commencing there. He was not an active participant, but he was aware of the effort and perceived its importance. The next year, Hellman left for MIT, where he discussed information theory with Peter Elias, a former colleague of Shannon. Elias showed Hellman a paper Shannon had written in 1945 and published in 1949, which made clear to Hellman several subtle connections between cryptography and information theory. Hellman began to suspect that Shannon's information theory was motivated by his work on cryptography. When he returned to Stanford in 1971, Hellman began to muse about cryptography.

Hellman worked on several different areas of information theory. With support from the Foundation, he examined error-detecting codes, data compression, and multiaccess channels (*aloha channels* that were predecessors to the now-popular Ethernet). This work stimulated ideas about encryption, which he considered in his spare time. He soon had enough material to publish a small paper on the subject and, pleased with his results, began to think about working on cryptography full time. Many of his colleagues advised him against this. At the time, the National Security Agency (NSA) had a virtual monopoly on cryptography research. A researcher working outside the NSA could reasonably expect to inadvertently duplicate NSA research or have any new results immediately appropriated and classified by the NSA. In either case, the prospects of making an impact in this field did not appear promising. Undaunted, Hellman approached the Foundation and described to them the sort of work he wanted to do. The agency allowed him to change research topics mid-grant. He turned to cryptography full-time and in 1976 published with

Whitfield Diffie, one of his students, a landmark paper entitled “New Directions in Cryptography.”

In this paper, Hellman and Diffie suggested an alternative to the basic system by which all coded messages were then encrypted. A message to be coded was traditionally operated on by a key cipher, which would transform the elements of the message in some way. The recipient of the coded message would then use the same key to unscramble the message into its original form. Only someone with the correct key could decipher the message, but anybody who knew the key could do so. The Digital Encryption Standard, created by IBM and accepted as the standard for electronic information encryption by the National Bureau of Standards in 1976, was based on this principle. The system had two important shortcomings. One was an issue of communication. It was necessary to distribute copies of the key to everyone who should have access to the coded messages, while at the same time denying access to people for whom the message should remain secret. The military had traditionally communicated keys by courier, but this was infeasible for large communities of people communicating on electronic networks. The other problem was authentication. It was always possible that some interloper with the key could intercept a coded message, introduce his own message into the original, and send the tampered version along to the intended recipient. The conventional systems offered no way to adjudicate disputes between the sender and receiver about the contents of a message. Hellman’s paper proposed an ingenious encryption scheme that solved both of these problems.

Hellman’s system used two different keys on the message: one to encode, the other to decode. The two keys were inverses of each other; that is, they had the opposite effect on the message. They were chosen so that it would be practically impossible to derive one key by knowing the other. One of the keys could be made public, allowing anyone to encrypt a message in such a way that it was only decipherable to the one person who had the partner key, which would be kept private. Because the keys were not inverses of themselves, knowledge of the public key alone would be of no use in decoding. Anyone desiring to send a message to someone else could simply look up the recipient’s public key in a directory, encode the message using the public key, and feel confident that only the recipient, with the private key, could decipher it. The plan eliminated the complications of distributing shared keys. It also guaranteed the authenticity of a message. If a person encrypted something with his or her own private key, no one, even if they could decode the message using the public deciphering key, would be able to add to or change the message in the code in which it had originally been ciphered.

The cipher system seemed foolproof. The question of the security of a public key system was of paramount importance, and the call to investigate it generated considerable research in related disciplines. A system for creating pairs of keys had to be invented. To guarantee security, this system had to have the property of easy confirmation of a public key and great difficulty of deriving the private key from the public one. Computer theorists working on computational complexity had been studying problems with exactly this property when exploring problems that required too much computer time to be practically solved even though their answers could easily be checked for correctness. Prior to Hellman's paper, the practical value of computational complexity was unknown. The public key system offered a problem of practical significance, which invigorated the field and gave it a new focus. It was essential to prove that the calculation of a private key from a public one belonged to that class of intractable problems, if one were to have any confidence in the coding system's security.

Hellman himself contributed to some of this work, not wholly by choice. His research agenda was circumscribed by the authority of the agencies that were funding him. Because the National Security Agency had control of all military spending in cryptography, money that Hellman received from the Air Force Office of Scientific Research initially could not go to cryptography research of any sort. Only later did Hellman spend Air Force money on mathematics research related to cryptography. At the Foundation, program director Frederick Weingarten resisted NSA pressure to curtail funding for Hellman's cryptography work. Some problems Hellman could not work on directly at all, regardless of the funding source. He was interested in appraising the security of IBM's Data Encryption Standard. He believed that the NSA had deliberately weakened it during the standards adoption process in order to ensure that it could break any message encrypted with this standard. The Foundation could not fund work related to a standard, however, and Hellman's funding from the military could not be used for cryptographic research. He could only approach the question obliquely, doing general research supported by the Foundation on the security of codes. He looked for short-cuts to exhaustive, trial-and-error methods to break codes and came up with cryptanalytic time-memory-processor tradeoffs that cut the resource cost of code-breaking.

There was some delay between the proposal of public key cryptosystems and their adoption by commercial interests. A mathematical function that could implement one was first proposed by Ronald Rivest, Adi Shamir, and Leonard Adleman of MIT in 1978. Their research, also supported by the Foundation, resulted in the creation of the RSA cipher, a public key cryptosystem based on modular exponentiation. That same year, Hellman and his student Ralph Merkle proposed a different

public key cryptosystem based on the so-called “knapsack” problem, a well-known problem considered computationally too complex to solve.

Despite these successes, business clients were slow to employ public key systems. This changed over time, however, as the explosive growth of electronic communications increased the need for privacy and authenticity. Banks, in particular, soon accepted public key cryptosystems as a means to conduct electronic funds transfers securely. Widespread implementation of public key cryptosystems in the mid-1980s coincided with the end of Foundation support for Hellman’s work. The Foundation had thus supported basic theoretical research and initial work on an important application of it. Once the basic research was complete, the Foundation bowed out and left development and implementation to others.

7.6 Woodrow Bledsoe, Bertram Raphael, and Artificial Intelligence Research¹⁶

Artificial intelligence (AI) has historically encompassed several different types of research. Some researchers focused on systems that recognized sounds or images; other researchers developed systems that solved problems and learned, exhibited linguistic capabilities, or had various combinations of these and other “human” abilities. The challenges associated with building these systems have been enormous. AI researchers drew heavily on, and stimulated considerable work in, more fundamental areas of computer science, such as computer theory, architecture, software, and graphics. Some of their approaches were hardware-intensive, calling for the construction of expensive experimental equipment; others were more theoretical, requiring only time on commercially available general-purpose computers.

While the military, particularly DARPA, lavished multi-million-dollar grants on laboratories working on artificial intelligence with state-of-the-art machinery, the Foundation tended to fund smaller-scale, theoretical work. Even this work demanded significant computing resources. The Foundation distributed nearly \$19,000,000 for AI research through more than 300 grants between 1963 and 1980; the average Foundation grant provided nearly \$63,000. This represented the fourth highest level of average monthly funding for research (behind architecture, robotics, and graphics) among the twelve categories of computer science research.

NSF funding for artificial intelligence may be divided into five categories. The largest was pattern recognition. Graphics, which included research in scene analysis and image processing, was supported to a comparable degree. These two areas together constituted approximately half of all funding for artificial intelligence. They drew particularly strong support in the mid-1970s, but this interest faded




Figure 7.8 NSF funding for research in artificial intelligence, 1955–1980.

after just a few years; and by 1980, a rough parity of funding had been achieved among all five areas of artificial intelligence. The rest of the Foundation support for AI was distributed fairly evenly among artificial intelligence processes (including automatic theorem proving, automatic problem solving, automatic programming, and learning systems), languages (including research on knowledge representation and natural language processing), and miscellaneous topics.

During the 1960s, the Foundation invested more money into artificial intelligence than almost any other research area. Although the absolute number of dollars granted in this period was fairly small (only 7% of total funding for artificial intelligence was granted before 1971), it was still enough to boost the field at a critical early age. Support went principally to languages and processes. During the 1970s, newer areas of research, such as databases and software, tended to compete more strongly with artificial intelligence for Foundation funding.

Automatic theorem proving, the use of computer systems to prove mathematical theorems, was an important early area of artificial intelligence research. Interest in this topic was not stimulated by indolence on the part of mathematicians. The

achievements of the theorem provers were far too modest for that. More likely, it was “[a feeling that] understanding these problem solving processes is an important step toward the programming of more complex and general problem-solving processes for a variety of intellectual tasks.”¹⁷ The Foundation strongly supported this enterprise.

Donald W. Loveland, the developer of one of the first important automatic theorem provers, drew a useful distinction between two different design philosophies: the logic approach and the human simulation approach. The basic idea is that the human simulation approach deliberately modeled human problem-solving techniques to prove theorems, while the logic approach used whatever technique the computer could effectively exploit.

Most influential early automatic theorem proving systems were of the human simulation type. In 1959, Harvard mathematician Hao Wang disrupted that approach’s hegemony, however, and research efforts shifted to the logic approach. Wang wrote an algorithm that improved the efficiency with which the computer manipulated logical statements while developing a proof. Work by a philosopher, John Alan Robinson, reinforced the trend to logic type provers by dramatically reducing the number of missteps the computer might take in probing for a correct proof. He published the *resolution principle* in 1965.

By the late 1960s, almost all efforts at automatic theorem proving involved the resolution principle. Provers based on this principle were successful, but somewhat unintuitive. The principle was generally employed through a “backward chaining” method. A prover assumed the negation of the theorem to be proved was true, and then looked for contradictions between the negated statement and a set of axioms that the operator supplied to the automatic theorem prover. Because no true statement in a formal system may contradict the system’s axioms, the discovery of any contradiction disproved the negated statement and hence proved the original statement. This was the most effective way to apply resolution. The human simulation approach was all but abandoned in the late 1960s.

Woodrow Bledsoe, a mathematician at the University of Texas at Austin, sought to reverse this trend. He first came to study artificial intelligence through an interest in pattern recognition, but his research focus shifted in the late 1960s to automatic theorem proving. He was not confident that resolution type provers could be capable of proving difficult theorems. He wanted to incorporate the instincts of the mathematician into automatic theorem proving. With Foundation support, he led a lively campaign to steer the field in that direction.

Although eliminating resolution was foremost on Bledsoe’s agenda, he was not blind to its value. His first effort at a theorem prover employed the resolution

principle, albeit begrudgingly. Bledsoe introduced a set of heuristics for reducing the scope of the problem, making it easier for the resolution solvers to tackle it. Heuristics were controversial tools because their use limited the generality of an automatic theorem prover. Theorems of a particular type required one particular set of heuristics. Heuristics thus tended to make the theorem prover a special-purpose device, opposing the trend toward general-purpose computing. Although the use of heuristics in automatic theorem proving was by no means unknown, Bledsoe's work departed from mainstream research in this area.

In 1972, Bledsoe published a paper describing IMPLY, his first prover that completely avoided the resolution principle. IMPLY preserved the original structure of the theorem being proved. It applied a set of rules to the theorem in question and broke the theorem's propositions into component propositions, which it would either recognize as valid or re-divide into still smaller propositions. This straightforward approach, called the *implication method* by Bledsoe, was congenial to mathematicians. Although IMPLY could execute proofs by itself, Bledsoe wrote it with the expectation that a mathematician would work with it interactively in constructing a proof. The similarity of IMPLY's tactics to those of mathematicians was its greatest asset. Bledsoe prepared a set of limit-theorem heuristics for IMPLY, and with them IMPLY proved some important limit-theorems from calculus.

The National Institute of Health (NIH) funded Bledsoe's initial work on IMPLY in the early 1970s, as it had his earlier AI research on biological systems and artificial evolution. When his research interest switched to automatic theorem proving, however, his work lost its relevance to NIH. Regarding the Foundation as a more appropriate funder, NIH forwarded his most recent grant proposal to the Foundation in 1971. Program manager Val Tareski at the Foundation agreed that Bledsoe's grant was a candidate for the "dropout" program, a mechanism to accept grants dropped from other federal agencies in response to the Mansfield Amendment. Based on strong peer reviews, the Foundation awarded Bledsoe a grant in 1972. He continued work on automatic theorem proving, using IMPLY, with Foundation support and disseminated his results in a 1972 article with his students, Robert Boyer and W. H. Henneman.

Bledsoe was not satisfied with the early limit-proving capabilities of IMPLY. The underlying mathematical tool (first-order predicate calculus) limited the sorts of proofs that IMPLY could attempt. For example, he was unable to apply IMPLY to theorems from general topology. After 1973, he started work on a different prover, using more powerful higher-order logics, better suited to topological reasoning. He was not alone in recognizing the added power higher-order logic could bring to his prover; several efforts were made at exploiting it for automatic theorem proving.

Bledsoe was distinguished from the others, however, by his dedication to direct proof. His new prover, named Set-Var, was a non-resolution prover in the style of IMPLY.

Writing a prover that worked in higher-order logic proved to be difficult, and he settled for an enriched first-order logic as the basis for Set-Var. He used Set-Var to prove some theorems in topology and in real analysis, such as the intermediate value theorem, which he had been unable to prove with IMPLY. Set-Var faltered, though, over inequalities. They were a stumbling block that also had impaired IMPLY's performance. In 1978, Bledsoe began work on a solver more adept at handling inequalities.

The principal problem caused by inequalities is the explosion in the number of logical alternatives they create. Faced with an inflated search space, the existing provers had difficulties finding their way to a complete proof. Bledsoe's strategy to handle this problem was to create a solver that utilized the logical operations *shielding term removal* (STR) and *variable elimination* (VE). Employing these operations, logical expressions could be simplified and analysis made easier. Bledsoe's operating procedure was to be extremely selective in generating expressions in order to ensure that the solver did not become overwhelmed by choices. Using a so-called "large-step" strategy, along with heuristics, Bledsoe created a new prover called STR+VE (pronounced "strive").

Bledsoe's work on STR+VE launched a fruitful period in his research. In 1980, he wrote an enhanced version of STR+VE that coupled resolution techniques with variable elimination. Meanwhile, he combined key concepts from Set-Var and STR+VE to create a new solver capable of proving difficult topological theorems, such as the paracompactness theorem and the normality of compact Hausdorff spaces. By 1990, he had prepared a resolution version of Set-Var, as well.

The conversion of Bledsoe's solvers to resolution type did not weaken his resolve that solvers should proceed as mathematicians do. Thus, he turned to solvers that operate by analogy. Taking guidance from his own theorem-proving experience as a mathematician, he identified, as an indispensable element of the proof process, the capacity to recognize familiar situations and try methods and approaches that were effective on similar problems. The analogy solver he envisioned had a database of completed proofs. When attempting a new proof, the solver scanned the database for other proofs on which it could pattern its effort.

Other researchers had begun work along these lines, but their best results, such as a proof of the commutativity of multiplication, could already be proved easily by non-analogy methods. Bledsoe believed that the full potential of analogy proofs had not yet been demonstrated. Working with Larry Hines, he developed an analogy

prover that could prove that the limit of a series of products is equal to the product of the limits when supplied with the proof of the theorem that the limit of a series of sums is equal to the sum of each of the individual limits. This proof pleased Bledsoe because this result had been quite difficult for all previous types of general-purpose automatic theorem provers to achieve.

Bledsoe and his students continued this research in several directions. They prepared a system to simplify proofs mechanically so as to make them clearer and more concise. They returned to the IMPLY system, outfitting it with variable elimination capabilities. They worked with the resolution version of Set-Var and the systems created by the hybridization of STR+VE and Set-Var. They also continued to pursue theorem proving by analogy. All of this work was supported by the Foundation.

When automatic theorem provers receive a proposition, they compare it against a database of things they know to be true (i.e., based on their axioms), and then determine the truth value of the proposition. This in itself is a challenging assignment for a computer; but for some applications, it is not sufficient. Computers in robots and expert systems need not only to decide whether a particular proposition is true or false, but also to find a correction to any false proposition. Only by producing correct information can they make the decisions necessary to complete their complex tasks. Such systems utilize answer-retrieval techniques, a subject of artificial intelligence closely allied with automatic theorem proving. The Foundation was instrumental in supporting work in this area. While the military assumed much of the burden of purchasing the expensive hardware needed to research robot systems, the Foundation established its niche as a funder of their decision-making computerized brains.

Bertram Raphael of the Stanford Research Institute (SRI) was a leading researcher in this area, supported by the Foundation. His 1964 dissertation at MIT on automatic theorem proving provided procedures for using computers to automatically evaluate the truthfulness of propositions by applying techniques from classical logic and predicate calculus to sets of axioms represented symbolically in the computer. This work built directly on the pioneering research on problem solving, such as the General Problem Solver developed at Carnegie Mellon during the 1960s.

In the mid-1960s, Raphael continued his research on theorem proving as a faculty member at UCLA. Charles Rosen, head of the learning machine/pattern recognition project at SRI, attended lectures that Raphael presented on the pro-

programming language LISP. Rosen saw potential applications for LISP to his robotics research, especially for expanding robot capabilities to make decisions based on sensory input, and recruited Raphael to work at SRI.

The pattern recognition group at SRI was working on hardware development for picture processing. By 1966, they simulated an image processor on a digital computer. This was a breakthrough because it allowed them to simulate various designs on a computer rather than build a new piece of hardware to experiment with each new design. The computer simulation proved to be more reliable as well as less expensive and less time consuming. This experience led the SRI group to adopt a new approach to robotics, one that emphasized the use of general-purpose computers. They believed they could avoid some of the existing difficulties by avoiding highly specialized components in favor of an exceptionally well-integrated performance of the many standard parts of the robot system. Instead of concentrating on developing a top-quality camera or a powerful computer, they sought to develop an especially strong underlying logical structure. The robot that resulted from this work, Shakey, was a landmark achievement in robotics.

Work on the Shakey project began in 1967, and the group completed the first version in 1969. The robot was originally limited in the complexity of instructions it could carry out by the XDS-940 computer that controlled it. Shakey accepted instructions in conversational English and translated them into statements of first-order predicate calculus. It executed commands to move to a particular location or push some number of objects into a specified configuration. It could not monitor its own progress while executing a task, however, so it was unable to make adjustments when it started to go astray. A second version of the robot, completed in 1971, featured many performance improvements. It received much of its computational power from a Digital Equipment PDP-10 computer that communicated by radio with a small on-board PDP-11 computer.

Raphael progressed from leader of the Shakey software team, to project head in 1971, to manager of the entire artificial intelligence laboratory in 1972. He managed and participated in the development of Shakey's problem-solving system, known as STRIPS (Stanford Research Institute Problem Solver). Axioms in predicate calculus represented Shakey's physical surroundings. Tasks, entered as goal statements, were represented as propositions. To perform an instruction, Shakey would attempt to prove the goal statement proposition from the axioms describing its environment. If the proof was successful, the task was accomplished. For instance, if the goal statement was that a red block should be on top of a table, and Shakey could prove from its sensory input that the red block was already on the table, then it had nothing to do to accomplish its job. In the more likely event that the goal was not

implied by the present conditions, STRIPS identified the difference between the goal state and the initial state and determined the changes necessary to reduce the difference between the current state and the goal. The program translated these changes into concrete actions that lower-level routines in Shakey's system could execute.

In developing STRIPS, Raphael applied his expertise in theorem proving. He proposed as goals theorems that were generally *not* true, and STRIPS would reason what axioms needed to be changed to make the theorem true. STRIPS's ability to generate answers to the question "what needs to be different?" was the key to its operation.

SRI research in artificial intelligence both supported and went beyond the Shakey project. Automatic theorem proving research led to the development of STRIPS, but it also led to the development of QA3 and QA4, two demonstration question/answer systems, which showed how classical logic and predicate calculus could be used automatically by a computer and applied to everyday problems. This research influenced Robert Kowalski in his work on the logical programming language Prolog. The Shakey project also benefited from the significant work being done at SRI on machine perception. Raphael worked on systems that enabled machines to interpret real-world pictures. Based on the work of Larry Roberts at MIT's Lincoln Laboratories, Raphael developed algorithms that analyzed regions, not simply lines and edges. This was useful not only for Shakey, but also for military applications such as aerial photography.

Many different organizations funded SRI. In the late 1960s and early 1970s, DARPA, the Office of Naval Research, and the Foundation supported Raphael's AI group. DoD provided funding for algorithms to recognize particular patterns and build communications-interface hardware. Foundation funding, applied to the work on theorem proving and machine reasoning, as well as to bits of hardware that functioned as custom interfaces between the major components of Shakey, was central in this research coalition.

Interest in continuing the Shakey project eventually waned. Scientists at SRI and other major robot research labs (e.g., Hitachi, MIT, Stanford, and Edinburgh) believed that the robot systems then under development had reached the limits of the technology then employed. These researchers redirected their work toward developing new hardware and software to incorporate into robotic systems. Raphael refocused SRI's robot effort on industrial-purpose, assembly-line robots. This research was less concerned with artificial intelligence. Since the project had commercial objectives, Raphael turned to industry for support. Corporations such as the Ford Motor Company largely replaced the government as sponsors.

The Foundation nevertheless continued to support those parts of the research program organized by Raphael and his collaborators Cordell Green, Marty Teitlebaum, Peter Hart, and Charles Rosen, that favored theory over application. Raphael eventually split SRI's man-machine systems group, which conducted experiments in *groupware*—enabling groups of people to share a common database of information and collaborate—in two: one group for basic research and another for applications. Foundation support stayed with the basic research group. Thus, continuously from 1967 to 1980 the Foundation supported one of the leading robotics research centers in the world. The work done at this center greatly advanced the theory and practice of robotics and eventually had payoff for American industry.

7.7 Donald Greenberg, Charles Csuri, and Computer Graphics Research¹⁸

The Foundation played an influential role early in the development of computer graphics by providing substantial grants to a small number of promising research projects. Almost no researchers worked on computer graphics prior to the late 1960s. Only the University of Utah, supported with DARPA funding, distinguished itself in this area. The Foundation moved promptly to initiate work, making its first grants in 1967, when the OCA was organized.

The Foundation's funding for graphics may be divided into four categories: hardware for graphics systems, research on algorithms for image production, development of applications for graphics, and miscellaneous other topics. Of these, imaging work garnered the most Foundation funds, with applications trailing slightly behind. Graphics systems and miscellaneous other projects each received significantly less support.

The Foundation funded graphics research by supporting a small number of researchers with moderately large grants. Between 1967 and 1980, only 167 grants were issued for graphics research. These grants, however, totaling almost \$11,000,000, averaged over \$50,000. Only grants to study computer architecture were higher. The largest grants were awarded to a small group of researchers, most of whom received more than one Foundation grant. Although 104 different researchers were named principal investigator on graphics grants, only 31 of them received more than a total of \$100,000 of Foundation funds by 1980. Two researchers, Donald Greenberg of Cornell and Charles Csuri at Ohio State University, together accounted for \$1,970,000, almost 20% of all graphics research funds granted by the Foundation.




Figure 7.9 Annual funding for graphics research, 1955–1980.

Greenberg received ongoing Foundation support for graphics research. With the Foundation's seed money, he founded a graphics research center at Cornell University to enhance the computer's role as a design tool. This was to be accomplished by improving human-machine interactivity through better graphical input and output techniques. His work was successful enough to attract industrial sponsorship. Hewlett-Packard and Digital Equipment Corporation, for example, believed that Greenberg's laboratory could contribute to the development of advanced computer-aided design systems, so they invested millions of dollars in equipment donations. These resources lifted the Cornell lab to preeminence in graphics research. Because of the center's success, the Foundation asked Greenberg to co-author the grant proposal in 1985 that brought one of five national supercomputing facilities to Cornell.

Cornell graphics grew around Greenberg's own effort to develop an image-synthesis system to assist architects. An architect himself, Greenberg began his work with graphics in 1965 when he went to Cornell to teach structural engineering. At that time there was no graphics equipment at Cornell; instead, he traveled to a General Electric Laboratory in Syracuse, where scientists were developing

systems for NASA's Apollo program. There, he produced results impressive enough to convince the Foundation in 1974 to support a computer graphics program at Cornell.

Working with faculty and student collaborators, Greenberg investigated fundamental problems in image synthesis. The group focused on several critical problems of image synthesis: the calculation of hidden surfaces and shadows and tones produced by the illumination. This work on graphical display, although challenging, was only a means to an end. Greenberg's objective was to produce systems that utilized graphics as a tool for simplifying design. He worked on systems to generate images for a variety of fields, including medical applications and water resource planning. Most of his effort was devoted to architecture. He developed models of how foliage and other shadows affect the escape of heat energy from buildings, for example. With these models, he created interactive graphics systems to assist architects in considering these factors in their designs.

This work on energy simulation brought Greenberg in touch with the literature on graphics for thermal engineering. Considering the laws of radiant heat energy, Greenberg realized that light, another form of radiant energy, obeyed the same laws. He applied this fact to develop a new way of mathematically describing the illumination of scenes. He calculated the illumination at any point by summing the light reflected to it from all other points. This approach effectively simulated the way light reflects diffusely from unpolished surfaces. A typical scene involved thousands of points. Calculating all of the equations consumed considerable computer time, but once it was done, the illumination value for each point was determined absolutely; it never required recalculation (as long as every element in the scene remained stationary). Greenberg presented this technique, which he called *radiosity*, in 1984.

Radiosity provided an alternative to ray-tracing, the other significant approach to general illumination in computer graphics. Ray-tracing was an inversion of the way illumination works in the real world. The computer mathematically tracks the path of each ray of light in a scene from the image in the observer's eye back to its source, taking note of the reflections made along the way. This method produced realistic images of polished surfaces, but it was less effective when dealing with diffuse light sources and reflections off dull surfaces. The numerous calculations demanded in ray-tracing were dependent on the location of the observer's vantage point. To display the scene from a different angle required running the entire ray-tracing procedure over again.

Radiosity and ray-tracing each had strengths and weaknesses. Ray-tracing brought exceptional realism in some circumstances, but radiosity reproduced subtle effects such as the bleeding of color from one surface to another and penumbras

along shadow boundaries. Because of his interest in architecture, Greenberg had special reason to favor radiosity. Very few additional calculations were needed to alter the perspective of scenes rendered with radiosity, and he could make quick transitions between different views of a radiosity scene. This simplified the simulation of movement through the space modeled by the computer, giving architects the power to envision moving through an as-yet unbuilt design. Greenberg's system was an enormous improvement over working with successions of static pictures or small-scale models.

Persistent work with radiosity paid off handsomely in improvements in realism. In 1985, Greenberg added the illumination contributions of hidden surfaces to his radiosity calculations. The following year, he proposed a radiosity method for non-diffuse surfaces. These additions enhanced the quality of the radiosity technique, but the method still failed to match ray-tracing for realism in certain situations. Eager to achieve true photographic realism, Greenberg began work on a synthesis of the two approaches. He created a system that used radiosity analysis to plot the general illumination of a scene and used ray-tracing routines to refine the image. In 1988, he improved his radiosity algorithm so that, in addition to consuming less memory, it was able to gracefully refine its own images. With this ability, he was able to enhance the interactivity of his system by removing the view-dependent ray-tracing routines.

Work on radiosity and other research on display fundamentals increased the value of computer graphics for scientists in many disciplines. Along with Greenberg's structural engineering applications, such as graphically simulating the effects of earthquakes on building structures, other Cornell researchers used his graphics equipment to better visualize their experimental results. Working with colleagues in other departments, Greenberg conducted research in biotechnology, molecular modeling, fluid flow, and general systems for visualization of scientific data.

The Foundation wholeheartedly supported this work. Seeking to address the rapid advances in graphics during the 1980s, the Foundation considered the effectiveness of its own support for graphics. A panel on graphics, image processing, and workstations, sponsored by the Foundation's Division of Advanced Scientific Computing, decided in 1986 that:

computer graphics and image processing are within computer science; the application of computers to the discipline sciences is called computational science. Applying graphics and imaging techniques to computational science is a whole new area of endeavor, which Panel members termed Visualization in Scientific Computing.

A Foundation-sponsored workshop in February 1987 concluded that allocations for scientific visualization were insufficient. A new initiative was needed “to get visualization tools into ‘the hands and minds’ of scientists.”

Thomas A. DeFanti, of Ohio State University, served as one of the editors of the Foundation’s report on the 1987 workshop. He had extensive experience with computer graphics at the Ohio Supercomputing Center. For example, he had produced numerous videotapes that explained how computer animation might be applied to engineering, physics, chemistry, geography, architecture, and medicine. Research in his laboratory led in 1987 to a sophisticated computer environment for graphical visualization of scientific data called apE. It was distributed to over 700 sites worldwide by late 1990.

That this sort of work should come out of Ohio State University is at the same time expected and surprising. In the 1970s, Ohio State computer scientists, like those at Cornell, developed a research program in graphics by means of steady Foundation support. The leader was Charles Csuri, an art professor with an interest in computers as an artistic medium. In a groundbreaking paper in 1974, Csuri set forth a conceptual framework to aid evaluation of the interactive art objects computers could generate. His research began with aesthetic issues, a perspective far removed from the desire for a graphical computer-aided design system that motivated Greenberg. Certain common issues, however, such as the need for mathematical descriptions of forms and the question of image display, brought both of their programs to study scientific visualization.

When Csuri first applied to the Foundation for funding in 1968, the program officers looked at his proposal with curiosity. They did not expect that Csuri would make a significant artistic contribution but did recognize that he was stimulating an interest in mathematics and computers among his students. His activities brought life to computer graphics at Ohio State; colleagues in the computer science department began experimenting with imaging systems. Although he concentrated on images that were abstract, others at the university took an interest in more traditional forms, particularly the human body. A face-drawing system developed by computer scientists Mark Gillenson and B. Chandrasekaran in 1973 was the beginning of a long-term effort at Ohio State to employ knowledge of physiology to the production and animation of human images.

The Foundation made decisive contributions to that research area through its support of the Ohio State group. David Zeltzer and Michael Girard were two researchers who benefited from this support. Zeltzer, a student working under

Chandrasekaran with the Ohio State Computer Graphics Research Group between 1978 and 1984, modeled the movements of skeletons in order to render human motions more accurately. He produced computer-animated sequences of complete skeletons walking over both level and uneven terrain. Girard, a student at Ohio State during the mid-1980s, conducted similar research exploring the mechanics of joints to animate walking figures with increased realism. Their work offered several promises. The success of the mathematical manipulations used in animating limbs suggested possibilities for computer control of robot arms and manipulators, while the techniques of three-dimensional animation were more widely applicable to all imaging applications. The work of the Computer Graphics Research Group was of fundamental importance in developing methods for the display of complex three-dimensional data, which was of increasing importance in almost every science.

7.8 Conclusions

As early as the 1950s, the Foundation was supporting research in computer science, especially on databases and information retrieval in connection with its active science information program. Especially after the creation of the OCA in 1967, the Foundation supported research across many areas of computer science: computer theory, circuits, computer architecture, software, numerical analysis, computer engineering theory, artificial intelligence, and computer graphics. Many of the leading computer scientists in the United States were supported by the Foundation in carrying out this research. Foundation support also had a benefit, not seen in the funding by the federal mission-oriented agencies such as DARPA and the DoE, of helping to broaden the academic base of computing in the United States. One program in particular, CER, helped a number of universities to broaden to include experimental computer science research, along with theoretical computer science, in their portfolios.

Notes

1. David M. Young, Jr. April 10, 1991. Private communication to Andrew Goldstein.
2. Donald P. Greenberg. January 30, 1991. Private communication to Andrew Goldstein.
3. Woodrow W. Bledsoe. June 13, 1991. Private communication to Andrew Goldstein.
4. Charles A. Csuri. March 4, 1991. Private communication to Andrew Goldstein.
5. These numbers are based on examining all of the computing-based research grants in the period. If one looks at only the Computer Research division and sections, as Chapter 1 does in this volume, one gets significantly different numbers.

6. The average Foundation grant size was remarkably consistent over time. In the early 1960s, investigators were given close to \$40,000 for their work. This figure increased to \$60,000 by 1980. This suggests, after adjusting for inflation, some decline in the real value of Foundation research grants over time. It is important to remember, however, that computing power, purchased either as hardware or as time leased on someone else's computer, grew cheaper year after year, in opposition to the trend of prices of items that constitute the consumer price index. Because computing power represented a significant cost in research budgets, standard inflation measures are an unreliable adjustment for the purchasing power of the grants.
7. One important area of computer science research—one that NSF was particularly well known for—was theoretical computer science. We do not discuss NSF funding for theoretical computer science in this chapter, but it is covered in: W. Aspray, B. O. Williams, and A. Goldstein. 1996. The social and intellectual shaping of a new mathematical discipline: The role of the National Science Foundation in the rise of theoretical computer science and engineering. In R. Calinger, ed. *Vita Mathematica: Historical Research and Integration with Teaching*. Mathematical Association of America Notes Series.
8. Case studies and statistical analysis of awards are not all that well suited to answering what the overall impact of Foundation funding was on research. In the Aspray, Williams, and Goldstein project, the researchers queried many present and former staff members and computer researchers by mail, telephone, and in person to try to gain a better assessment. As indicated in the remainder of this text, some of them made some general remarks, but it was hard to forge them into a general assessment.
9. This material is taken from Aspray, Williams, and Goldstein, 1992, and was lightly edited in 2018, but there was no effort to bring the story up to date. Over time, we lost the links between footnotes and the places in the text where the footnotes were located. So instead, we simply give all of the references for this case study here in one place:
 1. R. K. Brayton, G. D. Hachtel, and A. Sangiovanni-Vincentelli. 1990. Multilevel logic synthesis. *Proceedings of the IEEE*, 78(2).
 2. M. A. Breuer. October 1966. General survey of design automation of digital computers. *Proceedings of the IEEE*, 54(12): 1708–1709.
 3. J. N. Culliney. February 1975. Program Manual: NOR Network Transduction Based on Connectable and Disconnectable Conditions, in UIUCDCS-R-75-698, Department of Computer Science, University of Illinois, Urbana, IL.
 4. Hellerman, L. 1963. A catalog of three variable OR-Invert and AND-Invert logical circuits. *IEEE Transactions on Electronic Computers*, EC-12(3): 198–223.
 5. Y. Kambayashi and S. Muroga. June 1976. Network transduction based on permissible functions UIUCDCS-R-76-804. Department of Computer Science, University of Illinois, Urbana, IL.
 6. Oral history, Saburo Muroga, interviewed by Andrew Goldstein, July 23, 1991. IEEE History Center. http://ethw.org/Oral-History:Saburo_Muroga.
 7. S. Muroga. 1970. Logical design of optimal digital networks by integer programming. In J. T. Tou, ed. *Advances in Information Science*, vol. 3. New York: Plenum Press.
 8. S. Muroga. 1991. Computer-aided logic synthesis for VLSI chips (draft) to appear in M. C. Yovits, ed. *Advances in Computing*, vol. 31. Academic Press, p. 6.
 9. S. Muroga. 1971. *Threshold Logic and Its Application*. New York: John Wiley and Sons.

10. S. Muroga. "Threshold Logic," lecture notes of EE 497 and EE 498. 1965–1966 school year edition. Digital Computer Lab, University of Illinois, Urbana, IL.
 11. S. Muroga, K. Yahiko, C. L. Hung, and J. N. Culliney. 1989. The transduction method-design of logic networks based on permissible functions. *IEEE Transactions on Computers*, 38(10).
 12. S. Muroga, Y. Kambayashi, H. C. Lai, and J. Niel Culliney. 1989. The transduction method-design of logic networks based on permissible functions. *IEEE Transactions on Computers*, 38(10): 1404.
 13. S. Muroga and H. C. Lai. September 1976. Minimization of logic networks under a generalized cost function. *IEEE Transactions on Computers*, C-25: 893.
 14. L. Trevillyan. An overview of logic synthesis systems. *Proceedings of the 24th ACM/IEEE Design Automation Conference*, 167.
10. This material is taken from Aspray, Williams, and Goldstein, 1992, and was lightly edited in 2018, but there was no effort to bring the story up to date. Over time, we lost the links between footnotes and the places in the text where the footnotes were located. So instead, we simply give all of the references for this case study here in one place:
 1. J. L. Baer and G. Estrin. 1969. Bounds for maximum parallelism in a bilogic graph model of computations. *IEEE Transactions on Computers*, C-18: 1012–1014. DOI: [10.1109/T-C.1969.222572](https://doi.org/10.1109/T-C.1969.222572).
 2. G. Bekey and W. Karplus. 1968. *Hybrid Computing*. New York: John Wiley & Sons. 154.
 3. D. Gajski, D. Kuck, D. Lawrie, and A. Sameh. August 1983. CEDAR—A large scale multiprocessor. *Proceedings of the 1983 International Conference on Parallel Processing* (Bellaire, MI August, 1983), 524–529. DOI: [10.1145/859526.859527](https://doi.org/10.1145/859526.859527).
 4. K. Gallivan, W. Jalby, and U. Meier. November 1987. The use of BLAS3 in linear algebra on a parallel processor with a hierarchical memory system. *SIAM Journal on Scientific and Statistical Computing*, 8(6): 1079–1084.
 5. Oral history, W. Karplus, interviewed by Andrew Goldstein, July 26, 1991. IEEE Center for the History of Electrical Engineering. http://ethw.org/Oral-History:Walter_Karplus.
 6. W. Karplus. October 1955. The use of electronic analog computers with resistance network analogs. *British Journal of Applied Physics*, 6: 356–357. <https://iopscience.iop.org/article/10.1088/0508-3443/6/10/305/meta>.
 7. W. Karplus and O. J. M. Smith. March 1956. The application of electrical transients to well logging. *Journal of Petroleum Technology*, 8: 53–56. DOI: [10.2118/487-G](https://doi.org/10.2118/487-G).
 8. W. Karplus. 1956. Water-coning before breakthrough: An electronic analog treatment. American Institution of Mining Engineers, *Petroleum Transactions*, 207: 240–245. <https://www.onepetro.org/general/SPE-639-G>.
 9. W. Karplus, G. Bekey, and P. Pekrul. November 1958. Atmospheric diffusion of air pollutants. *Industrial and Engineering Chemistry* 50: 1657–1661. DOI: [10.1021/ie50587a037](https://doi.org/10.1021/ie50587a037).
 10. W. Karplus. 1962. The discrete-space-discrete-time computer. *Automatic and Remote Control*, vol. 1. London: Butterworth, 137–140.
 11. W. Karplus and G. Bekey. August 1986. The changing role of analog and hybrid computer systems. *Proceedings of the 1968 Congress, International Federation of Information Processing* (Edinburgh, Scotland). 681–691.
 12. Oral history, D. Kuck, interviewed by Andrew Goldstein, August 15, 1991. IEEE History Center. http://ethw.org/Oral-History:David_Kuck.

13. D. Kuck, Y. Muraoka, and S. C. Chen. December 1972. On the number of operations simultaneously executable in FORTRAN-like programs and their resulting speed-up. *IEEE Transactions on Computers* C-21: 1293–1310. DOI: [10.1109/T-C.1972.223501](https://doi.org/10.1109/T-C.1972.223501).
 14. D. Lawrie. 1975. Access and alignment of data in an array processor. *IEEE Transactions on Computers* C-24: 1145–1155. DOI: [10.1109/T-C.1975.224157](https://doi.org/10.1109/T-C.1975.224157).
 15. W. C. Tam and W. Karplus. December 1973. The use of the digital simulation language PDEL in hydrological studies. *Water Resources Bulletin*, 9: 1100–1111. DOI: [10.1111/j.1752-1688.1973.tb05849.x](https://doi.org/10.1111/j.1752-1688.1973.tb05849.x).
 16. P. C. Yew, D. A. Padua, and D. H. Lawrie. 1982. Stochastic properties of a multiple-layer single-stage shuffle-exchange network in a message switching environment. In J-L. Baer, ed. *Journal of Digital Systems*, 4(4): 387–410.
11. This is not universally true. For example, Westinghouse operated a large analog computer facility until 1991.
 12. This material is taken from Aspray, Goldstein, and Williams, 1992, and was lightly edited in 2018, but there was no effort to bring the story up to date. Over time, we lost the links between footnotes and the places in the text where the footnotes were located. So instead, we simply give all of the references for this case study here in one place:
 1. E. W. Dijkstra. March 1968. Go To statement considered harmful. *Communications of the ACM*, 11: 147–148. <https://homepages.cwi.nl/~storm/teaching/reader/Dijkstra68.pdf>.
 2. C. A. R. Hoare. 1972. Proof of the correctness of data representations. *Acta Informatica*, 1(4): 271–281. DOI: [10.1007/BF00289507](https://doi.org/10.1007/BF00289507).
 3. B. Liskov. 1986. *Abstraction and Specification in Program Development*. Cambridge, MA: MIT Press, p. xiv.
 4. B. Liskov. May 1982. On linguistic support for distributed programs. *IEEE Transactions on Software Engineering* SE-8(3): 203. DOI: [10.1109/TSE.1982.235250](https://doi.org/10.1109/TSE.1982.235250).
 5. B. Liskov. 1988. Distributed programming in Argus. *Communications of the ACM*, 31(3): 300. DOI: [10.1145/42392.42399](https://doi.org/10.1145/42392.42399).
 6. Oral history, B. Liskov, interviewed by William Aspray, August 6, 1991. http://ethw.org/Oral-History:Barbara_Liskov_1991.
 7. B. Liskov, T. Bloom, D. Gifford, R. Scheifler, and W. Weihl. 1988. Communication in the Mercury System. *Proceedings of the 21st Hawaii International Conference on System Science*, pp. 178–187. <https://dl.acm.org/citation.cfm?id=54784>.
 8. B. Liskov, A. Snyder, R. Atkinson, and C. Schafert. August 1977. Abstraction mechanisms in CLU. *Communications of the ACM*, 8(20): 565–566. DOI: [10.1145/359763.359789](https://doi.org/10.1145/359763.359789).
 9. R. L. London. 1981. Remarks on the impact of verification on language design. In M. Shaw, ed. *Alphard: Form and Content*. New York: Springer-Verlag.
 10. P. Naur and B. Randell, eds. 1968. Software engineering. NATO, 1969, *Report on a Conference Sponsored by the NATO Science Committee* (Garmisch, Germany), Oct. 7–11, 1968. <http://homepages.cs.ncl.ac.uk/brian.randell/NATO/nato1968.PDF>.
 11. Oral history, Mary Shaw, interviewed by Andrew Goldstein, July 26, 1991. IEEE Center for the History of Electrical Engineering (not deposited).
 12. M. Shaw, ed. 1981. *Alphard: Form and Content*. New York: Springer-Verlag.
 13. M. Shaw. 1980. The impact of abstraction concerns on modern programming languages. *Proceedings of the IEEE*, 68(9): 1122. DOI: [10.1109/PROC.1980.11811](https://doi.org/10.1109/PROC.1980.11811).

14. N. Wirth. April 1971. Program development by stepwise refinement. *Communications of the ACM*, 14: 221–227. DOI: [10.1145/362575.362577](https://doi.org/10.1145/362575.362577).
15. W. A. Wulf et al., ed. February 1978. “An Informal Definition of Alphard (Preliminary),” Carnegie Mellon University Technical Report. In M. Shaw, ed. *Alphard: Form and Content*. New York: Springer-Verlag, 1981.
16. W. A. Wulf, R. L. London, and M. Shaw. December 1976. Abstraction and verification in Alphard: Introduction to language and methodology. *IEEE Transactions on Software Engineering* SE-2(4): 253–265. DOI: [10.1145/359763.359782](https://doi.org/10.1145/359763.359782).
13. C. A. R. Hoare. 1972. Proof of the correctness of data representations. *Acta Informatica*, 1(4): 271–281. DOI: [10.1007/BF00289507](https://doi.org/10.1007/BF00289507).
14. This material is taken from Aspray, Goldstein, and Williams, 1992, and was lightly edited in 2018, but there was no effort to bring the story up to date. Over time, we lost the links between footnotes and the places in the text where the footnotes were located. So instead, we simply give all of the references for this case study here in one place:
1. W. Aspray. 1990. *John von Neumann and the Origins of Modern Computing*. Cambridge, MA: MIT Press.
 2. J. M. Boyle, W. J. Cody, W. R. Cowell, B. S. Garbow, Y. Ikebe, C. B. Moler, and B. T. Smith. 1972. NATS, A collaborative effort to certify and disseminate mathematical software. *Proceedings 1972 National ACM Conference*, vol. 2, 630–635. DOI: [10.1145/800194.805838](https://doi.org/10.1145/800194.805838).
 3. W. J. Cody. 1984. Observations on the mathematical software effort. In W. Cowell, ed. *Sources and Development of Mathematical Software*. Englewood Cliffs, NJ: Prentice-Hall. 3.
 4. W. J. Cody. 1984. FUNPACK: A package of special function routines. In W. Cowell, ed. *Sources and Development of Mathematical Software*. Englewood Cliffs, NJ: Prentice-Hall, pp. 50–51.
 5. Oral history, Wayne Cowell, interviewed by Andrew Goldstein, July 31, 1991. IEEE History Center. http://ethw.org/Oral-History:Wayne_Cowell.
 6. J. J. Dongarra and C. B. Moler. 1984. EISPACK: A package for solving matrix eigenvalue problems. In W. Cowell, ed. *Sources and Development of Mathematical Software*. Englewood Cliffs, NJ: Prentice-Hall, pp. 68–87.
 7. J. J. Dongarra and G. W. Stewart. 1984. In Wayne Cowell, ed. *Sources and Development of Mathematical Software*, p. 45, *op. cit.*
 8. G. Forsythe. 1966. Algorithms for scientific computation. *Communications of the ACM*, 9: 255–256. DOI: [10.1145/365278.365294](https://doi.org/10.1145/365278.365294).
 9. G. Forsythe and C. B. Moler. 1967. *Computer Solution of Linear Algebra Systems*. Englewood Cliffs, NJ: Prentice-Hall.
 10. R. G. Grimes, D. R. Kincaid, W. I. MacGregor, and D. M. Young. 1978. “ITPACK report: Adaptive iterative algorithms using symmetric sparse storage (Report CNA-139).” Center for Numerical Analysis. DOI: [10.2118/7687-MS](https://doi.org/10.2118/7687-MS).
 11. D. R. Kincaid, J. Respass, D. M. Young, and R. Grimes. ITPack 2C: A FORTRAN package for solving large sparse linear systems by adaptive accelerated iterative methods. *AMC Transactions on Mathematical Software*, 8: 302–322. DOI: [10.1145/356004.356009](https://doi.org/10.1145/356004.356009).
 12. D. R. Kincaid and D. M. Young. 1988. “Report CNA-217.” Center for Numerical Analysis, The University of Texas at Austin.
 13. J. R. Rice and R. F. Boisvert, eds. 1985. *Solving Elliptic Problems Using ELLPACK*. New York: Springer-Verlag.

14. J. H. Wilkinson and C. Reinsch. 1971. *Linear Algebra*, vol. 2. In F. L. Bauer, series ed., "Handbook for Automatic Computation." New York: Springer-Verlag).
 15. D. M. Young. 1950. *Iterative Methods for Solving Partial Differential Equations of Elliptic Type*. Harvard University.
 16. D. M. Young. 1989. A historical overview of iterative methods. *Computer Physics Communications*, 53: 12. DOI: [10.1145/87252.88086](https://doi.org/10.1145/87252.88086).
 17. Oral history, D. Young, interviewed by Andrew Goldstein, July 22, 1991. IEEE History Center. http://ethw.org/Oral-History:David_M._Young_Jr.
15. This material is taken from Aspray, Goldstein, and Williams, 1992, and was lightly edited in 2018, but there was no effort to bring the story up to date. Over time, we lost the links between footnotes and the places in the text where the footnotes were located. So instead, we simply give all of the references for this case study here in one place:
1. Oral history, W. Richards Adrion, interviewed by William Aspray, Amherst, MA, October 29, 1990. Charles Babbage Institute. <https://conservancy.umn.edu/bitstream/handle/11299/104300/oh211wra.pdf?sequence=1&isAllowed=y>.
 2. J. Bamford. 1982. *The Puzzle Palace: A Report on America's Most Secret Agency*. Boston: Houghton Mifflin.
 3. R. E. Bellman and L. Zadeh. 1970. Decision-making in a fuzzy environment. *Management Science*, 17: B-141–B-164. DOI: [10.1287/mnsc.17.4.B141](https://doi.org/10.1287/mnsc.17.4.B141).
 4. Oral history, J. Cherniavsky, interviewed by William Aspray, Washington, DC. September 28, 1991. IEEE History Center (not deposited).
 5. W. Diffie and M. Hellman. November 1976. New directions in cryptography. *IEEE Transactions on Information Theory* IT-22(6): 644–654. DOI: [10.1109/TIT.1976.1055638](https://doi.org/10.1109/TIT.1976.1055638).
 6. W. Diffie and M. Hellman. June 1977. Exhaustive cryptanalysis of the NBS data encryption standard. *Computer*, 10(6): 74–84. DOI: [10.1109/C-M.1977.217750](https://doi.org/10.1109/C-M.1977.217750).
 7. M. E. Hellman. July 1980. A cryptanalytic time-memory trade-off. *IEEE Transactions on Information Theory*, IT-26(4): 401–406. DOI: [10.1109/TIT.1980.1056220](https://doi.org/10.1109/TIT.1980.1056220).
 8. M. Hellman. April 1987. Commercial encryption. *IEEE Network Magazine*, 1(2): 6–10.
 9. Oral history, M. Hellman, interviewed by Andrew Goldstein, July 31, 1991. IEEE History Center. http://ethw.org/Oral-History:Martin_Hellman.
 10. M. Hellman, R. Merkle, L. Schroepfel, W. Washington, W. Diffie, S. Pohlig, and P. Schweitzer. 1976. "Results of an initial attempt to cryptanalyze the NBS data encryption standard." Electrical Engineering Department, Stanford University, Stanford, CA, SEL 76-042, September 9, 1976. https://ee.stanford.edu/~hellman/resources/1976_sel_des_report.pdf.
 11. R. C. Merkle and M. Hellman. September 1978. Hiding information and signatures in trapdoor knapsacks. *IEEE Transactions on Information Theory*, IT-24(5): 525–530. DOI: [10.1109/TIT.1978.1055927](https://doi.org/10.1109/TIT.1978.1055927).
 12. R. Rivest, A. Shamir, and L. Adleman. February 1978. On digital signatures and public key cryptosystems. *Communications of the ACM*, 21: 120–126. DOI: [10.1145/359340.359342](https://doi.org/10.1145/359340.359342).
 13. K. Self. November 1990. Designing with fuzzy logic. *IEEE Spectrum*, 27(11): 42–44, continued on 105. DOI: [10.1109/6.62215](https://doi.org/10.1109/6.62215).
 14. C. Shannon. October 1949. Communication theory of secrecy systems. *Bell System Technical Journal*, 28: 656–715. DOI: [10.1002/j.1538-7305.1949.tb00928.x](https://doi.org/10.1002/j.1538-7305.1949.tb00928.x).
 15. L. Zadeh. June 1965. Fuzzy sets. *Information and Control*, 8(3): 338–353. DOI: [10.2307/2272014](https://doi.org/10.2307/2272014).

16. L. Zadeh. January 1973. Outline of a new approach to the analysis of complex systems and decision processes. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-3(1): 28–44. DOI: [10.1109/TSMC.1973.5408575](https://doi.org/10.1109/TSMC.1973.5408575).
 17. L. Zadeh. 1968. Probability measures of fuzzy events. *Journal of Mathematical Analysis and Applications*, 23: 421–427. DOI: [10.1016/0022-247X\(68\)90078-4](https://doi.org/10.1016/0022-247X(68)90078-4).
 18. L. Zadeh. 1983. A computational approach to fuzzy quantifiers in natural languages. In N. Cercone, ed. *Computational Linguistics*. New York: Pergamon Press.
 19. L. Zadeh. January 1986. Test-score semantics as a basis for a computational approach to the representation of meaning. *Literary and Linguistic Computing*, 1(1): 24–N35. DOI: [10.1093/lc/1.1.24](https://doi.org/10.1093/lc/1.1.24).
 20. Oral history, Lotfi Zadeh, interviewed by Andrew Goldstein, July 24, 1991. IEEE History Center. http://ethw.org/Oral-History:Lotfi_Zadeh.
 21. L. Zadeh and C. A. Desoer. 1963. *Linear System Theory; The State Space Approach*. New York: McGraw-Hill.
 22. L. Zadeh and J. R. Ragazzini. July 1950. An extension of Wiener's Theory of Prediction. *Journal of Applied Physics*, 21: 645–655. DOI: [10.1063/1.1699725](https://doi.org/10.1063/1.1699725).
 23. L. Zadeh and J. R. Ragazzini. October 1952. Optimum filters for the detection of signals in noise. *Proceedings of the IRE*, 40(10): 1223–1231. DOI: [10.1109/JRPROC.1952.274117](https://doi.org/10.1109/JRPROC.1952.274117).
16. This material is taken from Aspray, Goldstein, and Williams, 1992, and was lightly edited in 2018, but there was no effort to bring the story up to date. Over time, we lost the links between footnotes and the places in the text where the footnotes were located. So instead, we simply give all of the references for this case study here in one place:
1. W. W. Bledsoe. 1971. Splitting and reduction heuristics in automatic theorem proving. *Artificial Intelligence*, 2: 55–77. [https://doi.org/10.1016/0004-3702\(71\)90004-X](https://doi.org/10.1016/0004-3702(71)90004-X).
 2. W. W. Bledsoe. 1971. Grant proposal P2J-132, submitted to the National Science Foundation, Office of Computing Activities, Computer Science and Engineering Section, program for Theoretical Science.
 3. Oral history, W. W. Bledsoe, interviewed by Andrew Goldstein, October 21, 1991. IEEE History Center. http://ethw.org/Oral-History:Woodrow_Wilson_Bledsoe.
 4. W. W. Bledsoe, R. S. Boyer, and W. H. Henneman. 1972. Computer proofs of limit theorems. *Artificial Intelligence*, 3: 27–60. DOI: [10.1016/0004-3702\(72\)90041-0](https://doi.org/10.1016/0004-3702(72)90041-0).
 5. E. A. Feigenbaum and Julian Feldman, eds. 1963. *Computers and Thought*. New York: McGraw-Hill.
 6. D. W. Loveland. 1984. Automated theorem-proving: A quarter-century review. *Contemporary Mathematics*, 29: 3. DOI: [10.1090/conm/029](https://doi.org/10.1090/conm/029).
 7. A. L. Norberg and J. E. O'Neill. 2000. *Transforming Computer Technology: Information Processing for the Pentagon, 1962–1986*. Johns Hopkins University Press.
 8. B. Raphael. 1976. *The Thinking Computer: Mind Inside Matter*. San Francisco, CA: W. H. Freeman and Company.
 9. Oral history, Bertram Raphael, interviewed by Andrew Goldstein, July 25, 1991. IEEE History Center. http://ethw.org/Oral-History:Bertram_Raphael.
 10. J. A. Robinson. 1965. A machine-oriented logic based on the resolution principle. *Journal of the Association of Computing Machinery*, vol. 12: 23–41. DOI: [10.1145/321250.321253](https://doi.org/10.1145/321250.321253).
 11. R. J. Solomonoff. 1966. Some recent work in Artificial Intelligence. *Proceedings of the IEEE*, 54(12): 1687–1697. DOI: [10.1109/PROC.1966.5250](https://doi.org/10.1109/PROC.1966.5250).

17. D. W. Loveland. 1984. Automated theorem-proving: A quarter-century review. *Contemporary Mathematics*, 29: 3.
18. This material is taken from Aspray, Goldstein, and Williams, 1992, and was lightly edited in 2018, but there was no effort to bring the story up to date. Over time, we lost the links between footnotes and the places in the text where the footnotes were located. So instead, we simply give all of the references for this case study here in one place:
 1. M. F. Cohen, S. E. Chen, John R. Wallace and D. P. Greenberg. August 1988. A progressive refinement approach to fast radiosity image generation. *SIGGRAPH Proceedings*, 22(4): 75–84. DOI: [10.1145/378456.378487](https://doi.org/10.1145/378456.378487).
 2. M. F. Cohen and D. P. Greenberg. July 1985. The Hemi-Cube: A radiosity solution for complex environments. *SIGGRAPH Proceedings*, 19(3): 31–40. DOI: [10.1145/325334.325171](https://doi.org/10.1145/325334.325171).
 3. C. Csuri. April 1974. Computer graphics and art. *Proceedings of the IEEE*, 62(4): 503–515. DOI: [10.1109/PROC.1974.9455](https://doi.org/10.1109/PROC.1974.9455).
 4. C. Csuri. March 4, 1991. Personal correspondence to William Aspray.
 5. P. N. French, L. E. Johnson, D. P. Loucks, and D. P. Greenberg. 1978. Water resource planning using computer graphics. *ASCE Convention and Exposition*, Preprint 3342. Chicago (October 16–20, 1978).
 6. M. Girard and A. A. Maciejewski. 1985. Computational modeling for the computer animation of legged figures. *SIGGRAPH '85 Proceedings of the 12th Annual Conference on Computer Graphics and Interactive Techniques*, 19(3): 263–269. DOI: [10.1145/325165.325244](https://doi.org/10.1145/325165.325244).
 7. C. Goral, K. E. Torrence, D. P. Greenberg, and B. Battaile. July 1984. Modeling the interaction of light between diffuse surfaces. *SIGGRAPH Proceedings*, 18(3): 213–222. DOI: [10.1145/964965.808601](https://doi.org/10.1145/964965.808601).
 8. D. S. Immel, M. F. Cohen, and D. P. Greenberg. August 1986. A radiosity method for non-diffuse environments. *SIGGRAPH Proceedings*, 20(4): 133–142. DOI: [10.1145/15922.15901](https://doi.org/10.1145/15922.15901).
 9. Glenn R. Ingram. January 16, 1969. Deputy Head of OCA. Memorandum to Fred Weingarten, National Science Foundation.
 10. B. H. McCormick et al. November 1987. Visualization in scientific computing. *Computer Graphics*, 21(6): vii–ix. DOI: [10.1109/MCG.1987.276849](https://doi.org/10.1109/MCG.1987.276849).
 11. J. H. Pittman and D. P. Greenberg. July 1982. An interactive graphics environment for architectural energy simulation. *SIGGRAPH Proceedings*, 16(3): 233–241. DOI: [10.1145/800064.801285](https://doi.org/10.1145/800064.801285).
 12. B. Raphael, 1976, *op. cit.*
 13. R. Rogers, P. Atherton, D. Nall, and D. P. Greenberg. 1978. A means for including shadowing in a building's thermal analysis. *3rd International Conference on Computers in Engineering and Building Design, CAD '78 Conference* (March 14–16, 1978), 97–109.
 14. A. Sunguroff and D. P. Greenberg. August 1978. Computer generated images for medical applications. *SIGGRAPH*, 12(3): 196–202. DOI: [10.1145/965139.807390](https://doi.org/10.1145/965139.807390).
 15. J. R. Wallace, M. F. Cohen, and D. P. Greenberg. July 1987. A two-pass solution to the rendering equation: A synthesis of ray tracing and radiosity methods. *SIGGRAPH*, 21(4): 311–320. DOI: [10.1145/37402.37438](https://doi.org/10.1145/37402.37438).
 16. D. Zeltzer and C. Csuri. June 10–12, 1981. Goal directed movement simulation. *Proceedings of the Canadian Man-Computer Communications Society*, 7th Conference, Waterloo, Ontario, 274. DOI: [10.20380/GI1981.41](https://doi.org/10.20380/GI1981.41).



Information Technology Research

W. Richards Adrion

The NSF Information Technology Research (ITR) program greatly increased the funding base for CISE, opened up new research areas within computer science, and pushed CISE to be more interdisciplinary. Its extreme increases in workloads and difficulty in finding qualified reviewers for interdisciplinary proposals, however, also led to sustained stress for NSF staff.

ITR can be traced back to the High Performance Computing and Communications Initiative (HPCCI) and the resulting High Performance Computing Act of 1991¹ (HPCA), which established the President's Information Technology Advisory Committee (PITAC). As a National Research Council (NRC/CSTB) report indicates, the “difficulty of explaining and justifying federal IT research spending influenced the evolution and eventual transformation” of HPCCI, which led to “federal proposals for new and larger research programs, notably the 1999 Information Technology for the Twenty-First Century (IT²) initiative.”²

The effort of “explaining and justifying” federal information technology research fell to PITAC, led by co-chairs Bill Joy and Ken Kennedy. Their report, *Information Technology Research: Investing in Our Future*, recommended “a strategic initiative to support long-term research in fundamental issues in computing, information, and communications [that would] increase the total funding base by \$1.37 billion per year by FY 2004 . . . [and] use the resulting budget increases to encourage research that is visionary and high-risk.”³

Concurrently, a working group of the National Science and Technology Council (NTSC) was developing the 1999 Information Technology for the Twenty-First Century (IT²) federal initiative:

Table 8.1 Proposed allocations in president's FY 2000 budget implementation plan

Agency	Fundamental Information Technology Research and Development	Advanced Computing for Science, Engineering and the Nation	Social, Economic, and Workforce Implications of Information Technology	Total
DoD	\$100 million	—	—	\$100 million
DoE	\$6 million	\$62 million	\$2 million	\$70 million
NASA	\$18 million	\$19 million	\$1 million	\$38 million
NIH	\$2 million	\$2 million	\$2 million	\$6 million
NOAA	\$2 million	\$4 million	—	\$6 million
NSF	\$100 million	\$36 million	\$10 million	\$146 million
Total	\$228 million	\$123 million	\$15 million	\$366 million

the Federal Government is making an important re-commitment to fundamental research in information technology. The IT² initiative proposes \$366 million in increased investments in computing, information, and communications research and development (R&D) to help expand the knowledge base in fundamental information science, advance the Nation's capabilities in cutting edge research, and train the next generation of researchers who will sustain the Information Revolution well into the 21st Century. . . .

IT² builds on the Government's previous accomplishments and existing investments in High Performance Computing and Communications (HPCC), including the Next Generation Internet (NGI) and the DoE's Accelerated Strategic Computing Initiative. The IT² research agenda responds directly to the findings and recommendations of the President's Congressionally-chartered Information Technology Advisory Committee (PITAC), which concluded in a report released in February 1999.⁴

Six federal agencies (see Table 8.1) participated in IT²: the Department of Defense (DoD), Department of Energy (DoE), National Aeronautics and Space Agency (NASA), National Institutes of Health (NIH), National Oceanic and Atmospheric Administration (NOAA), and National Science Foundation (NSF). The Defense Advanced Research Projects Agency (DARPA) jump-started its efforts with a broad agency announcement (BAA) in late 1998 of Expeditions into the 21st Century, created "to encourage vigorous and revolutionary research in information technology (IT)."⁵ Among the DARPA-funded projects were Project Oxygen (MIT),

Endeavor (University of California at Berkeley), and Portolano/Workscape (University of Washington).

The NSF FY 2000 Budget⁶ request included \$110 million for research in software systems, scalable information infrastructure, high-end computing, and socioeconomic and workforce issues. Additional IT² funds supported centers and aimed at terascale computing systems. In late 1999, NSF asked for proposals under an agency-wide Information Technology Research program, which focused on eight areas: software, IT education and workforce, human-computer interfaces, information management, advanced computational science, scalable information infrastructure, social and economic implications of computing and communications, and revolutionary computing. For FY 2000, the NSF placed all of the \$90 million NSF ITR funds in CISE. This amounted to a large increase in both the CISE and NSF budgets. In addition, the CISE budget included \$36 million for a terascale computer system, bringing the total IT²-related increase in CISE funding to \$126 million.

For this, Ruzena Bajcsy, CISE Assistant Director (1991–2001), credited strong support from the community as well as a strong relationship with NSF Director Rita Colwell.⁷ (Bajcsy also cited three important members of the computer science community: Ed Lazowska, Andy van Dam, and Richard Newton.) While still the NSF Director Designate, Colwell had given a presentation entitled “Turning the Clock Forward” at the 1998 Computing Research Association Snowbird Conference. In her speech, she said “. . . computer science and engineering stands in singular stead today, as the science of creating, processing, and transforming information. It’s truly breathtaking to note the speed with which ideas in computer science spin out into the marketplace . . . in perhaps a third or a quarter of the time the process takes in most other disciplines.”⁸ Lazowska and several others at Snowbird made a strong case directly to Colwell for funding for the IT² initiative and for CISE leadership.

The ITR program addressed the areas singled out by PITAC (see Figure 8.1). The ITR program encouraged small projects of up to \$500,000 total over three years and large projects of up to \$15 million total over five years. Letters of intent were required for all proposals, with preproposals required for proposals over \$500,000. CISE received 1,800 letters of intent as well as 1,154 “small” proposals, and made 156 awards (13.5% of proposals) to 81 institutions. NSF received 1,350 letters of intent, 980 preproposals, and 263 proposals for more than \$500,000 (a screening review encouraged only 133 proposals). NSF made 75 awards (7.7% of preproposals) to 41 institutions. *Making IT Better* cited “the stresses on NSF program management caused by a large influx of researcher communications (e.g., letters of intent, preproposals, and proposals) that need to be evaluated and responded to by a fixed

staff already busy with ongoing responsibilities. Extraordinary efforts were made to recruit experts to participate in the necessary peer review.”⁹ This concern about staff workload came up in almost all reports and assessments—whether from a given Committee of Visitors, GPRA,¹⁰ or the CISE Advisory Committee.

Bajcsy assigned her division directors to coordinate directly with other directorates and related agencies. With this assignment and as Division Director for Experimental and Integrative Activities (EIA), W. Richards Adrion was the interface to Mary Clutter (AD of Biological Sciences), Margaret Leinen (AD of Geosciences), and Judith Sunley (AD of Education and Human Resources, succeeded by Judith Ramaley). Adrion also worked with the NIH Biomedical Information Science and Technology Initiative (BISTI) and the DARPA BioComputational Systems (BioCOMP) program. Michael Evangelist (DD for Computer—Communication Research) and Robert Borchers (DD for Advanced Computational Infrastructure and Research) were assigned to coordinate with Engineering and Mathematical and Physical Sciences (MPS). Michael Lesk (DD for Information and Intelligent Systems) was the first ITR program director.

Figure 8.1 shows the change in emphasis areas for FY 2001. Lesk commented in 2002:

the overwhelming considerations in allocation of money were (a) proposal pressure and (b) NSF inter-directorate budgetary issues. Within CISE, the money flowed to areas that were attracting an unusual number of proposals. This included subjects such as the “digital divide” and quantum computing, which are recent “hot” areas. It also included areas like natural language processing, where the rise of statistical techniques in place of rule-based techniques is new, rapid progress. Across directorates, MPS [Mathematical and Physical Sciences Directorate] managed to get a major share of the total ITR budget, resulting in considerable emphasis on proposals for applications of IT in the physical sciences and astronomy. GEO [Geosciences Directorate] was very efficient at getting its PIs to submit proposals, resulting in substantial funding for geosciences related proposals.¹¹

In reflecting on how MPS and GEO managed to get a large share of the ITR funds, Margaret Leinen, AD/GEO, deserves credit for seeing that ITR was an opportunity to expand their investments. GEO’s PIs highlighted their large-scale sensor networks and “big data” demands. MPS represents the oldest scientific disciplines whose PIs are highly skilled at obtaining funding. One of the largest FY 2000 awards was a \$12.3 million grant to the University of Florida for the GriPhyN (Grid Physics Network) team of seven IT research groups and four frontier physics experiments: the CMS and ATLAS experiments at the Large Hadron Collider, the Laser Inter-

FY 2000:

- Software: writing programs that work
- Scalable Information Infrastructure: making networks faster and accessible
- Information Management: finding and using information
- Revolutionary Computing: building new kinds of computers
- Human-Computer Interaction: helping all people to use machines and information
- Advanced Computational Science: using IT to advance the sciences
- IT Education and Workforce: helping train students of all ages at all levels
- Social Implications of IT: understanding how to maximize societal benefit

FY 2001:

- Systems Design and Implementation
- People and Social Groups Interacting with Computers and Infrastructure
- Information Management
- Applications in Science and Engineering
- Scalable Information Infrastructure for Pervasive Computing and Access

FY 2002:

- Software and Hardware Systems
- Augmenting Individuals and Transforming Society
- Advancement of the Frontiers of Science via Information Technology

FY 2003: No defined areas of interest

FY04: (National Priorities)

- Advances in Science and Engineering (ASE)
- Economic Prosperity and Vibrant Civil Society (ECS)
- National and Homeland Security (NHS)

Figure 8.1 Research focus areas defined in ITR solicitations.

ferometer Gravitational-wave Observatory (LIGO), and the Sloan Digital Sky Survey (SDSS). The *Report of Review Committee of NSF's High Performance International Internet Services (HPIIS) Project*¹² cited the ITR GryPhyN project as recognizing needs at the application level. In addition to grants involving GEO and MPS, the CISE Experimental and Integrated Activities Division's Biological Information Technologies and Systems (BITS) ITR grants led to longer-term programs in bioinformatics, computational biology and neuroscience, and biologically inspired computing.

Bajcsy¹³ had hoped that “small” proposals would attract junior faculty, but many senior faculty researchers applied and were successful. While some at NSF had wanted the large grants to be “centers,” Bajcsy did not want them to be managed in the same way as traditional NSF centers—with periodic reviews and significant

education and outreach components. “I really wanted the people [to] use the money for research,” she said, and research was clearly prioritized. The 2005 ITR Committee of Visitors cautioned:

the relatively large-scale nature of some ITR projects appeared inconsistent with the level of evaluation and oversight given them . . . management plans should be required . . . clear timelines and metrics of success [should] be established and linked to these management plans, and . . . these timelines and metrics of success [used] for oversight during the lifetime of the project.¹⁴

Frank Anger, who succeeded Lesk as ITR program manager, submitted a five-page ITR report¹⁵ for a CISE Division Director Retreat in 2002, pointing out the difficulty of comparing a \$1 million project to a \$10 million project as a reason for moving to a three-level competition for FY 2001: small proposals (less than \$500,000 over 3 years), medium or group proposals (\$500,000 to \$5 million over 5 years) and large proposals (up to \$15 million over 5 years).

Anger said that the FY 2001 structure was continued in FY 2002, dropping preproposals for medium proposals. This decision was based on the excessive number of panelists needed to handle preproposals, the extra burden placed on NSF staff, the already compressed timetable, and the belief that the difference in effort to prepare a good preproposal and a good full proposal was not that great. In short, preproposals for \$5 million projects caused a burden with little benefit.

A CISE Government Performance and Results Act (GPRA) report in 2001 noted that the NSF ITR program had expanded its emphasis on fundamental, high-risk R&D and on research and education activities that apply information, and was enabling research and education in multidisciplinary areas and emerging opportunities. “CISE is now using ITR funds to develop new thrust areas in CISE; they are indirectly helping to define a new core of programs in the five divisions. . . .”¹⁶ Anger pointed out, however, that “as a five-year program that opens new areas of research and increases the number of larger, multi-PI projects, there is a concern that PIs and areas may not continue to receive support as the ITR funds move into the base.”¹⁷

The interdisciplinarity of many ITR proposals made it difficult to find the proper expertise for review. The review process put a significant strain on the reviewing community and on NSF program officers. As the ITR COV report stated:

The size and interdisciplinary nature of ITR proposals challenge NSF’s traditional review and oversight procedures. The panels were required to be broader than usual, the proposals incorporated components (research, infrastructure, education, dissemination) which require different evaluation models, and the

medium and large-scale proposals require a greater degree of management and accountability. Given the breadth of the community involved it is difficult to assemble a strong, diverse, and conflict-of-interest-free pool of reviewers. While great efforts were made to ensure a sufficient number of appropriate reviewers, there was a general consensus that an increased use of quality mail reviews would have been beneficial.¹⁸

Anger argued that apportioning money to areas/panels/divisions/awards was problematic for several reasons: holding funds centrally (in CISE or later by a cross-directorate ITR committee) left panel award decisions dependent on multiple decision makers. Certain strategies for allocating funds to panels were reactive, slowed the award process, and caused “uncertainty, rivalry and gaming of the system.” He said that while much had been done to streamline ITR, the turnaround time for non-ITR proposals suffered. Proposals were not directly assigned to program directors but were often assigned to a panel on “the basis of the title and perhaps a scan of the Project Summary.”¹⁹ The COV added, “the sheer volume of proposals combined with the lack of NSF staff assigned to the ITR program led to some concern about the level of feedback provided to Principal Investigators.”

In FY 2003, the limit for medium proposals was lowered from \$5 million to \$4 million. In FY 2002 only one medium award was \$4 million or more, and only 10 of 94 awarded projects were \$3 million or more. The FY 2003 ITR Announcement abandoned prescribed areas and instead focused on goals and outcomes of the research, setting out more than a dozen objectives. For 2003, a 75-member coordinating committee (ITRCC) was involved in processing more than 2,500 ITR proposals. The program received 1,110 small proposals (889 in CISE and 301 in other directorates). CISE received 1,485 medium proposals and 67 preproposals and 106 full proposals. More than 800 proposals received in response to the fiscal 2003 ITR solicitation were related to homeland security.

The CISE advisory committee resolved to study the unprecedented nature of the ITR program.²⁰ Barbara Grosz (Harvard), Leonard Krishtalka (Kansas), Ralph Roskies (Pittsburgh), and Fred Schneider (Cornell, chair) served on the panel and carried out a review similar to a Committee of Visitors. Among their concerns were that when a specific panel had a particular focus, proposals not central to the focus might not be given adequate attention. In other words, proposal pressure by panel focus areas was fundamentally different from proposal pressure by subdiscipline. They were also concerned that guidelines setting a percentage of proposals in each panel to be rated “highly competitive” might skew the outcomes—for example, multiple panels in the same focus area might lead to more highly rated proposals




Figure 8.2 CISE funding FY 1999–2008. (Source: NSF Annual Requests to Congress FY 2001–2010 (using the “actual” budget numbers))

in that area. The panel was concerned about the lower success rate in ITR versus “regular” programs (see Figure 8.3). And not least, the doubling of proposals was accompanied by only a 5% increase in program staff.

Two important aspects of the ITR program—its interdisciplinarity and its openness to high-risk research—set it apart from other programs, but also present particular challenges to the review process. Within the context of the review process, we must seek more formal means (metrics) to evaluate the “high-risk/high payoff” nature of a proposal as well as to separate those proposals that are truly interdisciplinary from those that merely offer lip service to this goal. These metrics need to be provided to reviewers, panel members and program directors so that a clear set of criteria are established under which to evaluate these ITR-specific aspects of project review.²¹

Figure 8.2 shows the budgets for CISE divisions and the ITR program. In FY 2004, the last year of the ITR competition, the solicitation encouraged the submission of proposals targeting one or more of the “National Priorities.”²² These priorities encompassed a broad range of science and engineering research and education topics in which information technology (IT) plays a critical role.

Just before ITR ended in FY 2004, Peter Freeman reorganized CISE (see Chapter 4). He combined the two divisions with large infrastructure programs into a single Shared Cyberinfrastructure (SCI) Division. He moved the research activities

in ANIR and ACIR, combined them with programs in the remaining three divisions, restructured those divisions (as CCF, CNS, and IIS), and created program clusters. In FY 2005, the cyberinfrastructure programs moved from being a division in CISE to being an office reporting to the NSF Director: the Office of Cyberinfrastructure (OCI).

After FY 2004, a percentage of ITR funds were committed to ongoing obligations to FY 2000–2004 grants, some funds were invested in expanding programs in the research divisions, and some were directed toward new initiatives. For the FY 2004 ITR competition, the “national priorities” were realized by investments in Cyber Trust, Science of Design, and Information Integration. Of the 488 Cyber Trust proposals, CISE made 50 awards and used \$5 million in co-funding from DARPA. Of the 238 Information Integration proposals, CISE made 33 awards. In Science of Design, NSF received 182 proposals and made 24 awards.

For FY 2005²³ the emphasis areas included continuing Cyber Trust, Science of Design, and Information Integration and adding new emphases²⁴ on Broadening Participation in Computing (BPC) and Computational Science/High End Computing—Dynamic Data Driven Application Systems. For FY 2005, some ITR funds were used for the GENI Project (see Chapter 9). From FY 2008 forward, CISE used ITR funds partially to begin the Expeditions in Computing²⁵ program.

The CISE funding (success) rate continued to decline (See Figure 8.3), even with a substantial infusion of ITR funding. The magnitude and scale of the ITR program contributed to the declining success rates.




Figure 8.3 CISE funding rate.

ITR funding was a major event in the overall story of NSF investments in computing research. NSF recognized that the interaction among research disciplines—for example, studies of the social and psychological factors of human-machine interaction; the intersection of biological and computing research in genomics, neuroscience, and biologically inspired devices; the relationship of an “internet of things” to environmental, geological and social systems; and many more—created more opportunities for core computer science research. Accordingly, NSF gained awareness of the positive interaction between serious, investigator-driven, fundamental research and research done in the context of real applications.²⁶

Notes

1. The High Performance Computing Act of 1991 (HPCA) is an Act of Congress promulgated in the 102nd United States Congress as (Pub.L. 102–194) on December 9, 1991. Often referred to as the “Gore Bill,” it was created and introduced by then Senator Al Gore, and led to the development of the National Information Infrastructure (NII) and the funding of the National Research and Education Network (NREN).
2. National Research Council. 2000. *Making IT Better: Expanding Information Technology Research to Meet Society's Needs* (Washington, DC: The National Academies Press). <https://doi.org/10.17226/9829>.
3. B. Joy and K. Kennedy. February 1999. *Information Technology Research: Investing in Our Future*. National Coordination Office for Computing, Information, and Communications. Arlington, VA.
4. National Science and Technology Council IT² Working Group. 1999. *Information Technology for the Twenty-First Century: A Bold Investment in America's Future*. Proposed in the President's FY 2000 Budget Implementation Plan, June 1999. <https://www.nitrd.gov/historical/it2/it2-ip.pdf>.
5. National Research Council, 2000, *op. cit.*
6. See <https://www.nsf.gov/about/budget/fy2000/rps.htm>; last accessed 17 February 2019.
7. Oral history, Ruzena Bajcsy, interviewed by William Aspray, March 19, 2017. Charles Babbage Institute.
8. R. Colwell. July 27, 1998. “Turning the Clock Forward.” Computing Research Association Conference, Snowbird, Utah. NSF Office of Legislative and Public Affairs. <https://www.nsf.gov/news/speeches/colwell/rc80727.htm>.
9. National Research Council, 2000, *op. cit.*
10. Government Performance and Results Act.
11. M. Lesk. March 2002. “ITR—The First Two Years.” Draft internal report. Charles Babbage Institute.
12. Report of the 2004 High Performance International Internet Services (HPIIS) Review Panel 3, March 2004.
13. Bajcsy interview, 2017, *op. cit.*

14. M. J. Irwin. March 8–10, 2005. *FY 2005 ITR Committees of Visitors (COVs) Report*. Submitted to Peter Freeman by the CISE Advisory Committee chair, Alfred Z. Spector, on March 16, 2005.
15. Frank Anger. October 3, 2002. “ITR Program Five-Pager.” Internal document presented at CISE DD Retreat. Charles Babbage Institute.
16. National Science Foundation. October 2001. *CISE FY 2001 GPRA Performance Report*. <https://www.nsf.gov/pubs/2002/nsf02105/nsf02105.pdf>.
17. Anger, 2002, *op. cit.*
18. Irwin, 2005, *op. cit.*, p. 2.
19. Anger, 2002, *op. cit.*
20. B. Grosz, L. Krishtalka, F. Schneider (chair), and R. Roskies. March 2003. “Preliminary Study of ITR.” NSF CISE Advisory Committee. Charles Babbage Institute.
21. Irwin, 2005, *op. cit.*, p. 2.
22. FY 2004 NSF priorities included: Advances in Science and Engineering (ASE), Economic Prosperity and Vibrant Civil Society (ECS), and National and Homeland Security (NHS).
23. P. Freeman. August 2004. “Strategic Directions of the National Science Foundation and the CISE Directorate.” Presentation to National Science Foundation Computer and Information Science and Engineering. <https://slideplayer.com/slide/7463224/>.
24. D. Crawford. February 2005. “CISE Directorate—Overview and Funding Opportunities.” Presentation to National Science Foundation Computer and Information Science and Engineering. <https://ceas.uc.edu/content/dam/ceas/documents/College/Research/Crawford.pdf>.
25. The Expeditions in Computing program awards represent the largest individual research investments made by CISE, with each award providing up to \$10 million over five years. CISE funded 22 Expeditions projects that have yielded transformative breakthroughs in computing and information technology, catalyzed partnerships with industry to bring these technologies to bear, and given rise to entirely new sectors of the U.S. economy. See NSF Media Advisory 18-015, *Expeditions in Computing: 10 Years Transforming Science and Society*. https://www.nsf.gov/news/news_summ.jsp?cntn_id=297412&org=NSF&from=news; last accessed 22 March 2019.
26. See, for example: D. E. Stokes. August 2017. *Pasteur's Quadrant: Basic Science and Technological Innovation*. Washington, DC: Brookings Institution Press.



Networking Research and Deployment

W. Richards Adrion, Peter A. Freeman

In barely 25 years, the general public gained the ability to communicate easily and remotely by emails, video, and audio recordings; to exchange files of text and numerical data; and to access online services and information worldwide. This expanded capability has already revolutionized many major areas of modern life and has even become routine and an essential part of modern society's infrastructure. An "information society" is now recognized by social scientists of many disciplines (historians, economists, political scientists, and others) as well as diverse commentators, politicians, and many others.¹

This chapter provides a deeper coverage of NSF's support of networking research and initial deployment. The focus is on activities primarily related to the Computer and Information Science and Engineering (CISE) Directorate and its immediate predecessors. During the years covered by this chapter, the NSF Engineering Directorate supported fundamental optical and wireless technology developments that are essential to networking, the Math and Physical Sciences (MPS) Directorate supported research in physics and materials that is integral to chip technology, and more recently the Education and Human Resources (EHR) Directorate has supported various activities related to networking. We do not cover these non-CISE activities here.

NSF's role involved research *on* concepts and mechanisms of networking, and deployment *of* operational networks. In general, the "deployment" aspect is clear. By charter, NSF is not an operational agency responsible for providing services to citizens or carrying out mandated functions itself. The answer often given by NSF to those seeking support for an ongoing service such as a general campus computing facility is "NSF doesn't do that." In certain instances where the service is one that

no one else provides, such as operating a very large and expensive telescope, or is strictly in support of scientific research, such as providing training for teachers of science, NSF does deploy a service and sometimes continues to support an operational network.² The Computer Science Network (CSNET) and the National Science Foundation Network (NSFNET) are both examples where NSF for a time ran an operational network. NSF entered both of these projects with a specific goal to hand off the networks to self-supporting entities within a limited timeframe.

In his well-known 1945 description of Memex,³ Vannevar Bush envisioned something similar to today's World Wide Web, based on his wartime leadership of scientific research for the U.S. government.⁴ The earliest connections between multiple computers were in air defense systems such as SAGE.⁵ The earliest non-military network of computers was the Sabre⁶ air reservation and ticketing system developed by IBM and American Airlines around 1960.⁷ Almost all civilian⁸ connections to computers from remote locations through the mid-1970s were connections of remote terminals to a central computer over dial-up telephone lines.⁹

The technical foundation of current computer networking can be traced back to the late 1950s. At that time, support for networking in the U.S.¹⁰ was largely through the military and sometimes through their contractors. The common understanding today is that the origin of today's Internet was the ARPANET program of the Advanced Research Projects Agency (ARPA) of the U.S. Department of Defense, as created in the mid-1960s. ARPANET was designed as a packet-switched network¹¹ based on ideas first published independently by Donald Davies,¹² Len Kleinrock,¹³ and Paul Baran.¹⁴

The ARPANET project began after Bob Taylor moved from NASA to ARPA,¹⁵ replacing Ivan Sutherland, who had replaced the visionary J.C.R. Licklider. Taylor would soon hire 29-year old Lincoln Labs engineer Larry Roberts¹⁶ to head up the ARPANET project. Roberts' initial idea was to connect all ARPA-sponsored computers directly over telephone lines, which would place significant demands on the computers, especially considering the wide variety of operating systems and hardware. Initially, some ARPA-funded researchers did not see how they would benefit from sharing resources with other researchers.¹⁷ In 1967, Wes Clark suggested the idea of using separate small computers to standardize the network interface and reduce the load on the local large computers.¹⁸ By 1969, Bolt Beranek and Newman (BBN) was developing the Interface Message Processor (IMP), a small computer that handled communication between computers at UCLA, Stanford Research Institute (SRI), University of California Santa Barbara (UCSB), and University of Utah—the initial 4-node ARPANET. Further work for the Defense Department on ARPANET is an important story detailed elsewhere.¹⁹

By the mid 1970s, ARPANET had spread across the U.S. with international connections to the United Kingdom and Norway. Vinton Cerf and Robert Kahn proposed the Transmission Control Protocol (TCP),²⁰ which would later be split into TCP (Transmission Control Protocol) and IP (Internet Protocol). In 1975, the Defense Communications Agency (DCA) took over management of ARPANET. ARPA created one of the first modern internets and supported some of the earliest consequential innovations of which the TCP/IP protocol suite is the best known. As described in Chapter 2, by 1985 DCA had begun to use ARPANET as an operational military network following the cancellation of the new command, control, communications and intelligence (C3I) network, AUTODIN II. As a result, the DoD began a process of creating a military-only MILNET and moving “research” sites to a limited ARPANET research network. This action by DoD boosted CSNET and later the rapid expansion of NSFNET.

9.1 Early NSF-Supported Research on Networks (1950–1980)

ARPANET was developed by military agencies (largely ARPA through leading academic and industrial computer science research centers such as Stanford, UCLA, CMU, MIT, USC-ISI, BBN, and a few others), NSF’s support of a growing number of academic computer centers in the 1960s showed the power of computing and remote access for a wide variety of applications. It also helped to develop computer science as an academic discipline.²¹

After NSF support of campus computing facilities concluded in the early 1970s,²² support for basic research in computer science increased slowly. An early NSF-funded network research project²³ was awarded to the University of California, Irvine in 1971.²⁴ This created the first functioning, local-area network anywhere. The NSF Office of Computing Activities (OCA) in the early 1970s tried to develop national and regional research networks. As described in Chapter 2, OCA’s Don Aufenkamp announced plans at EDUCOM²⁵ and detailed the NSF networking plans at another conference.²⁶ What happened to this effort is not at all clear. Historian Janet Abbate²⁷ has suggested the reason for not pursuing these networks was the limited budget of the Office of Computing Activities (and its successor). The importance of networking to researchers took many years to be recognized.

By 1976, NSF leadership was not interested in anything like an ARPANET-like national network and was unconvinced that a network for sharing resources and for collaboration had sufficient value given the high cost. As detailed in Chapter 2, an opportunity for NSF to be involved in networking arose when four academic computer scientists (Lawrence Landweber, Richard Lipton, Richard DeMillo, and

Edward Robertson), with encouragement from NSF staffers Rick Weingarten and Rick Adrion, proposed Theorynet, a project that would use a locally developed mailbox system on the University of Wisconsin computer system and access it over a commercial packet-switched network to support researchers in theoretical computer science. Research collaboration rose steadily, and social scientist Starr Roxanne Hiltz found positive outcomes in terms of collaboration and jointly published papers resulting from the project.²⁸ Theorynet served over 100 theoretical computer scientists. The modest success of Theorynet lent credibility to the future CSNET and NSFNET projects.

9.2 NSF Leads Public Networking (1980–1995)

The confluence of three events in the late 1970s and early 1980s—the success of Theorynet in facilitating communication among theory researchers, the growing NSF investments to support experimental computer science research, and an increasing demand for access to high-performance computing—led to CSNET and NSFNET. Theorynet helped convince the NSF leadership that networks could increase research collaboration. The many reports and studies²⁹ that led to the NSF Coordinated Experimental Research (CER) initiative called for a network for computing researchers. While email, limited file transfer, and remote access services were beginning to be available commercially (examples include Compuserve, Tymnet, and Telenet), through academic computer centers (BITNET), and through research collaborations (the Bell Labs UUCP-based network, which would become the basis for Usenet), none of these met the goals of the CER initiative or the hopes of leading academic computer science departments to connect NSF-funded experimental computer science researchers with the ARPANET community and each other.

In parallel, Kent Curtis and Rick Adrion discussed with DARPA and others the possibilities for expanding ARPANET, and Larry Landweber and his colleagues began looking at alternatives. Cost and management issues were sticking points. NSF could not afford to replicate or expand ARPANET, and NSF believed it was not positioned to manage the development of an alternative. Eventually, in January 1981, the National Science Board (NSB) approved a five-year proposal for CSNET. Under Project Director C. William Kern, NSF would manage the project using a series of contracts, but only for two years (through 1983). The expectation was that NSF management would focus on setting up an organization that could collect and disburse funds so that the CSNET organization could become self-supporting within five years—funded by user fees.³⁰ The last NSF payment for CSNET opera-

tions was made in mid-1985. Under the leadership of the University Corporation for Atmospheric Research, Bolt Beranek and Newman, and a strong executive board, CSNET connected more than 165 university, industrial, and government computer research groups—serving more than 50,000 researchers and students—and provided accounts for 1000 Internet hosts. Numerous networks outside the U.S. were connected.³¹ CSNET became self-supporting and had significant industry funding. CSNET clearly had demonstrated, for the first time, that users were willing to pay for network services. Chapter 2 provides more details.

NSF created the Office of Advanced Scientific Computing (OASC) in May 1984, with John Connolly from the Materials Research division as director, Larry Lee from Mathematics as program director for centers, and Rick Adrion (on loan from the Division of Computing Research) as program director for networking. The OASC Networking Advisory Committee recommended the establishment of a “Sciencenet Phase 1”³² using available technology such as expanding and interconnecting ARPANET, BITNET, and commercial network services. A report by Landweber and David Farber recommended that NSF should (1) add a Sciencenet³³ manager and management team (or contract for such services); (2) establish a working group representing the centers, networks, and NSF management; and (3) establish a permanent Technical Advisory Committee (TAC).

Dennis Jennings, Director of the BITNET-based European Academic Research Network (EARN) and Computing Center Director at University College Dublin, accepted the NSFNET directorship and began work in January 1985. Jennings arrived with a clear vision for NSFNET.³⁴ He focused on developing a general-purpose network for science and engineering research rather than a network only for accessing the supercomputer centers. The difference of opinion between John Connolly and Gordon Bell about this issue led to the networking program being split off as a separate division in the CISE Directorate, as described in Chapters 2 and 3. Jennings adopted a “network of networks” approach incorporating “tiered” networks that included campus local area networks (LANs), regional networks, and a national backbone. The initial 56KB backbone was based on Dave Mill’s “fuzz-ball” PDP-11-based routers due to the high cost of ARPANET Interface Message Processors and the lack of commercial alternatives. The backbone connected the five OASC centers (Cornell, Illinois, Pittsburgh, Princeton, and San Diego) and NCAR with regional networks, which included among others NYSERnet, SURAnet, BARRnet, MIDnet, Westnet, MERIT, NorthWestNet, and NEARnet.

Jennings is best known for the adoption of the DoD TCP/IP and related ARPANET protocols as the standards for NSFNET. This decision enabled the NSFNET program to ask DARPA³⁵ to enable *all* users on campuses with ARPANET sites to access

the centers with ARPANET connections, via campus-wide networks, expanding availability beyond computer science and engineering departments.

Jennings left at the end of March 1986 and was replaced by Steven Wolff, who joined NSF from the Ballistic Research Laboratory at Aberdeen Proving Grounds. He led the NSFNET project and served as the Division Director for the CISE Networking and Communications Research and Infrastructure (NCRI) Division, which Gordon Bell created by moving NSFNET from the Division of Advanced Scientific Computing and adding communications and networking research programs from other CISE divisions. As Wolff recalls, “shortly after CISE was formed, the networking division was formed and, with me . . . as division director . . . it sticks in my mind that was April Fool’s Day of 1986 that that happened.”³⁶ Wolff also credited the strong group of NSFNET program directors³⁷ he was able to recruit to NSF and key members of the staff³⁸ originally recruited by John Connolly to OASC.³⁹

According to the Internet2 timeline,⁴⁰ in February 1987 the Southeastern Universities Research Association network (SURAnet) became the first operational regional network. Soon afterward, NYSERNet deployed a 56KB statewide network connecting New York’s leading universities and corporations to the Cornell and Princeton centers and became the first entity outside of the U.S. government to use the TCP/IP Internet protocols. NYSERNet was reported to be the inspiration for Jennings’s tiered network strategy.

In 1987, NSF provided \$14 million⁴¹, which would grow to \$58 million with amendments,⁴² to a consortium led by the Michigan MERIT⁴³ network, and including IBM and MCI, to re-engineer and manage the NSFNET, the first national high-speed Internet backbone.⁴⁴ A year later, the MERIT-MCI-IBM consortium had a T1 (1.5MB) backbone in place and was connecting the regional networks. The MERIT consortium cooperative agreement with NSF left open the possibility for opening up the network to the private sector. “It had to come,” noted Wolff,

because it was obvious that if it didn’t come in a coordinated way, it would come in a haphazard way, and the academic community would remain aloof, on the margin. That’s the wrong model—multiple networks again, rather than a single Internet. There had to be commercial activity to help support networking, to help build volume on the network. That would get the cost down for everybody, including the academic community, which is what NSF was supposed to be doing.⁴⁵

Wolff and his team at NSF designed the 1989 solicitation in “a way that would enable bidding companies to gain technical experience for the future.”⁴⁶ This was a decision that would bring questions from Congress as described below.

In 1988 NSF announced its intention to end support for NSFNET. In response organizations were created in anticipation of the potential of the Internet as a commercial as well as a research network. NYSERNet formed a commercial company, Performance Systems International (PSI, later PSINet), to manage its statewide network. Rick Adams, who had founded UUNET (built on the informal Usenet), and Bill Schrader, who led PSINet, created the first two commercial Internet Service Providers (ISPs). The MERIT-IBM-MCI partners created a non-profit organization, Advanced Network & Services, Inc. (ANS), to run the network infrastructure for the soon-to-be-upgraded NSFNET Backbone Service. In same period, ARPANET was decommissioned,⁴⁷ and DARPA researchers were moved to the NSFNET regional networks. Some controversies arose from the way in which ANS was allowed to take over management of the NSF backbone (discussed below).

A new T3 service, inaugurated by ANS in 1991, represented a thirty-fold increase in the bandwidth on the backbone. By 1992, over 6,000 networks were connected, one-third of them outside the United States. By 1995, NSFNET had spurred dramatic Internet growth. NSF's \$58 million investment in NSFNET, complemented by in-kind and other investments by IBM and MCI, resulted in 100,000 public and private networks in operation around the country. "The efforts to privatize the backbone functions had been successful," announced Paul Young, then head of CISE, "and the existing backbone was no longer necessary."⁴⁸ On April 30, 1995, NSF decommissioned NSFNET, turning over backbone services to commercial providers. The same year NSF funded MCI to bring up the very high-speed Backbone Network Service (vBNS), connecting a limited number of research universities with 155 mbps performance and experimental native Asynchronous Transfer Mode (ATM) functionality. The vBNS was later critically valuable as the initial backbone for the early Internet2 gigabit points-of-presence since the vBNS connections program expanded access for research and use.

When asked about his biggest disappointments, Steve Wolff identified "two mistakes":

One was trying to privatize the net prematurely. It was just bad judgment on my part. I thought that the commercial market was ready, and it wasn't. And in some sense, it still isn't, although that could be argued . . . but the greatest intellectual disappointment was [in] maybe 1989 or 1990, somewhere in there. Darleen [Fisher] kept pestering me with these proposals [for a] wireless internet and I said, "No, no, no. A, there's nothing there and B, we've got to get the wired stuff working first, then we can worry about that." That was a big mistake because we could have jump started the wireless work . . . by a couple of years anyway.⁴⁹

Asked why he left NSF, Wolff said,

. . . they [Rick Adams of UUNET and Bill Schrader of PSINet] objected to NSFNET on principle, that it was taking business away from them. And then, when we struck the deal with . . . Al Weiss and ANS [Advanced Network and Services] to share the circuits with the commercial enterprise, they just . . . you know, they went apoplectic and that was the . . . beginning of . . . all the FOIA [Freedom of Information Act] requests from Gordon Cook [the ‘Cook Report’]⁵⁰ and which of course then I think . . . raised the interest of the IG [NSF Inspector General]. And I think probably my last year there . . . I was mostly not running the division. I was answering FOIA requests and responding [to] the IG. So, in a sense, it was the right time for me to leave.

9.3 Transitioning to the Commercial Internet

Partly due to the events Wolff cites and to inquiries from Congressman Rick Boucher (D-VA),⁵¹ Chairman of the Subcommittee on Science of the House Science, Space, and Technology Committee,⁵² the NSF Office of the Inspector General investigated⁵³ the NSFNET. Both Congress and the public had raised questions about: (1) the NSFNET solicitation, evaluation of the proposals, and the award; (2) expansion of NSFNET and conversion to T3; (3) spinning off ANS & CO+RE; (4) administrative issues (form of contract, conditions, prior approval, compliance, accessibility, funding strategy); and (5) the future of NSFNET.

As background, the NSFNET Solicitation closed on August 14, 1987, with the receipt of proposals from six applicants, including MERIT. While three were considered technically responsive to the solicitation, NSF decided that on the basis of cost, the MERIT proposal should be funded, and the National Science Board approved a five-year cooperative agreement to implement and manage the NSFNET backbone in November 1987. The NSF Office of the Inspector General (OIG) found “NSF’s decision to award the Cooperative Agreement for NSFNET to MERIT was reasonable.”⁵⁴

In addition, the Division of Networking and Communications Research and Infrastructure (DNCRI) asked MERIT to expand the NSFNET backbone to T3 (45 mbps). In November 1990, the NSB extended the authorization limit so the entire backbone could provide T3 speeds, completed in mid-1992.

As part of its investigation of NSFNET, the NSF Office of the Inspector General found:

NSF’s decision to upgrade NSFNET to T3 before the T1 network was saturated was reasonable. [The OIG] also believed the price was not unreasonable. Nonetheless, it is legitimate to question whether NSF should have issued a new solicitation for

the node expansion and/or the T3 conversion, rather than increasing the existing award . . . [however,] the NSFNET solicitation explicitly envisioned expansions of and improvements to the network, and the public was on notice that the successful offeror would be responsible for expansions of the NSFNET backbone within the period of the award.⁵⁵

As previously mentioned, MERIT, IBM, and MCI had formed a non-profit corporation called Advanced Network & Services, Inc. (ANS) to take over management and operation of NSFNET. NSF specifically allowed ANS to “solicit and attach to the NSFNET Backbone new users, including commercial users, and may connect them to new or existing nodes on the Backbone.” Commercial users would reimburse ANS for connection costs, the added traffic, and related support, and reimbursements would be used to enhance the network infrastructure and services so the level of service to NSF would not be diminished. In May 1991, ANS created ANS CO+RE Systems, Inc., a for-profit corporation, to engage in activities beyond scientific research and education. The NSFNET Acceptable Use Policy (AUP) prohibited purely commercial traffic from using NSFNET. When CO+RE was formed, NSF program staff agreed with MERIT and ANS that the MERIT/ANS network operations center, IBM-provided routers, and MCI-provided lines were not subject to the AUP because NSF was not paying for equipment and facilities but was instead paying for the conveyance of NSFNET traffic and the provision of network support services.

The NSF OIG found that “NSF reasonably concluded that allowing commercial use of the network—with the conditions NSF imposed—is consistent with NSF’s overall statutory mandate.” The OIG further concluded that “it was not unreasonable for NSF to decide that allowing MERIT to permit some commercial traffic over the network created by MERIT was consistent with the objectives of NSF and the NSFNET program.”

In response to concerns that CO+RE granted an advantage over other network providers by virtue of ANS’s relationship with NSF, the OIG recommended that “for the remaining period of the amended Cooperative Agreement, NSF ensure that other network providers continue to be offered access to the T3 network on the same terms as CO+RE, and, if the offer is accepted, then access is provided fully and fairly.”⁵⁶

The OIG interviewed Wolff and his staff, the MERIT staff, and others to reconstruct the decision-making process. Consistent with Wolff’s comments above, the Inspector General found that “the record is utterly barren of documentation of NSF’s reasoning for allowing commercial use of the network.”⁵⁷ NSF staff had acted appropriately, the OIG said, but without adequate documentation and record

keeping and the OIG had particular concerns with MERIT being allowed prior approval rights under the Federal Demonstration Project⁵⁸ and compliance with Circular A-110. The OIG also found that because NSF intended other commercial network providers to have access to the backbone on the same terms accorded to CO+RE, NSF should have affirmatively announced this to the networking community.

As the Internet evolved, by 1992, it included several government or government subsidized backbones or regional networks, a couple dozen regional/mid-level networks, and thousands of private (industry, university, and institutional) networks including private for-profit commercial mid-level and wide-area nets (commercial backbones).

The US portion of the Internet is made up of . . . Federally subsidized components such as NSFNET, NASA Science Internet (NSINET), Energy Sciences NET (ESNET) and DARPA Test Net . . . that have agreed to interconnect and carry each other's traffic . . . commercial networks (PSINet, CERFnet, UUNET/ALTERNET) that are linked together via a commercial internet exchange (CIX) and, via some of its members, linked to the NSFNET backbone. . . . International connections have been established through government agreements or through business negotiations by the commercial networks.⁵⁹

On June 15, 1992, NSF published a draft solicitation for a Network Access Point (NAP) Manager and Routing Authority (RA) and a provider of very high speed Backbone Network Services (vBNS) and began coordinating with other federal networks. This was the path toward decommissioning NSFNET in 1995.

By 1993, a controversy arose concerning NSF's role in second-level Internet domain name registrations. In 1983, Jon Postel, Paul Mockapetris, Craig Partridge, and others contributed to the design, testing, and implementation of a domain name system (DNS) to make Internet navigation easier.⁶⁰ Shortly thereafter, the Internet Engineering Task Force defined seven "top level domains,"⁶¹ and the USC Information Science Institute and Jon Postel began managing Internet domain registration and allocation of Internet Protocol (IP) numbers. In November 1987, the Defense Communications Agency transferred control of IP numbers from Postel and ISI to the Network Information Center (NIC) at SRI International. In 1991, Network Solutions, LLC, was awarded the contract to operate the domain name registry (for .com, .org, .mil, .gov, .edu, and .net) on behalf of the U.S. Defense Information Systems Agency (DCA's successor)—free of cost to customers.⁶² Two years later, Network Solutions (NSI) was the sole bidder for the contract for operating domain registry service for the National Science Foundation for "second

level domains,” within the .com, .net and .org domains on NSFNET (e.g., some-company.com). Initially, and also when most of the registrants were educational institutions, NSF paid the entire cost of Internet second-level domain name registration. By 1995, the Internet experienced substantial growth in commercial participation that would have overwhelmed NSF. An independent panel recommended that NSI and NSF amend their “cooperative agreement to create a self-sustaining fee-based system, and [include] a provision for 30% of registration fees to be placed in a custodial account.”⁶³ The fee was \$50 with \$15 going to an Internet Intellectual Infrastructure Fund, which totaled \$45.5 million by 1998.⁶⁴ These funds were used to support NSF research and infrastructure grants. Following a number of legal actions⁶⁵ (e.g., alleging the fee was a “tax” or that NSI had a monopoly) and decisions by the Clinton administration, the agreement ended in 1998. Fees were reduced and eventually NSI lost its sole control of all Internet registries.

In this time period, two related NSF-funded activities would come to have significant impact on the Internet.⁶⁶ The first was the development of the leading early web browser, Mosaic, at the National Center for Supercomputing Activities (NCSA) at the University of Illinois at Urbana-Champaign. The second was the development of a powerful search engine developed by Google. Both emerged from large NSF-supported projects.

Tim Berners-Lee, a researcher with the European Organization for Nuclear Research (CERN), developed the World Wide Web, its hypertext structure, and the protocols necessary to access web pages on diverse computers. Berners-Lee implemented the first web server and web pages, but CERN quickly adopted the Berkeley Viola browser.⁶⁷ Viola was limited to running within the Unix X windows system, opening the door for Marc Andreessen and Eric Bina at Illinois and NCSA. They developed Mosaic with funding from Larry Brandt in NSF/OASC. Mosaic successors such as Netscape Navigator (by a company founded by Andreessen and Jim Clark), Internet Explorer, Safari, Firefox, and Chrome built on the Mosaic graphical user interface (GUI) characteristics such as the URL address bar, back/forward/reload buttons, and other interactive elements.

Prior to Google, there were a number of Internet search engines: Excite (Stanford), Yahoo! (Stanford, initially a directory), Webcrawler (University of Washington, bought by Excite), Lycos (Carnegie Mellon), Infoseek (based on Inquiry from the University of Massachusetts Amherst), AltaVista (Digital), Inktomi (Berkeley), and many others. Several of these can be traced to NSF funding. Google⁶⁸ was born in the Stanford InfoLab and supported by an NSF grant for the Stanford Integrated Digital Library Project (IRI-9411306) under the direction of Hector Garcia-Molina. Sergey Brin and Lawrence Page were graduate students in the InfoLab that included

Garcia-Molina, Rajeev Motwani, Jeff Ullman, and Terry Winograd. Google was built on the PageRank algorithm⁶⁹ that Brin, Page, Motwani, and Winograd developed. Google before long eclipsed all other search engines.

While NSF was supporting NSFNET, it helped fund the Gigabit Network Testbed Initiative that ran from 1989 to 1995. In 1989, NSF and DARPA provided \$20 million for five testbeds to explore long-distance networking issues and applications a thousand times faster than the NSFNET backbone at 1.5 megabits per second bandwidth.⁷⁰ This initiative became a joint activity with industry. Network service providers and technology companies contributed an estimated \$400 million and—along with NSF's NCSA, SDSC, and Pittsburgh supercomputing centers—deployed the testbeds and participated in the research. The testbeds explored advanced networking issues, investigated architectural alternatives, and carried out experimental applications in diverse areas such as weather modeling, chemical dynamics, radiation oncology, and geophysics data exploration.

The Gigabit Testbed Initiative provided a new type of research collaboration among network and application researchers, the computer science and telecommunications communities, and academia/industry/government research teams. It also leveraged government investments with substantial contributions from industry. Three statewide high-speed networks resulted: the North Carolina Information Highway (NCIH) formed by BellSouth and GTE based on the Vistanet testbed, the NYNET experimental network formed by NYNEX as a result of their Aurora testbed involvement, and the California Research and Education Network (CalREN) created by Pacific Bell as a result of their Casa testbed participation.⁷¹

9.4 Network Research after NSFNET (1996–2001)

As NSF began to privatize NSFNET in the mid-1990s, the Foundation contracted with MCI to establish the very high-speed Backbone Network Service (vBNS). It was to serve as an infrastructure for advanced networking research and to support scientific research without interacting with general Internet traffic. The vBNS operated in parallel to the commercial backbone networks that replaced the NSFNET backbone. Researchers accessed vBNS via the NSF supercomputer centers and NSF-specified Network Access Points, where the vBNS connected to other federal research networks.

Once the vBNS was in place, NSF established a High-Performance Network Connections program to support universities and colleges to connect to high-performance networks. By 1999, 150 institutions across all 50 states were connected to the vBNS. The vBNS operating speed of 155 megabits per second far exceeded the

45 mbps offered by commercial Internet service providers (ISPs) and supported new network technologies, such as IPv6, to meet the special needs of advanced applications. By 2000, the vBNS backbone was upgraded to 2.4 gigabits per second (OC-48). Commercial connections to vBNS were also offered for the first time in 2000.

President Clinton and Vice President Gore had announced the Next-Generation Internet (NGI) initiative back in October 1996. It included \$300 million over three years to connect universities and national laboratories with high-performance networks and to promote next-generation networking technologies. The vBNS became a key part of the NGI, and the NSF Connections program helped more than 150 institutions connect to the vBNS and the Internet2⁷² consortium's Abilene network—exceeding the NGI goal of 100 institutions. NSF also supported advanced networking applications development and research on high-performance networking. Under NGI, the Science, Technology, and Research Transit Access Point (STAR TAP) in Chicago connected six U.S. research networks, including vBNS and Abilene, and 12 international research networks⁷³ by 2000 when the NGI initiative ended successfully.

Concurrent with vBNS, NSF initiated a research program to provide technical and engineering support and overall coordination of the vBNS connections. The National Laboratory for Advanced Network Research (NLANR) was created in 1995 as a collaboration among the NSF supercomputer centers. As the vBNS evolved into a stable leading-edge platform and other high-speed networks were formed, NLANR expanded its focus and served as technical support for High Performance Network Service Providers such as Internet2 and STAR TAP.

A related activity began in March 2002, when Larry Peterson and David Culler convened a meeting of researchers interested in “planetary-scale network services” and proposed PlanetLab⁷⁴ as a community testbed. With support from Intel, PlanetLab grew to 100 nodes at 42 sites within six months. In February 2003, PlanetLab was online via Internet2's Abilene backbone. NSF announced a \$4.5 million award to Princeton, Berkeley, and Washington for enhancing PlanetLab in September 2003. As PlanetLab evolved, it was used to develop VINI⁷⁵ and deployed on the National Lambda Rail (NLR) and Internet2's NewNet backbone. It also contributed to Measurement Lab,⁷⁶ an open source project to provide an open, verifiable measurement platform for global network performance. PlanetLab was one of the ideas in the GENI Project (described below).

In 2003 a consortium of leading U.S. research universities and private sector technology companies deployed the National LambdaRail (NLR), a national networking infrastructure to foster next-generation network-based applications in science, engineering, and medicine. The 12,000-mile, high-speed national NLR

computer network was originally owned and operated by the U.S. research and education consortium that created it without corporate partners. Unfortunately, NLR struggled to provide reliable services, underwent several changes in leadership, and failed to pursue “integrated applications, systems, and network research.”⁷⁷ From 2006 to 2007, Internet2 and NLR discussed a merger, but “they couldn’t get past three main points of contention: the transfer of assets to the merged entity, its commitment to research, and the role of regional network groups under the new organization.”⁷⁸ Eventually, NLR was purchased by billionaire Patrick Soon-Shiong in November 2011 to use for healthcare applications and was shut down in March 2014.

9.5 Networking Research (2002–2004)⁷⁹

When Peter Freeman became AD/CISE in early 2002, one of his objectives was to emphasize networking research. In his early assessment of the NSF networking research activities it quickly became clear that some new directions were needed. As described in a near-contemporaneous description,⁸⁰ four important steps were taken beginning in early 2003 that started to reshape CISE support for networking research:

[1] Supported by one of the workshops⁸¹ mentioned [earlier in the referenced source], the first was the announcement of twin testbed funding programs: Experimental Infrastructure Network⁸² (EIN) program and the Networking Research Testbeds (NRT) program. The program solicitation for EIN notes that the purpose is to: “establish, address, explore, and experiment with next generation network infrastructure technologies to meet the rapidly emerging requirements of e-Science and other advanced applications which are not being addressed by today’s research networks (e.g., Abilene or vBNS) or the Internet.” Concurrently, NRT⁸³ set out to create a new generation of networking technologies through the process of ideation, realization and experimentation carried out on a diverse set of research testbeds.

Together they funded several smaller scale networking testbeds including PlanetLab, ORBIT, Emulab, DETER. These testbeds have been playing a very important role in shaping GENI.

[2] A major reorganization of the CISE Directorate officially took place in November 2003. One result was to separate the networking research activities from the operational networking responsibilities.

[3] As a new strategic direction was being set for the new CNS Division, more emphasis was placed on research in the networking domain than had previously

been the case. The decoupling of network research from the operational and near-term activities laid the foundation for future focus not coupled to any particular operational paradigm.

[4] Finally, the reorganization, increased and focused budget, and initial programmatic activities under EIN and NRT enabled the last, key initial event—the hiring of Guru Parulkar.

The next major phase of CISE involvement in networking can be dated to early 2004 when Guru Parulkar integrated a number of ideas from the networking research community into a proposal to CISE management.

9.6 Initiation of the GENI Project (2004–2006)

By 2004 the Internet was having great impact worldwide. The NSF-supported research described just above⁸⁴ and a few other funders and companies were extending the original functionality of the Internet. At the same time, serious concerns were being raised in the technical community about the Internet’s ability to meet not only expansive visions of a networked future but also the critical daily activities of millions of people. The situation was qualitatively different from that existing when the original technological developments were made, tested, and deployed.

Scott Shenker, in an ACM SIGCOMM 2002 keynote address, lamented “that the SIGCOMM community had been so successful in building the Internet that it was now locked in a box in terms of longer-range, larger improvements.”⁸⁵ The number of opportunities for networking researchers to have impact was shrinking. Shenker further opined that the success of the networking research community in aiding the birth of the Internet had three vital components: “Intellectual Depth, Transformative Impact, and Community.” He then addressed the future (for the networking research community) and addressed the central question: “How can we retain the three important concepts?” His answer: “Focus on transformative community projects that can engage a community and transform the world.”

Shenker’s comments mirrored several NSF-supported workshops between 1996 and 2000⁸⁶ and private comments made by several respected networking experts to Peter Freeman. This situation in the research community matched the situation internally in CISE in 2002: there was effectively no overarching strategy.

Outside the technical community—from Congress and the executive branch of the government, to business and industrial leaders, to ordinary citizens—people were becoming aware of problems of security, privacy, capacity to transmit large volumes of data, real-time response for control of critical infrastructures, and

support of new network-based applications. Industry was starting to respond, but because of the operational importance of the existing networks and their tendency to address only their own concerns—sometimes in ad hoc ways—it was difficult to obtain general solutions to the problems. Technical leaders in the research and operational communities understood that basic research and experimentation with a variety of solutions would be needed, and that only the government could conduct the kind of long-range, pre-competitive research that was needed. Some people also understood that, especially in the political environment after 9/11, only NSF would have the latitude to undertake such research.

The driving forces for developing a more comprehensive approach to network research included the fundamental needs to build security and robustness into the designs of networks, bridge the gap between mobile devices and stationary networks, provide for real-time and highly reliable control of critical machinery; and enable new classes of services. The technical community observed several additional drivers: ossification of the then-current Internet architecture, inability to experiment with and test proposed new technologies at scale and under realistic traffic loads, the “push” of initially unrelated technological developments such as optical switching and mobile devices capable of network operations, and the “pull” of potential new networked applications on the horizon, such as telemedicine.

The needs and promises of a major new project were clear enough to the leadership of CISE, but as of early 2004, no one had yet proposed such a project, nor were they actively searching for one.⁸⁷

9.7 The Origin of GENI (2004)

GENI didn’t just suddenly pop into existence. Indeed, the name “GENI” wasn’t created until over a year after the effort started at NSF; it was initially called *CIRI—Clean-slate Internet Re-invention Initiative*.

The start of GENI was in April 2004, when Guru Parulkar made a presentation to Peter Freeman and Deborah Crawford. He reviewed problems with the current Internet, including lack of security, capacity, and capabilities (e.g., real-time control of distant facilities). He then outlined the difficulty or impossibility of creating the much-anticipated digital world. Parulkar identified a “Gang of Four” dealing with the looming “brick wall” that the Internet faced.⁸⁸

At the end of his presentation, Parulkar asked for guidance. Freeman and Crawford’s immediate response was that they believed the direction he had just described was an exciting idea that fit extremely well into several broad national concerns, including cybersecurity, innovation, and competitiveness. They encour-

aged him to focus on this activity and to provide planning grants to the Gang of Four, with the idea of developing the ideas to the point that it could become a full-fledged project. Within a month, a request for a planning grant had been submitted, properly reviewed, and granted.

From this very first meeting, Freeman and his colleagues tried to move this as far and as fast as possible. Through the end of 2006, the majority of GENI activity was internal planning and planning grants to small teams of academic researchers.⁸⁹ This form of direct, top-down action was fairly rare in NSF. Indeed, the CISE leadership told Parulkar immediately that they were going to push this concept in a “DARPA-like” fashion, but within the framework of NSF protocols of community involvement in setting directions and having impartial reviews. Crawford’s guidance based on her senior roles at NSF was invaluable.

Freeman and Crawford knew of Parulkar’s record of research and practical innovation but were not actively searching for a new major project nor looking to him to supply it. Indeed, it took Freeman almost a month to find time in his schedule to meet the first time with Parulkar for an hour. Nonetheless, his presentation captivated both with its clarity, vision, innovativeness, inclusiveness, and attention to many of the drivers noted above. Further, it immediately suggested an important generalization, which, when overlaid on the technical aspects, transformed them into a major research infrastructure project at NSF. It suggested to them how the development of the overall idea could potentially provide valuable research infrastructure for other research areas beyond networking.

9.8 Conceptual Design for GENI (2004–2006)⁹⁰

The ideas that Parulkar had pulled together from research going on in the community, especially by Tom Anderson (University of Washington), Larry Peterson (Princeton), Scott Shenker (Berkeley), and Jon Turner (Washington University, St. Louis), together with his own research, had three key ideas: *slicing* (having multiple networks using the same routers, with the ability to switch between them dynamically), *virtualization* of a network, and *programmability* of routers. All had existed, but Parulkar combined them in a new way.

The diagram he presented to Freeman and Crawford that day⁹¹ became an avatar for GENI (see Figure 9.1). The overall concept was of a continental-scale network for supporting *experimental*, network-based, research *at scale* (i.e., traffic loads and demands) and in a manner that would allow graceful degradation of operations to “commodity” Internet service when needed or desired. Moreover,

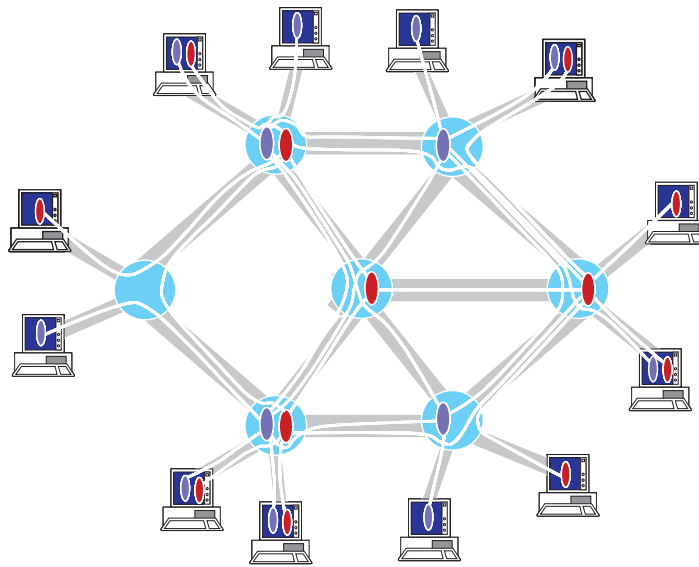


Figure 9.1 Virtualized overlay networking concept. Red and blue represent two potentially very different networks sharing the same physical infrastructure.

multiple networks could utilize the same equipment at the same time and a specific network could be *reconfigured dynamically*.

With weekly coordination meetings and other direct organizational support from the Office of the AD, led by Peter Freeman and Deborah Crawford, the next 20 months became an extremely active period. By the end of the summer of 2004, proposals for six planning grants had been reviewed and funded, and the research groups were already at work; these groups involved several dozen researchers (faculty, graduate students, and professional staff). Starting in January 2005, a series of workshops explored possible uses of such a facility and sought ideas relevant to the conceptual design. These workshops ultimately engaged well over a hundred members of the CISE research community.

Concurrently, CISE began briefing NSF management, OMB, and Congressional staff; holding several (five or six) ad hoc group meetings of senior computer scientists not involved in the GENI work to get their opinions and suggestions; and discussing the project with leaders in industry and elsewhere in the government. CISE hired a certified project manager, an experienced network designer, and administrative support to pull together the above efforts and help fashion a coherent, conceptual design.

In August 2005, the leadership of CISE convened an ad hoc panel of senior computer scientists not involved in the project on the day before SIGCOMM 2005, in which we presented the conceptual design and requested feedback. The group included Paul Baran,⁹² Alan Kay, Ed Lazowska, Bob Kahn, Vint Cert, David Farber and several others. The discussion was robust and helpful.

At the SIGCOMM2005 meeting, an announcement of the GENI Project⁹³ was distributed to all attendees, noting that the facility would require the following: a rich set of link technologies, including IP tunnels, guaranteed bandwidth packet switched paths, dedicated circuits, and optical lightpaths;⁹⁴ a flexible set of node configurations, including virtual machines running on commodity processors, dedicated processors, customizable hardware, and ultimately integrated optical switching systems;⁹⁵ enhanced diversity of devices at the network's edge;⁹⁶ and a significant development effort to translate the results of various CISE research programs into heavily instrumented architectures and services that could be deployed and evaluated on the experimental infrastructure.

The remainder of 2006 was spent wrapping up the conceptual design, continuing outreach to interested parties, and planning to transition the project to complete community control through long-term cooperative agreements. This last activity involved the creation of a Computing Community Consortium (CCC)⁹⁷ and a GENI Project Office (GPO).

9.9 Implementation of GENI (2007–2016)

The CCC is a platform, similar to some of the boards of the National Academies, where members of the research community can come together, study critical areas, and suggest research agendas to various agencies (not just NSF). The first area studied was how to best utilize a GENI facility. The CCC has continued up to the present and issued a number of additional reports, two of which have served as the basis for Presidential Initiatives.

The solicitation for a contractor to staff and run a GPO in cooperation with CISE resulted in several high-quality responses.⁹⁸ Bolt, Beranek, and Newman (BBN, now a division of Raytheon) was chosen and began operation on July 1, 2007, under a 5-year cooperative agreement with CISE. After a mid-term review, the initial award was later extended for another five years.⁹⁹

Summarizing¹⁰⁰ the first 10 years, the GPO, together with the research community, has accomplished the following:

- utilized a spiral-development model involving many research groups in the community;

- successfully built a continental-scale, GENI experimental facility (despite the very small staff size of the GPO);
- held numerous workshops introducing researchers to GENI, as well as to how to use it to run experiments;
- enabled the publication of hundreds of peer-reviewed articles and a book; and
- engaged the campus enterprise computing centers (generally distinct from computers for research) on over 50 campuses, placing “GENI racks” in many of their machine rooms to enable wide usage of GENI on those campuses.

The GENI effort, including a related program in CISE¹⁰¹ called the Future Internet Architecture program, occupied much of the networking research community; and by some reports, it went on to produce a new generation of researchers focused on discovering new mechanisms and processes.

As NSF networking support increasingly went to GENI-related projects beginning in late 2007, GENI naturally dominated academic networking research.¹⁰² As planned, CISE ended its support for the GPO in mid-2017.

9.10 Conclusion

This chapter does not cover everything NSF has done in networking. It does, however, provide an overview of NSF’s involvement in networking with references to more detailed accounts. We believe that the contributions in this area made by NSF-funded researchers are basic and foundational.

NSF’s role in bringing about the networked world we live in today is a good example of how investigator-driven, basic research sometimes leads to unexpected results. A more focused example is the comment of a faculty member that, around 1985, he and a graduate student had published a small, somewhat obscure result in computer science theory. Five years later, he learned that a developer at Google was using their result in the auction mechanism for placing advertisements—reaping Google untold revenue.¹⁰³

Research results may differ considerably from what was imagined. While the original objective of funded work may be achieved, for example, connecting scientists remotely to their equipment, the unanticipated results may have much broader results—in this case, by showing how to deal with massive amounts of data. Big data is now changing the modern world far beyond science and engineering.

We add that most basic research has never been so successful and impactful on the broader world. There is one exception to that maxim, however, and that is the educational benefit to students at all levels of learning to explore the world about them, collect data about it, and attempt to use that to devise new knowledge.

Notes

1. For example: W. Aspray and P. Ceruzzi, eds. 2010. *The Internet and American Business*. Cambridge, MA: MIT Press; M. Castells. September 8, 2014. The impact of the Internet on society: A global perspective. *MIT Technology Review*. <https://www.technologyreview.com/s/530566/the-impact-of-the-internet-on-society-a-global-perspective/>; J. Manyika and C. Roxburgh. October 2011. *The Great Transformer: The Impact of the Internet on Economic Growth and Prosperity*, McKinsey Global Institute. https://www.mckinsey.com/~ /media/McKinsey/Industries/High%20Tech/Our%20Insights/The%20great%20transformer/MGI_Impact_of_Internet_on_economic_growth.ashx.
2. An example of this is the ongoing support for XSEDE. <https://www.xsede.org>.
3. V. Bush. July 1945. As we may think. *Atlantic Monthly*.
4. G. Z. Pascal. May 1999. *Endless Frontier*. Cambridge, MA: MIT Press.
5. K. Schaffel. 1991. *The Emerging Shield*. Washington DC: Office of Air Force History, United States Air Force; K. C. Redmond and T. M. Smith. October 2000. *From Whirlwind to MITRE: The R&D Story of the SAGE Air Defense Computer*. Cambridge, MA: MIT Press.
6. R. V. Head. 2002. Getting Sabre off the ground. *IEEE Annals of the History of Computing*, 24(4): 32–39. DOI: [10.1109/MAHC.2002.1114868](https://doi.org/10.1109/MAHC.2002.1114868).
7. M. Campbell-Kelly. 2004. *From Airline Reservation Systems to Sonic the Hedgehog*. Cambridge, MA: MIT Press.
8. Military and very large commercial applications could afford dedicated telephone circuits, and some of those applications had direct connections between different computers in near, or perhaps remote, locations.
9. There were several examples of much earlier connections between computers. The Sabre system was driven by two mainframe computers connected by cable (for redundancy), and military systems undoubtedly also had such direct connections. Another type was a mainframe connected to a minicomputer, which was connected to terminals that permitted limited real-time interaction with terminals and then was able to submit batch jobs to the mainframe and return results to the minicomputer and its terminals. Peter Freeman built such a system in 1985 at the University of Texas at Austin, running an interactive FORTRAN compiler that provided real-time syntax checking to users at terminals. See *Design Considerations for a Time-shared Computer* by P. A. Freeman and *SYCK—63, A Fortran-63 Syntax Checker* by R. E. Fikes, both 1965 M.S. theses, University of Texas at Austin. Their work was supported by the University of Texas Computation Center, which had significant support from NSF at the time.
10. Similar support was underway in England (<http://www.historyofcomputercommunications.info/Book/4/4.11-NPLNetworkDonaldDavies66-71.html>; last accessed 13 January 2019) and the Soviet Union (“Why the Soviet Internet Was Doomed from the Start,” <http://www.bbc.com/future/story/20161026-why-the-forgotten-soviet-internet-was-doomed-from-the-start>; last accessed 13 January 2019), and perhaps elsewhere.
11. In-depth coverage of the concepts of packet switching and Baran’s contributions is included in an 11-volume series of technical reports: Rand Corporation. 2009. *Paul Baran and the Origins of the Internet*. Santa Monica, CA, <https://www.rand.org/about/history/baran.html>, last accessed 13 January 2019; also see: J. Abbate. 2000. *Inventing the Internet*. Cambridge, MA: MIT Press.

12. Donald Davies of the British National Physical Laboratories independently developed the idea for and had named the approach *packet switching*. Davies delivered a paper, co-authored with K. A. Bartlett, R. A. Scantlebury, and P. T. Wilkenson, at the ACM Conference on Operating Systems Principles in 1967 that first alerted Larry Roberts to packet switching and to Paul Baran's work at RAND.
13. L. Kleinrock. 1961. Information flow in large communication nets. *RLE Quarterly Progress Report 1*, https://www.lk.cs.ucla.edu/bibliography-public_reports.html.
14. P. Baran. 1964. On distributed communications networks. *IEEE Transactions on Communications Systems*, 12(1): 1–9. DOI: [10.1109/TCOM.1964.1088883](https://doi.org/10.1109/TCOM.1964.1088883).
15. J. E. O'Neill. Winter 1995. The role of ARPA in the development of the ARPANET, 1961–1972. *IEEE Annals of the History of Computing*, 17(4): 76–81. DOI: [10.1109/85.477437](https://doi.org/10.1109/85.477437).
16. See Internet Pioneers: Larry Roberts. <https://www.ibiblio.org/pioneers/roberts.html>; last accessed 20 February 2019.
17. Abbate, 2000, *op. cit.*
18. K. Hafner and M. Lyon. 1996. *Where Wizards Stay Up Late: The Origins of the Internet*. New York: Simon & Schuster.
19. A. L. Norberg and J. E. O'Neill. 2000. *Transforming Computer Technology: Information for the Pentagon, 1962–1986*. Johns Hopkins University Press.
20. V. Cerf and R. Kahn. 1974. A protocol for packet network intercommunication. *IEEE Transactions on Communications*, 22(5): 637–648. DOI: [10.1109/TCOM.1974.1092259](https://doi.org/10.1109/TCOM.1974.1092259).
21. Chapters 1 and 6 discuss some of this material.
22. See Chapter 1.
23. NSF Award #7001116, “Design of Distributed Computing Systems.” PI: D. J. Farber, start date 7 January 1971.
24. D. J. Farber. 1973. The three faces of computer networks. *Computer*, 6(8): 10–12. DOI: [10.1109/MC.1973.6541669](https://doi.org/10.1109/MC.1973.6541669).
25. D. Leavitt. May 3, 1972. NSF-sponsored research could lead to national science computer network. *Computerworld*, p. 10.
26. D. Aufenkamp and E. C. Weiss. 1972. NSF activities related to a national science computer network. *Computer Communications: Impacts and Implications*, vol. 226.
27. Abbate, 2000, *op. cit.*
28. S. R. Hiltz. June 1981. *The Impact of a Computerized Conferencing System on Scientific Research Communities*. Final Report NSF-MCS-77-27813.
29. The Feldman and Snowbird reports, and articles in ACM *Communications* by Peter Denning, which described the serious problems facing experimental computer science research, are discussed in Chapter 2.
30. D. Comer. 1983. The computer science research network CSNET: A history and status report. *Communications of the ACM*, 26(10): 747–753. DOI: [10.1145/358413.358423](https://doi.org/10.1145/358413.358423).
31. L. Landweber. 2012. InternetHistory.Asia. <http://InternetHistory.Asia>. Downloaded from <https://sites.google.com/site/internethistoryasia/book1/3-2-csne>; last accessed 4 March 2018.

32. W. R. Adrion, D. J. Farber, F. F. Kuo, L. H. Landweber, D. C. Nagel, and J. B. Wyatt. 1984. *SCIENCENET: Report on the Evolution of a National Supercomputer Access Network*. Washington, DC: National Science Foundation.
33. Because someone previously had trademarked “Sciencenet,” the network quickly became known as NSFNET.
34. D. M. Jennings, L. H. Landweber, I. H. Fuchs, D. J. Farber, and W. R. Adrion. 1986. Computer networking for scientists. *Science*, 231(4741): 943–950. DOI: [10.1126/science.231.4741.943](https://doi.org/10.1126/science.231.4741.943).
35. Jennings et al., 1986, *op. cit.*, p. 946.
36. Oral history, S. Wolff, interviewed by Rick Adrion, July 20, 2017. Charles Babbage Institute.
37. Jane Caviness, Doug Gale, Priscilla Huston, and George Strawn served in various roles in NCRI and with NSFNET. Doug Gale created the Internet Legacy Institute where one can find materials and interviews about the Internet. See <http://www.internetlegacyinstitute.org>; last accessed 25 March 2019.
38. Including Dave Stout, Dan van Bellingham, Don Mitchell, and Darleen Fisher.
39. Wolff interview, 2017, *op. cit.*
40. Internet2 Community Timeline. <https://www.internet2.edu/about-us/internet2-community-timeline/>; last accessed 23 February 2019.
41. The MERIT cooperative agreement included cost sharing of \$5,000,000 from the State of Michigan for facilities and personnel, approximately \$6,000,000 from MCI in reduced communication charges, and \$10,000,000 from IBM in equipment, installation, maintenance, and operation.
42. NSF 8720904 Merit Operations and Management of NSFNET Backbone. https://www.nsf.gov/awardsearch/showAward?AWD_ID=8720904&HistoricalAwards=false; last accessed 25 March 2019.
43. MERIT was the Michigan Educational Research Information Triad, a non-profit corporation managed by a consortium of eight Michigan universities: University of Michigan, Central Michigan University, Eastern Michigan University, Michigan State University, Michigan Technological University, Oakland University, Wayne State University, and Western Michigan University. The company’s name has been legally changed to MERIT Network, Inc.
44. Internet2 Community Timeline.
45. NSF. *The Internet—Changing the Way We Communicate*. <https://www.nsf.gov/about/history/nsf0050/pdf/internet.pdf>; last accessed 24 February 2019.
46. Ibid.
47. Mark Pullen in a presentation at the 20th NSFNET Anniversary Celebration took credit for decommissioning ARPANET and moving the DARPA PIs to the NSF regional networks, thus creating a larger networking research community. See *Panel—NSFNET: The Community*. Moderator: Doug Gale. Panelists: Sidney Karin, Richard Mandelbaum, J. Mark Pullen, Glenn Ricart, Henry E. Shaffer, Jim Williams. <https://web.archive.org/web/20170202190232/http://nsfnet-legacy.org/archive.php>; last accessed 24 February 2019.
48. NSF. *The Internet—Changing the Way We Communicate*. <https://www.nsf.gov/about/history/nsf0050/pdf/internet.pdf>; last accessed 24 February 2019.

49. Wolff interview, 2017, *op. cit.*
50. <http://www.cookreport.com/>.
51. Boucher authored the legislation that permitted the first commercial use of the Internet (“the Boucher Amendment”)—the relevant language from which was enacted as Section 4 of Public Law 102-476 (102d Congress), October 23, 1992, 106 STAT. 2297. This language amended the National Science Foundation’s “acceptable use policy,” which only permitted the NSFNET to carry research and educationally related material, and authorized commercially oriented traffic on the NSFNET.
52. Management of NSFNET. Hearing before the Subcommittee on Science of the Committee on Science, Space, and Technology, U.S. House of Representatives, One Hundred Second Congress, Second Session. Congress of the U.S., Washington, DC. House Committee on Science, Space and Technology. March 12, 1992. <https://files.eric.ed.gov/fulltext/ED350986.pdf>; last accessed 24 February 2019.
53. NSF Office of the Inspector General. April 23, 1993. Review of NSFNET. OIG9301. National Science Foundation. <https://www.nsf.gov/pubs/stis1993/oig9301/oig9301.txt>; last accessed 24 April 2019.
54. OIG9301.
55. OIG9301.
56. OIG9301.
57. OIG9301.
58. See <http://thefdp.org/default/about/history/>; last accessed 24 February 2019.
59. A. Oldehoeft. February 1992. *Foundations of a Security Policy for Use of the National Research and Educational Network*. NIST.
60. P. Mockapetris. November 1987. *IETF Request for Comments 1035: Domain Names—Implementation and Specification*. Network Working Group. Information Sciences Institute. <https://tools.ietf.org/html/rfc1035>; last accessed 24 March 2019.
61. The top-level domains included .com, .net, .org, and .gov.
62. ICANNWiki entry on Network Solutions. https://icannwiki.org/Network_Solutions; last accessed 4 March 2019.
63. Lawrence Rudolph. March 24, 1999. General Counsel, National Science Foundation. Testimony before the House Committee on Science Subcommittee on Basic Research.
64. National Science Foundation. March 16, 1998. *NSF and NSI End Internet Intellectual Infrastructure Fund Portion of Domain Name Registration Fees*. News Release 98-017.
65. American Association for the Advancement of Science. 1998. NSF hits Internet jackpot. *Science*, 280(5365): 813. <https://doi.org/10.1126/science.280.5365.813b>.
66. See Chapter 3.
67. M. Lasar. 2019. Before Netscape: The forgotten web browsers of the early 1990s. *ArsTechnica*. <https://arstechnica.com/information-technology/2019/05/before-netscape-forgotten-web-browsers-of-the-early-1990s/>; last accessed 10 October 2019.

68. S. Brin and L. Page. 1998. The anatomy of a large-scale hypertextual web search engine. *Computer Networks and ISDN Systems*, 30(1-7): 107–117. DOI: [10.1016/S0169-7552\(98\)00110-X](https://doi.org/10.1016/S0169-7552(98)00110-X).
69. L. Page, S. Brin, R. Motwani, and T. Winograd. 1999. *The PageRank Citation Ranking: Bringing Order to the Web*. Stanford InfoLab.
70. Corporation for National Research Initiatives. December 1996. *The Gigabit Testbed Initiative: Final Report*. <http://www.cnri.reston.va.us/gigafr/>.
71. The activities from 1985 to 1995 are described in Chapters 2 and 3 in this book as well as in several readily accessible sources: B. M. Leiner, V. G. Cerf, D. D. Clark, R. E. Kahn, L. Kleinrock, D. C. Lynch, J. Postel, L. G. Roberts, and S. Wolff. 2009. A brief history of the Internet. *ACM SIGCOMM Computer Communication Review*, 39(5): 22–31, DOI: [10.1145/1629607.1629613](https://doi.org/10.1145/1629607.1629613); B. M. Leiner, V. G. Cerf, D. D. Clark, R. E. Kahn, L. Kleinrock, D. C. Lynch, J. Postel, L. G. Roberts, and S. S. Wolff. 1997. The past and future history of the Internet. *Communications of the ACM*, 40(2): 102–109, DOI: [10.1145/253671.253741](https://doi.org/10.1145/253671.253741). Several recent interviews of key players, conducted mostly in 2017 and 2018, are available at the Charles Babbage Institute and provide additional information and points of view. Among these are Larry Landweber, Gordon Bell, Rick Adrion, Charles Brownstein, Steve Wolff, Dennis Jennings, Irene Lombardo, and George Strawn. A highly informative description of NSFNET and some of what came before can be found in the *NSFNET Final Report* published by the MERIT regional networking consortium: <https://www.MERIT.edu/about/> explains the organization, and a section can be found on the site pointing to their history (now 50+ years long).
72. A group of universities formed the not-for-profit Internet2 consortium in 1996 to develop new Internet technologies and capabilities. Today, Internet2 has more than 220 university members; more than 60 corporate sponsors, partners, and members; and more than 40 affiliate members, including NSF.
73. Supported by the \$4.5 million NSF High Performance International Internet Services (HPIIS) program.
74. B. Chun, D. Culler, T. Roscoe, A. Bavier, L. Peterson, M. Wawrzoniak, and M. Bowman. 2003. Planetlab: An overlay testbed for broad-coverage services. *ACM SIGCOMM Computer Communication Review*, 33(3): 3–12. DOI: [10.1145/956993.956995](https://doi.org/10.1145/956993.956995).
75. A. Bavier, N. Feamster, M. Huang, L. Peterson, and J. Rexford. 2006. In VINI veritas: realistic and controlled network experimentation. *ACM SIGCOMM Computer Communication Review*, 36(4): 3–14. DOI: [10.1145/1159913.1159916](https://doi.org/10.1145/1159913.1159916).
76. C. Dovrolis, K. Gummadi, A. Kuzmanovic, and S. D. Meinrath. 2010. Measurement lab: Overview and an invitation to the research community. *ACM SIGCOMM Computer Communication Review*, 40(3): 53–56. DOI: [10.1145/1823844.1823853](https://doi.org/10.1145/1823844.1823853).
77. R. J. Aiken, J. Boroumand, and S. Wolff. 2004. Network and computing research infrastructure: Back to the future. *Communications of the ACM*, 47(1): 93. DOI: [10.1145/962081.962086](https://doi.org/10.1145/962081.962086).
78. O. P. Malik. November 4, 2007. *National Lambda Rail, Internet2 Merger Is Off. Again*. GIGAOM Report (blog). <https://gigaom.com/2007/11/04/nlr-internet2-merger-off/>.
79. These sections, down to “Implementation of GENI (2007–2016),” draw heavily on the September 6, 2007, unpublished, private report, “The Origins of GENI,” in Charles

Babbage Institute ; OAD/CISE. August 1, 2005. "GENI: Global Environment for Networking Investigations," internal memo in Charles Babbage Institute; GENI Planning Group. September 2006. "GENI Design Principles," *IEEE Computer*, 102–106; P. A. Freeman. August 28, 2014. "The GENI Vision," technical report in Charles Babbage Institute; and on personal recollections of Peter Freeman, who authored these sections. Direct quotations will be identified, as well as other specific sources.

80. Freeman, 2007, *op. cit.*
81. Several NSF-supported workshops are referenced in Freeman, 2007, *op. cit.* Because such reports are often difficult to locate, in 2007 Freeman made copies of those that were still accessible. That collection is in Charles Babbage Institute.
82. "EIN Program Solicitation," 2003, <http://www.nsf.gov/pubs/2003/nsf03539/nsf03539.htm>.
83. NRT Program Solicitation, <http://www.nsf.gov/pubs/2003/nsf03538/nsf03538.htm>.
84. See also Chapter 4 of this book.
85. *Whither SIGCOMM? SIGCOMM Past, Present, and Future*. Keynote address by Scott Shenker at SIGCOMM2002. Charles Babbage Institute.
86. For example, *National Conversation on NSF Advanced Networking Infrastructure Support*, AAAS, Washington, DC, February 22–23, 1999, and Leesburg, VA, May 27, 1999, and *Future Priorities for NSF Networking Activities*, CNRI/XIWT, C. Brownstein (ed.), Reston, VA, May 12, 1999. Both in Charles Babbage Institute .
87. There were many demands on the senior leadership of CISE at that point in other areas; the internal organization relative to networking research was just beginning to stabilize with new leadership in CNS and the cluster of programs including networking; and no compelling, comprehensive ideas had yet come from the community.
88. *The Gang of Four Plus One*: The core ideas for GENI emerged from the work and ideas of Larry Peterson, Scott Shenker, Tom Anderson, and Jon Turner (and Parulkar). Internally, they were referred to as "The Gang of Four." The "Plus One" referred to Parulkar but also referred to David Clark of MIT in some contexts since he had provided valuable, private perspective to Freeman on several occasions.
89. Freeman, 2007, *op. cit.*
90. GENI Planning Group, 2006, *op. cit.*
91. Parulkar's original presentation file is not available today, but the diagram shown here was embedded in OAD/CISE, 2005, mostly likely by Parulkar, who was one of the staff responsible for the memo.
92. Paul Baran participated fully in the meeting by phone, but had to hang up before the meeting ended. The next day Peter Freeman received an email from Baran indicating his full support for the concept, and urging the team to continue the project immediately. (Personal remembrance of Peter Freeman.)
93. "GENI Initiative," announcement prepared by OAD/CISE, in Charles Babbage Institute.
94. A key attribute of these link technologies is that they expose fine-grain control over bandwidth aggregation and direct failure effects to the network architecture.

95. A key attribute of these node technologies is that they can be programmed (controlled) to support architecture-specific logic.
96. Including traditional desktop computers, low-cost handheld devices, sensor networks, and programmable RF devices. The devices at the edge of the network should allow researchers to emulate a wide range of end-system behavior.
97. *About the CCC*, <https://cra.org/ccc/about/>; last accessed 8 March 2019.
98. See Chapter 2.
99. *GENI Project Office, Phase 2*, September 22, 2011, NSF Awards Database, AWD_ID=1125515.
100. Based on Freeman's personal knowledge and brief analysis of the GENI Bibliography, <http://www.geni.net/about-geni/geni-bibliography/>; last accessed 11 January 2019 (a dynamic bibliography of known papers relating to GENI).
101. NSF Future Internet Architecture Project. <http://www.nets-fia.net>; last accessed 13 January 2019.
102. A comprehensive overview and analysis of the project is yet to be written but three sources are: (1) M. Berman, J. S. Chase, L. Landweber, A. Nakao, M. Ott, D. Raychaudhuri, R. Ricci, and I. Seskar. March 2014. GENI: A federated testbed for innovative network experiments. *Computer Networks*, 61: pp. 5–23, (a paper that describes the overall approach and architecture); (2) R. McGeer, M. Berman, C. Elliott, and R. Ricci. 2016. *The GENI Book*. Springer International Publishing; and (3) the GENI Bibliography, <http://www.geni.net/about-geni/geni-bibliography/>.
103. Private comment to Peter Freeman in 2000 by a faculty member at Georgia Tech.



NSF Support of High-Performance Computation¹

Peter A. Freeman

High-performance computers are among the most complex engineered artifacts yet created, resting on fundamental research in physics, electrical engineering, and computer science supported by NSF and other sponsors. Yet the design of the machines themselves has almost always been an endeavor funded by industry. NSF has focused its support instead on the use of these powerful machines in the natural and social sciences to tackle some of the most challenging, fundamental problems facing humankind.

Compounding the problem of understanding the intertwined histories of the *creation* of very advanced computers and their *use* is the fact that both histories are connected to the process of creating digital computing devices in general. At NSF this has led to a certain tension between the user community (initially physical, geological, and atmospheric sciences and more recently the biological and social sciences) and the research communities on which computer designers rely (primarily physics, engineering, and computer science).² That tension has caused NSF's organizational responsibility for providing advanced computing services to migrate among different directorates and offices, even to this day.

Merely describing these high-end machines can be complex, as multiple names have been used: “advanced computers,” “supercomputers,” and “high-performance computers,” among others. In the 1950's, the few computers in existence were simply called “computers” or even “computing machines”; other terms came into use later. “High-performance computer” is currently most frequently used and often just abbreviated “HPC”; further, it is often used as a noun

(e.g., who funds HPC?) and sometimes as an adjective (e.g., what is your HPC budget this year?). The accompanying systems and applications software, peripherals (especially storage), and other factors such as the computational problem being addressed also are relevant factors today. We will use “HPC” interchangeably with other terms and trust that the context will make it clear.³

This chapter is intended only to be a framework and set of references to permit one to dig more deeply into the history of HPC generally and at NSF in particular.

10.1 1950–1954

Extending the capability of humans to perform computations accurately and rapidly is an endeavor almost as old as civilization. As so often has been the case with invention and ingenuity, however, development of the modern computer was hastened by the needs of the military—initially in the Second World War, leading to dramatic increases in computing capability to aim weapons at rapidly moving targets, break communication codes, perform nuclear calculations, design new weapons, and carry out many other tasks.

Although most of this work was carried out in secrecy, by 1945 the usefulness of such devices for scientific and technological calculations and processing of information was starting to be more widely understood. Vannevar Bush’s article in the July 1945 *Atlantic Monthly*⁴ foresaw this future. The development and use of HPC during the war, especially in the further development of nuclear weapons and code breaking, continued and was supported by the various mission-oriented agencies of the United States government, as well as by the contractors that supported their efforts.

By 1950, when NSF finally came into existence,⁵ the scientific need for HPC was clear to those whose research demanded it. In 1953, the idea of using a computer to perform an experiment to understand a basic scientific issue was demonstrated at Los Alamos.⁶ The first four NSF annual reports barely mention computation or computers, aside from support of some workshops on the use of computers in geology and meteorology.⁷ However, an NSF study in 1953 did note the increasing use of computers in research of various kinds.⁸

10.2 1955–1983

In 1955, however, the NSF annual report noted: “In order to provide the Foundation with informed advice as to the computer needs of modern science, and its possible role in assisting universities in meeting these needs, an ad hoc Advisory Panel on University Computing Facilities was appointed in February 1955.”⁹ Members of the

panel included John von Neumann (chair), Edward Teller, S. M. Ulam, and J. Barkley Rosser, among other notable scientists from a variety of fields.

In May 1955, the National Science Board (NSB) determined that NSF should provide computers to universities. A 1956 study¹⁰ by the Math, Physics, and Engineering Science (MPE) Directorate foresaw an HPC machine for use by researchers supported by NSF, at a cost exceeding \$5 million. A formal funding program was not established until 1959, but a number of proposals for computers (and some research on the design of advanced computers) were funded in the interim. The number of very powerful computers in existence anywhere prior to the early to mid-1960s was fairly small, so differentiation between common computers and HPCs would be largely meaningless.¹¹

In 1962, NSF Director Alan Waterman requested the National Academy of Science (NAS) to study “the status and likely growth of computer uses” in the areas of research and education of relevance to NSF. The study was prompted by Philip Morse, a prominent physicist at MIT, calling for more computing power. J. Barkley Rosser headed the study, and the report was finished in 1966. In 1967, the President’s Science Advisory Committee (PSAC) commissioned a similar study headed by John R. Pierce, head of communications research at Bell Labs.

These reports, plus a directive from President Johnson, led NSF to create the Office of Computing Activities (OCA) in 1967. OCA provided an institutional home for NSF support of equipment grants, research on computers, and related theoretical studies. One of its early hires was Kent Curtis, recently from Lawrence Berkeley Laboratory, where he had overseen a “supercomputer center.” For a few years, NSF continued to support computing facilities on university campuses until the demand outstripped NSF’s budgets. By the early 1970s, OCA ended its support for general facilities and limited computing equipment to specific research efforts, leaving support for general scientific computing to other parts of NSF, other agencies such as the Department of Energy (DoE), or local, state, or campus support.

By the late 1970s many in the research community felt future scientific advances would be impeded by the lack of advanced computers. An interagency study group, led by Peter Lax of NYU, reported (“The Lax Report”) in late 1982 that access to HPC and the design of future HPC machines was woefully inadequate.¹² A four-part federal program was proposed. In mid-1983, an internal NSF working group, led by Marcel Bardon and Kent Curtis, recommended that NSF provide “supercomputer services for academic research and science education” and support “networks linking universities and laboratories with each other.”¹³

By the end of 1983, the stage was set for NSF to resume serving the scientific community with HPC. An Office of Advanced Scientific Computation (OASC) advisory

committee further strengthened the call for NSF action.¹⁴ (Chapters 1 and 7 in this volume provide a deeper and fuller description of NSF's activities during this period.) The appointment of Erich Bloch as NSF Director in June 1984 was an inspired choice on multiple dimensions, not the least of which was his knowledge of HPC.¹⁵

10.3 1984–1991

Two intertwined activities—providing HPC for academic scientists and designing future HPC systems—significantly changed the scientific provision of HPC assets. One was unique to NSF and short-term; the other involved NSF collaboration with other agencies, as well as with industry, and was much longer-term.

By the fall of 1984, Bloch had addressed the above HPC recommendations.¹⁶ Initially awards to several universities and commercial entities provided time on Cray machines to the scientific community. Then OASC quickly issued a solicitation for proposals, and in early 1985 the first three awards for NSF supercomputer centers were made¹⁷ to the San Diego Supercomputer Center (SDSC) at the University of California at San Diego (UCSD), the National Center for Supercomputer Applications (NCSA) at the University of Illinois at Urbana-Champaign, and the John von Neumann Center (JvNC) at Princeton University.

A fourth center at Cornell University developed an advanced prototype computing facility. Later known as the Cornell Theory Center, it was led by Kenneth Wilson, a Nobel laureate in physics and strong HPC proponent. In 1986, the Pittsburgh Supercomputing Center (PSC) was established by a consortium of Carnegie Mellon University, the University of Pittsburgh, and the Westinghouse Corporation. These were the original five NSF supercomputer centers.¹⁸

The JvNC at Princeton had ordered a supercomputer that in the end did not meet the contracted requirements, and the manufacturer went out of business.¹⁹ NSF gave the JvNC additional time to develop an alternate plan, but Erich Bloch withdrew NSF funding as of April 1990, and the Center closed down.²⁰ The Cornell Center for Advanced Computing continues today as an important center for research and education, but it also eventually lost center-scale NSF funding in the mid-1990s.

The four centers that still remained in 1990 had their support extended through 1995. In renewing the awards, the NSB asked for a report “. . . to investigate the future changes in the overall scientific environment due [to] the rapid advances occurring in the field of computers and scientific computing.” A blue ribbon panel²¹ chaired by Lewis Branscomb, former head of the National Bureau of Standards and later chief scientist of IBM, documented the scientific successes of the NSF centers and the substantial lead in HPC that the U.S. enjoyed at that time. The strong recommendation was for the NSF and the country to capitalize on these results.

As the end of the original ten years of support for the supercomputer centers approached, another distinguished review committee was appointed. This committee was headed by Edward F. Hayes, vice president of research at Ohio State and a former NSF program manager; it also included future NSF Director Arden Bement, the second Assistant Director (AD) of CISE Bill Wulf, and the then-current AD of CISE, Paul Young. Funding for the four remaining centers was extended two more years to permit the committee time to do its work. The report²² (known as the “Hayes Report”) further documented the advances that had been made, carefully considered various alternatives, made multiple recommendations for continuing the supercomputer program, and extended the findings of the previous reports.

This support for HPC that began in 1985 has continued in one form or another up to the present (2019). Advanced scientific computing has enabled countless scientific discoveries²³ and provided one of the original drivers for the creation of the Internet.²⁴

Before describing the longer-term activity that started in the mid-1980s, we should mention the Semiconductor Research Corporation (SRC), a non-profit consortium created in 1982 by the Semiconductor Industry Association. Erich Bloch, then at IBM, was the first chairman of the SRC Board. When he came to NSF in early 1984, he ensured that NSF would provide research funding to SRC. Its research was intended to help U.S. industry regain its leadership in semiconductors.²⁵

The longer-term activity began with a “systematic review of the status and directions of high performance computing” by the Federal Coordinating Council for Science, Engineering, and Technology (FCCSET) Committee. Their report was submitted in late 1987.²⁶ Gordon Bell, first AD/CISE, was involved in this report and other FCCSET matters.²⁷

The Office of Science and Technology Policy (OSTP) in late 1989 created the Program Plan called for in the strategy document.²⁸ This plan was produced by an interagency group representing multiple agencies—authored by Department of Energy (DoE), Department of Defense (DoD), and NSF representatives.²⁹ Drawing on the Program Plan, on May 18, 1989, Senator Al Gore (D-TN) introduced Senate Bill 1067: “To provide for a coordinated Federal research program to ensure continued United States leadership in high-performance computing.” A later version, known as the High Performance Computing Act of 1991, was enacted on December 9, 1991, and is colloquially known as the “Gore Bill.” It led to the development and funding of the National Research and Education Network (NREN) and advanced HPC.

Bloch, Bell, and later Wulf—all of whom were acknowledged experts in computer architecture—understood the details of advanced computation and shared a vision of what it could enable. This led to a period of unrivaled NSF leadership in

HPC, which moved NSF into the forefront of academic computing for science and engineering.

10.4 1992–2000

The rapid advance of computing, the emergence of the Internet, and the explosion of computer usage created an environment in which NSF often struggled to keep up with demand and research opportunity in HPC. In retrospect, it seems as though a new or revised HPC program was barely started before studies and panels were convened to recommend a successor.³⁰ In reality, this mirrored the turmoil in the larger world of enterprise computing, except that the technical issues in scientific research often exceeded commercial problems by an order of magnitude. Only a few years into the NSF Supercomputer Centers program, calls mounted for additional HPC capacity. This culminated in 1995 with the Hayes Report.³¹

It clearly called for continuation of the Centers program, but also for broadening access to computational and other related resources. NSF followed up with the Partnerships for Advanced Computational Infrastructure (PACI).³² Two initial awards were made in March 1997: to the National Computational Science Alliance, led by NCSA, and to the National Partnership for Advanced Computational Infrastructure (NPACI), led by the San Diego Supercomputing Center (SDSC). Both sites were among the original centers; and together they included 100 cooperating universities across the country with some overlap. NSF support was phased out for the Cornell and Pittsburgh centers.

The PACI partners were leaders in grid computing (multiple computers cooperating on a single problem), cloud computing (where computation, storage, and other resources are provided as a service), and very large databases (“big data”). These new applications, based on older concepts and prototypes, were driven by massive amounts of scientific data and computational architectures that could support simulations requiring vast amounts of processor time. This not only advanced science but showed the way for industrial and commercial applications that eventually became commonplace.

10.5 2000–2004

The seminal PITAC report, “Information Technology Research: Investing in Our Future,”³³ greatly increased NSF-CISE funding for computing research. Partly in response to the Information Technology Research (ITR) Report, PSC received an award in mid-2000 for a terascale computer.³⁴ NSF issued a solicitation for a Distributed Terascale Facility (DTF), calling for proposals to acquire terascale com-

puters, terabyte storage systems, and gigabit networks.³⁵ The first awards were made in August 2001.³⁶ Another award enabled PSC’s terascale machine to join the DTF, and in 2003 to add additional resources to the expanding grid—then called the Extensible Terascale Facility (ETF). In 2004 the grid entered full production mode, bringing to a turning point this expansion of NSF capabilities provided to the science community.³⁷ By then it was clear to most observers that, for many scientific problems, it was essential to have available the most advanced computational infrastructure possible; it was no longer a discretionary choice in order to be competitive.

10.6 2005–Present

An effort around 2000³⁸ deployed some of CISE’s expanded funding to support a state-of-the-art “cyberinfrastructure”³⁹ to advance computational science and engineering. Following NSF protocol, the NSB appointed a blue ribbon panel headed by Dan Atkins. The Atkins Report⁴⁰ was submitted in early 2003. While strongly arguing for HPC, the Atkins Report also advocated for the democratization of science and engineering and emphasized resources that allowed experimentation and analysis of data at scale in multiple ways.

Planning for implementation began while the Atkins Report was still in draft form. Ultimately NSF created a new Office of Cyberinfrastructure reporting directly to the NSF Director. The Office been reorganized several times, but to this day NSF continues to provide academic researchers with the latest computational resources.

The organizational and programmatic changes at NSF regarding HPC for the general scientific community are a result of two intertwined forces: the continuing, even accelerating pace of change in the technologies available and the ability of the general scientific community to utilize them. The pace of change is a well-known story. On the other hand, the adoption and utilization of new technology often takes surprisingly longer.

Usually research scientists do not immediately drop what they are doing just because more powerful computational tools appear. New and more powerful hardware might even be useless for those who are not computing experts. In the mid-1980s, for example, very powerful parallel machines entirely lacked basic algorithms, languages, compilers, and other software. This type of situation has been repeated because, even as the non-hardware tools are being developed, the hardware development continues apace.

NSF’s organizational changes, then, are often a direct result of this push-pull. The organizational home of HPC at NSF has oscillated between the Office of the

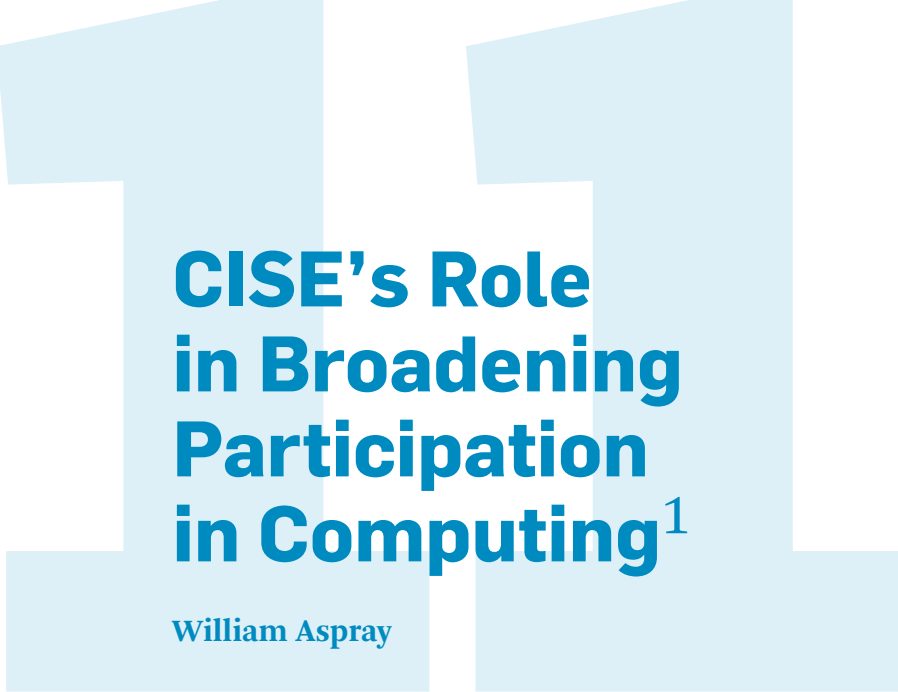
Director and CISE. The fluctuating interest is between *advocacy* for more service in the high-performance computing arena (which they believe the Director will ensure) and a focus on *utilization* of what is available (which leads the Director to move it back to CISE). This is a pattern that has been repeated since almost the founding of NSF and will continue as long as technology continues to improve.

Notes

1. “From Supercomputing to the TeraGrid,” National Science Foundation Online Fact Sheet, provides a shorter version of this chapter, especially for the period 1950–1980. https://www.nsf.gov/news/news_summ.jsp?cntn_id=106875; last accessed 15 February 2019.
2. This mirrors an enduring tension in the research community itself. It is often the basic sciences that generate the most complex computational/communication problems that often exceed the state of the art and thus provide new goals for computer scientists. Yet computer scientists do not want to be seen only as programmers and computer center operators; they aspire to be basic scientists who are pursuing entirely unrelated questions of computation and the structures that implement entirely new forms of computation.
3. See Chapters 1, 3, and 9 in this volume for additional information on HPC.
4. V. Bush. July 1945. “As We May Think.” *The Atlantic*, pp. 101–108. <https://www.theatlantic.com/magazine/archive/1945/07/as-we-may-think/303881/>; last accessed 1 October 2019.
5. National Science Foundation. 1951. *Annual Report of the National Science Foundation, 1950–51*. U.S. Government Printing Office. This report contains a short but excellent history of basic research and its increasing importance to society.
6. S. Strogatz. March 4, 2003. “The Real Scientific Hero of 1953.” <https://www.nytimes.com/2003/03/04/opinion/the-real-scientific-hero-of-1953.html>; last accessed 1 October 2019. One of the authors of the paper described in this article was John Pasta, later the head of computer science research funding for NSF.
7. Director of the Meteorological Office. July 1954. “Conference on High Speed Computing,” in *The Meteorological Magazine*, 83(985): 193–224.
8. For example, in its report to the National Science Board (Appendix VII in the 1953 NSF Annual Report), the Advisory Committee on Minerals Research lists “high-speed computers” as an important piece of instrumentation for Fundamental Geophysical Research. Although not NSF-supported, the Fermi-Pasta-Ulam experiment was conducted in 1953, as described in Strogatz, 2003, *op. cit.*
9. National Science Foundation. 1955. *Fifth Annual Report*, p. 55. Charles Babbage Institute, or <https://www.nsf.gov/about/history/annual-reports.jsp>; last accessed 1 October 2019.
10. “Preliminary Summary: Foreseeable Basic Research Facility Needs for MPE Disciplines.” July 16, 1956. Internal document for National Science Foundation use. Charles Babbage Institute.
11. The national security needs of the country were driving advanced development of true HPCs. One of the earliest was the IBM Stretch in the 1950s; one of the chief designers was Erich Bloch, later the NSF Director from 1984 to 1990.

12. Peter Lax. December 26, 1982. *Report of the Panel on Large Scale Computing in Science and Engineering*, under sponsorship of National Science Foundation and Department of Defense. Charles Babbage Institute.
13. NSF Working Group on Computers for Research, Kent K. Curtis, Chairman, under the direction of Marcel Bardon. 1983. *A National Computing Environment for Academic Research*. Internal report, National Science Foundation: Washington, DC. Charles Babbage Institute.
14. Rosalie Steier. July 1984. NSF takes the initiative. *Communications of the ACM*, 27(6): 528–618. Charles Babbage Institute.
15. Bloch had been a key figure in the design of one of the first supercomputers at IBM and then the design of IBM's System 360 family of computers. In addition to his personal expertise, at IBM he had been associated with a number of world-class scientists in other fields and thus had an appreciation for the computational challenges they faced.
16. Erich Bloch. September 11, 1984. "COSEPUP Panel on Computer Architecture," draft internal document. Charles Babbage Institute.
17. S. Karin. 1990. The evolving supercomputer environment at the San Diego Supercomputer Center. In: J. S. Kowalik, ed., *Supercomputing*. NATO ASI Series (Series F: Computer and Systems Sciences), vol. 62. Berlin, Heidelberg: Springer. "From Supercomputing to the TeraGrid," *op. cit.*, gives another overview. For more insights into the creation of the first NSF supercomputer centers, see: Oral history, Charles Brownstein, interviewed by William Aspray, June 23, 2017. Charles Babbage Institute; and Oral history, Rick Adrion, interviewed by William Aspray, March 14, 2017. Charles Babbage Institute.
18. The NSF-supported supercomputer at the National Center for Atmospheric Research (NCAR) was initially sometimes listed as one of the NSF supercomputers, but it was never a part of the supercomputer program and is supported out of the Geophysical Sciences Directorate.
19. T. Misa. 2013. *Digital State: The Story of Minnesota's Computing Industry*. Minneapolis: University of Minnesota Press, p. 320.
20. Oral history, Bill Wulf, interviewed by William Aspray, July 28, 2017. Charles Babbage Institute.
21. NSF Blue Ribbon Panel on High Performance Computing. 1993. *From Desktop to Teraflop: Exploiting the U.S. Lead in High Performance Computing*. Washington, DC: National Science Foundation. The panel was staffed by Nico Habermann, AD/CISE, and was his last activity before his untimely death; the report is dedicated to him.
22. E. F. Hayes et al. September 15, 1995. "Report of the Task Force on the Future of the NSF Super Computer Centers Program." <https://www.nsf.gov/pubs/1996/nsf9646/nsf9646.htm>; last accessed 1 October 2019.
23. Numerous studies have documented this fact, sometimes implicitly in the form of an acknowledgement in a scientific paper. The annual supplements to the President's Budget starting in 1991 (for FY 1992), especially in the early years, illustrate some of these discoveries; see <https://www.nitrd.gov/publications/index.aspx>; last accessed 1 October 2019.
24. See Chapter 9 in this volume.

25. SRC Timeline—Semiconductor Research Corporation, <https://www.src.org/src/story/timeline/>; last accessed 27 March 2019.
26. OSTP. November 20, 1987. *A Research and Development Strategy for High Performance Computing*. Charles Babbage Institute.
27. Bell's papers at the Computer History Museum, especially VBB3, provide additional information. See also: Oral history, Gordon Bell, interviewed by William Aspray, July 14, 2017. Charles Babbage Institute; and Brownstein interview, 2017, *op. cit.*
28. OSTP. Executive Office of the President. September 8, 1989. *The Federal High Performance Computing Program*. Charles Babbage Institute.
29. See Chapter 3 in this volume.
30. For example, see *From Desktop to Teraflop*, 1993, *op. cit.*; Committee to Study High Performance Computing and Communications, 1995, *Evolving the High Performance Computing and Communications Initiative to Support the Nation's Information Infrastructure*, National Academy Press.
31. Hayes et al., *op. cit.*
32. NSF. April 2001. "The PACI Program—Nifty 50." <https://www.nsf.gov/about/history/nifty50/paci.jsp>; last accessed 1 October 2019; and P. Young. April 9, 1997. "Testimony Before the House Science Committee." <https://www.nsf.gov/about/congress/105/paciyou.jsp>; last accessed 1 October 2019.
33. President's Information Technology Advisory Committee. February 1999. *Information Technology Research: Investing in Our Future*, PITAC Report to the President, https://www.nitrd.gov/pitac/report/pitac_report.pdf; last accessed 1 October 2019. Charles Babbage Institute.
34. In measuring the capability of a computer for scientific calculations, "floating point operations per second," or "flop," is often used. A "terascale computer" is thus a computer whose capability is in the range of trillions of flops. <https://www.nsf.gov/od/lpa/news/press/00/pr0053.htm>; last accessed 1 October 2019. On ITR, see Chapter 8 in this book.
35. NSF. 2001. "Distributed Terascale Facility (DTF)." <https://www.nsf.gov/pubs/2001/nsf0151/nsf0151.htm>; last accessed 1 October 2019.
36. NSF. August 9, 2001. "Distributed Terascale Facility to Commence with \$53 Million NSF Award." <https://www.nsf.gov/od/lpa/news/press/01/pr0167.htm>; last accessed 1 October 2019.
37. NSF. ca. 2005. "From Supercomputing to the TeraGrid." https://www.nsf.gov/news/special_reports/cyber/fromsctotg.jsp; last accessed 1 October 2019.
38. Bajcsy interview, 2017, *op. cit.*
39. The term *cyberinfrastructure* is used to denote a system that embodies the state-of-the-art in computing, communications, and digital information environments. While somewhat similar terms may have been used earlier, Bajcsy's use of the term brought it to prominence.
40. D. E. Atkins et al. January 2003. "Revolutionizing Science and Engineering Through Cyberinfrastructure." <https://www.nsf.gov/cise/sci/reports/atkins.pdf>; last accessed 1 October 2019.



CISE's Role in Broadening Participation in Computing¹

William Aspray

The National Science Foundation's charter directs it to ensure adequate human resources to carry out the nation's scientific research and education. Thus, from the creation of the Foundation in 1950, there has been an interest in having a sufficiently broad pool of scientists and science educators.

Indeed, there is a close affiliation between educational goals and human resource goals within the Foundation, and education and human resources are housed within the same NSF directorate. The big initiative taken by the Foundation in this general area in the 1950s was the response to Sputnik with a massive program of fellowships and traineeships. The Foundation also took great interest, increasing in the 1960s and 1970s, in helping to build capability to teach computer science.²

Yet because historically there was a strong identification of being a scientist with being white and male, much of the concern in the culture at large about "broadening participation" in science was skewed for many years toward providing a technical education to a sufficiently large pool of white males.

Substantial change came as a result of the women's rights and civil rights movements of the 1960s and 1970s, culminating in the Foundation's direct response to the Science and Technology Equal Opportunity Act of 1980. During the 1980s, the Foundation supported projects across the directorates concerning broadening participation. However, the Foundation found that these earnest efforts were

Table 11.1 NSF programs to broaden participation in science generally

1974	Women in Science Program
1978	Minority Graduate Fellowship Program
1987	Centers of Research Excellence in Science and Technology (CREST)
1990	Minority Postdoctoral Research Fellowship (MPRF) Program
1991	Louis Stokes Alliances for Minority Participation (LSAMP)
1992	Program for Persons with Disabilities (PPD)
1993	Program for Women and Girls [renamed Gender Diversity in Science, Technology, Engineering, and Mathematics in 1999]
1993	Urban Systemic Initiatives (USI) Program
1994	Rural Systems Initiative (RSI) Program
1994	Research in Disabilities Education (RDE) Program
1994	Facilitation Awards for Scientists and Engineers with Disabilities (FASED) Program
1996	Presidential Awards for Excellence in Science, Mathematics, and Engineering Mentoring (PAESMEM)
1997	Professional Opportunities for Women in Research and Education (POWRE) Program
1998	Historically Black Colleges and Universities Undergraduate Program (HBCU-UP)
1998	Alliances for Graduate Education and the Professoriate (AGEP) Program
1999	ADVANCE
2005	Research on Gender in Science and Engineering Program

not particularly effective at broadening participation of women or other under-represented groups. As a result, the Foundation rethought its efforts. Table 11.1 lists a number of Foundation-wide programs.

The early records are not complete, but the earliest broadening participation grant we found in the computing field was awarded in the late 1970s to the University of Texas at Austin's Nell Dale, who established a re-entry program for women in computer science under a grant from the NSF Women in Science program. Early efforts were also made through the Supercomputer Centers program (1985–1997), and these broadening participation efforts were continued in the Foundation's

follow-on Partnerships for Advanced Computational Infrastructure (1997–2004), Terascale Initiatives (2000–2004), and TeraGrid (2005–2010).

Two of the Foundation's programs for broadening participation in science and technology generally were particularly influential on CISE. The Program for Women and Girls, directed by Ruta Sevo, focused on funding replicable projects (known as Model Projects), disseminating scientific results about underrepresentation, identifying best practices, and conducting research (even with only modest funding). This program was aimed at K–12 and undergraduate education. While some computing projects were funded, including ones on the educational value of “pair programming” to teach girls to program, and teaching design principles to young women through game design, the most important impact on CISE was to build a community of researchers interested in underrepresentation. They quickly populated CISE's early research-oriented grant programs on broadening participation in computing.

The other influential early Foundation program was ADVANCE, established in 2001. Its director, Alice Hogan, effected organizational change in policies and practices to make the college environment more welcoming to women. ADVANCE also built a community of practice among college administrators who disseminated best practices. Caroline Wardle, who led CISE's IT Workforce program, and later Jan Cuny, who led CISE's Broadening Participation in Computing program, each had close ties to ADVANCE.

Three major CISE programs have specifically targeted broadening participation in computing: Information Technology Workforce (ITWF), Broadening Participation in Computing (BPC), and Computing Education for the Twenty-First Century (CE21). By the late 1990s, there was a widespread call for more IT workers in the United States, driven originally by the Y2K problem and amplified by the dot-com boom. The Presidential Information Technology Advisory Committee (PITAC) in 1999 called for broadened participation in IT careers.³ That same year, six professional computing societies together released their own study, which discussed the place of underrepresented groups in meeting the workforce demand.⁴

In response, CISE held a virtual workshop involving 234 participants to identify the causes of underrepresentation and create a research program that CISE might support. CISE also responded by creating the IT Workforce program, which ran from 2000 until 2004. At first, the program only supported research. Toward its end, Peter Freeman arrived as the CISE AD, and he emphasized implementation of research that had already been funded. An annual principal investigator meeting contributed substantially to building a research community and disseminating

research results. (Cohoon and Aspray's *Women and Information Technology* (2006) analyzes a number of these ITWF studies.⁵) Research on girls included recruitment of middle school girls into computing, the gendered high school curriculum, and study of the impact of race on computing career decisions. Research on post-secondary education examined retention of women in computer science degree programs, comparison of computer science and management information science programs, and critique of the commonly used "pipeline" metaphor for conceptualizing the departure of women from the computing field. Research on IT careers studied short-term job training programs for low-income women and the impact of gender on professional commitment of women entering the IT field. The ITWF research expanded under the auspices of the Information Technology Research (ITR) program funded by Congress as a direct response to the PITAC report.⁶ One notable grant supported with ITWF and ITR funds was to the ATLAS Institute at the University of Colorado Boulder, which led to the founding of the National Center for Women & Information Technology, also supported by CISE AD Peter Freeman.

In 2003, Freeman convinced Jan Cuny to lead CISE's efforts in increasing diversity in computing. From the University of Oregon, she had already been active in the Computing Research Association's well-regarded Committee on the Status of Women in Computing Research. When she arrived at the Foundation, Freeman and his Deputy Director Deborah Crawford provided an ample budget and ran political interference, which enabled Cuny to bring about extraordinary changes throughout Freeman's four-year term and continuing for another decade. The first BPC grants were awarded in 2006. This program focused on funding alliances:

. . . the CISE BP Initiative will focus on broad alliances (of academia, K-12 outreach, industry, and community-based organizations) across and within targeted groups to address issues spanning wide regions of space . . . the individual groups retain their identity and can continue to focus efforts on issues and challenges unique to their community, while at the same time they can come together to leverage work on common issues . . . ⁷

The BPC Alliances were constructed around four goals (see below), and the Alliances directed their efforts at making institutional change to educational institutions. The alliances often resulted from community contacts assembled during an ITWF program. Some alliances focused on a single population such as women (NCWIT), African-Americans (iAAMCS), or Hispanics (CAHSI), while others spanned geographical regions, such as Georgia Computes, the southeastern United States (STARS), and Massachusetts (CAITE).⁸ BPC funded more than 30 projects—most of them large. The 11 alliances that existed as of 2009 are shown in Table 11.2.

Table 11.2 Broadening participation in computing alliances as of 2009

Alliance Name	Principal Investigator	Lead Institution
A4RC	Gerry Dozier	North Carolina A&T
AccessComputing	Richard Ladner	U. Washington
CAHSI	Ann Gates	U. Texas El Paso
CRA-W/CDC Coalition	Lori Clarke	CRA-W, CDC
STARS	Teresa Dahlberg	U. North Carolina Charlotte
ARTSI	Andrew Williams	Spelman College
CAITE	Rick Adrion	U. Massachusetts Amherst
EL	Richard Tapia	Rice U.
Georgia Computes	Mark Guzdial	Georgia Tech
Into the Loop	Jane Margolis	UCLA
NCWIT	Lucy Sanders	NCWIT

Alliance participants (1) develop and implement interventions that support students and early career faculty, (2) create sustainable changes in culture and practices at the institutional, departmental, and organizational levels, (3) serve as models and contribute to repositories for effective practices to broaden participation, and (4) leverage the work of existing BP efforts and other Alliances.⁹

Table 11.3 lists activities carried out by the BPC Alliances.

While BPC was underway, CISE created Pathways to Revitalized Undergraduate Computing Education (CPATH) under the direction of Caroline Wardle and Harriet Taylor. Aiming to transform undergraduate computing education nationally, CPATH interfaced with BPC by providing a path to an undergraduate education in computing (and ultimately into a computing career) for high-school students who became excited about computing in BPC engagement programs. In 2010, BPC and CPATH merged into Computing Education for the Twenty-First Century (CE21). CE21 was intended to build on what had been learned by the BPC Alliances and be open to all students. The most important element of CE21 is the CS10K initiative to place 10,000 high-school teachers able to teach a rigorous, modern computing course.

In addition to bringing modern computing to high schools, CS10K introduced similar courses at the beginning college level. CISE's Jan Cuny carried this complex and ambitious project out in partnership with a number of organizations: the ACM Education Policy Committee (led by Robert Schnabel and Cameron Wilson), the

Table 11.3 Sample programs carried out by the BPC alliances

AccessComputing	Workshops for students with disabilities to learn about computing and encourage them to study it and prepare for computing careers
CAHSI	Mentor-Grad program to encourage undergraduate Hispanic students to apply to graduate school in computer science and train for careers as professors
Georgia Computes	Summer computing camps
CAITE	Helping community college students bridge to four-year degrees in computing
A4RC	School-year and summer research experiences to encourage African-American undergraduates to pursue doctoral degrees in computing
CRA-W/CDC Alliance	Discipline-specific mentoring workshops for women and minority doctoral students and recent postgraduate
STARS	Annual conference

Computer Science Teachers Association (led by Chris Stephenson), the College Board (led by Trevor Packer), and the start-up firm Code.org (initiated by one of the company's co-founders, Hadi Partovi, and carried out by Cameron Wilson who relocated from ACM).

At the time, most public high schools did not have substantive, modern courses in computer science; many had none at all. Some had keyboarding courses or instruction in off-the-shelf software. Strong computing courses were typically taught as an elective to only a few students, primarily in wealthy schools and most often to white male students. This high school reform involved many disparate tasks: creating new courses, developing curricula, training teachers, convincing high schools to offer these courses when computer science was not typically a part of the college prep curriculum, changing educational policies of local school districts and state education boards, and providing professional support communities for the teachers.

NSF's main focus was on creating new introductory computer science courses, intended originally for high school use and also used at the lower level in college. There were two major courses developed out of these efforts. One was Computer Science Principles, written by a group led by Owen Astrachan of Duke University and Amy Briggs of Middlebury College. They distilled the great concepts from computer science, as well as computer science ways of thinking, in a single year-long

Table 11.4 Organizations providing informal computer science education (sample)

Entrepreneurial efforts to teach children computer science:

GoldieBlox
 Black Girls Code
 CodeEd
 Iridescent
 Digigirlz
 CodeNow
 CoderDojo
 Girls Who Code

Entrepreneurial efforts to teach female college students and adult women computer science:

Geek Girl
 Girl Develop It
 PyLadies
 she++

high school course. Jane Margolis, who led one of the BPC Alliances, found that many of the students in Los Angeles public schools were not prepared for Computer Science Principles. Thus, in 2008 she and colleagues developed a curriculum more accessible to a diverse population of high-school students, entitled Exploring Computer Science. These courses or variants on them have been introduced in many high schools and colleges across the country. The cumulative efforts of NSF and partners succeeded in getting modern computer science taught to thousands of high-school students and college freshmen. These curricular projects also found favor with the CISE AD Jeanette Wing, who was keen on disseminating computational thinking to the general public.

The CS10K effort was designed to teach computer science through formal organizations such as public high schools and colleges. Yet simultaneously a movement arose to teach computer science in informal settings, often reaching those who are underrepresented in the computing field. Table 11.4 lists some of these organizations.

There are hundreds of these informal education efforts, and they are typically small, entrepreneurial, and geographically limited (although a few grew to be national in scope). CISE did play a critical role in the Computer Science Collaboration Project (CSCP), a large national organization with similar goals. CSCP built collaboration between the organizations and alliances that were participating in the BPC program. Focused on K-12, CSCP embodied principles developed by an earlier organization, the National Girls Collaborative Project (NGCP), organized by Karen Peterson and Brenda Britsch.

NGCP carries its work out through locally based Collaboratives. It supports the efforts of these Collaboratives through mini-grants to girl-focused STEM programs, professional training webinars and in-person seminars, the preparation and distribution of statistics about girls and STEM, and the construction of repositories of success stories and exemplary practices with relevant links to other resources and organizations.¹⁰

By 2014, NGCP worked with 12,800 organizations in 39 states and engaged more than eight million girls and four million boys. NGCP followed four principles: being student centered, connecting science learning to students' futures, facilitating a science identity among students, and encouraging a mindset that mental growth can be achieved through learning (disabusing students of the notion that these talents are innate). NSF funding was critical to both NGCP and CSCP. CISE funding enabled mini-grants to such organizations as the Girlstart Game Development program¹¹ in Austin, Texas, and the Learning Computer Science through the Lens of Culture and Science enrichment academy in Yucaipa, California.¹²

While there continues to be much work to do, NSF has supported research and implementation projects and worked hard to enhance both formal and educational programs that will enable the United States to broaden participation in computing.

Conclusions

In the 1980s and 1990s, CISE learned from Foundation-wide programs intended to broaden participation, especially ADVANCE and the Program for Women and Girls. Since then, CISE has been a leader in the Foundation, with its approaches of broadening participation in computing adopted widely across NSF. The BPC program has relied on research-based practice to be both effective and efficient. Best practices are spread by alliances involving multiple institutions. The BPC alliances are well placed, then, to succeed in their efforts.

Affiliated with the work on broadening participation have been efforts to reform high-school and college education in computer science. NSF has taken a

leading role in the effort to create new curricula that teach the fundamentals of computer science and to get these placed in a wide range of high schools and colleges, reaching a wide demographic profile of students. The bold effort to create 10,000 high-school teachers who are prepared and are teaching a rigorous computer science curriculum is well on its way to being achieved.

Notes

1. This paper is drawn from: W. Aspray. 2016. *Participation in Computing: The National Science Foundation's Expansionary Programs*. Cham, Switzerland: Springer.
2. See Chapter 6 for a discussion of education and facilities in the era prior to the formation of CISE.
3. Presidential Information Technology Advisory Committee (PITAC). 1999. *Information Technology Research: Investing in Our Future*. National Coordination Office for Computing, Information and Communications.
4. P. Freeman and W. Aspray. 1999. *The Supply of Information Technology Workers in the United States*. Washington, DC: Computing Research Association. Supported in part by NSF Grant No. EIA-98-12240.
5. J. Cohoon and W. Aspray. 2006. *Women and Information Technology: Research on Underrepresentation*. Cambridge, MA: MIT Press.
6. For more on the ITR program, see Chapter 8 of this volume.
7. P. Freeman and J. Cuny. 2005. Common ground: A diverse CS community benefits all of us. *Computing Research News*, 17(1).
8. D. E. Chubin and R. Y. Johnson. 2011. *Telling the Stories of the BPC Alliances: How One NSF Program Is Changing the Face of Computing*, Center for Advancing Science & Engineering Capacity, American Association for the Advancement of Science; "Broadening participation: A Program Greater than the Sum of Its Parts: The BPC Alliances." *Communications of the ACM*, 54(3): 35–37.
9. National Science Foundation. 2009. *Broadening Participation in Computing (BPC)*. http://www.nsf.gov/funding/pgm_summ.jsp?pims_id=13510; last accessed 20 October 2014.
10. W. Aspray. 2016. *Participation in Computing: The National Science Foundation's Expansionary Programs*. Cham, Switzerland: Springer.
11. <http://ngcproject.org/project-summary-girlstarts-game-development-series-girlstart-after-school-participants>.
12. <http://www.cscproject.org/project-summary-learning-computer-science-through-lens-culture-and-society>.



What Does an AD/CISE Do?

Peter A. Freeman

At the start of my fifth year as Assistant Director (AD) of the Computer and Information Science and Engineering (CISE) Directorate, I wrote a memo entitled “Reflections on the Position of CISE AD.” Later that year, I wrote a memo to my successor, not yet named, entitled “Issues, Lessons Learned, and Hints for the Future.” These two memos are reprinted in their entirety below with only minor editorial changes. They provide a contemporaneous account of some lessons learned. They may be of use to future ADs, or candidates, as well.

TO: Whomever Is Interested
FROM: Peter A. Freeman
RE: Reflections on the Position of CISE AD
DATE: April 27, 2006

Preface: This is written with no intention to self-aggrandize, boast, or otherwise comment—positively or negatively—on my own performance in the position the past four years. My objective is only to inform those trying to select among candidates as well as to inform the candidates themselves and anyone else who is interested about the position.

Generalities: For any who are not closely familiar with the position, it is primarily an executive leadership position, not a management position and not a scientific position—although it obviously includes some strong elements of both. The two primary functions of the job deal with strategy (in all its aspects) and

communications (in multiple ways). Those two functions probably consume (in one way or another) 80% of the AD's time—or they should if one is to be successful.

Given the nature of NSF and the CISE portfolio, the AD/CISE clearly should have a strong reputation in computer science while having knowledge of and interest in the other related fields that CISE supports—computer engineering, information systems, and so on. I emphasize strength in computer science not only because it provides the intellectual basis for all computing-related activities but because that is where the majority of the funding goes. Similarly, while the AD need not have come from academia, lack of direct and fairly recent and significant experience in academia would put the AD at a great disadvantage. At the same time, solid familiarity with the processes and expectations of a big organization is also essential.

It is a position that requires the ability to make decisions. While in many ways NSF operates like a collegial faculty in which everyone is equal, there is, in fact, a hierarchy that anyone who is at NSF for more than a brief period pays attention to. As the top of a major piece of the hierarchy, the AD must pay attention to that fact and its implications. There are also a number of implications that reflect on the choice of an AD—judgment, fairness, balance, consistency, and having a strong basis on which to make decisions.

If one looks at a book on leadership—my favorite is John Gardner's *On Leadership*¹—you typically find a list of the characteristics of a leader. This position fits such a definition well. You also find, of course, that different people have different styles of leadership and different balances among the various characteristics. That's fine, and the CISE ADs to date reflect that diversity—but lack of at least some strength in one or more of the characteristics of a leader could lead to far less than strong overall performance.

While it is not a management position in the sense of managing a team of software developers or even that of a department chairman who must deal with many operational problems, a good sense of organizations and management is critical. While others (the Deputy AD, Division Directors, other staff) have the primary responsibility for running the operations, the AD is responsible for setting objectives for them, making sure that they are doing their jobs, dealing with extraordinary situations, and making decisions on major issues that are brought to the AD for a final decision.

While not a scientific position, it requires being able in some cases to make decisions based on scientific judgments. It also requires some familiarity with

and appreciation for a wide range of scientific and engineering subjects. And, of course, the strategic directions for CISE that the AD is expected to set should be well-informed and reflective of the field as well as his or her own personal opinions. In short, the AD should be a scientific as well as an organizational leader.

Much of the work on strategy by the AD is expressed in the end by policies that he or she helps formulate and then enforce. While this is usually a team effort and there are policy experts on staff—indeed, NSF has a “Policy Office”—familiarity with and appreciation for policy and policy formulation is a valuable bit of experience on which to draw. In particular, understanding that public policy that affects thousands of people over a number of years can only be made or changed very carefully is important—and probably where someone whose experience has all been in the world of industry would have the toughest time.

Some other important aspects: At present and for the foreseeable future, the CISE AD is the single most visible person in Washington providing funding for CS research. This carries with it a level of responsibility for representing the entire CS research community in a number of ways. One operational way in which this shows up is that historically the AD/CISE chairs the primary interagency committee (NITRD) for coordinating all federally supported computing research. This is a fairly visible position to the Administration and to the Congress, and requires a good level of diplomacy and statesmanship (the AD/CISE is usually the highest ranking official on the committee by several levels of responsibility).

One of the implicit, but very important responsibilities of the job is to know what is going on in the rest of the world in computer science research and education as well as the broader aspect of computing. The AD/CISE is routinely turned to by the Director of NSF as the in-house expert on what is going on in other countries; I often get asked about offshoring, for example. More substantively, an important aspect of setting direction for CISE is to make sure the U.S. academic CS community is participating in the appropriate foreign venues and issues. An interest in and knowledge of international activities in the field is very important.

Personal characteristics: The position requires stamina, dedication, patience, and forbearance. The hours can be long and one is well advised to not go to meetings without doing one’s homework. Missing deadlines can have real consequences for the field if an opportunity is thus missed.

A good deal of travel is involved. While in many ways you are the “boss” and get to set your own agenda, in many other ways neither your time nor your scope of action is your own—it is the government, after all, and you most definitely have a

boss—actually, many bosses when you consider the various rules and procedures and independent entities that can cause you to drop everything to focus on what they are asking for (e.g., Congress).

The “gold standard” of NSF is objective peer review, of course. That ethos extends to almost all activities and situations at NSF—one of the things that makes it a great place to work. The AD must apply this principle scrupulously in all situations because of his or her visibility and the fact that decisions made by the AD typically have broad impact.

There are different styles, of course, but some degree of openness is important. The AD is the head of an organization of 100 people; so just on a daily basis, being open to communication, to new ideas, to feedback is important. More publicly, being open to the press, to the public, and to the research and education community is important. At the same time, the AD must be much more circumspect both in word and action, especially about some matters, than an academic is used to!

Communication ability is essential. Being open and being willing to listen is the passive, receptive side of communication. At the same time, the AD must be able to write and speak well in as many situations as possible—one-on-one, in small groups, to peers, those reporting to him or her, to superiors, the public, etc.

The AD is very much part of a team, so being willing and able to work in a variety of teamwork situations is essential—not only to get along but to be able to press the case for CISE and make things work.

All of the above should have already implied my last point—good judgment is essential. The AD must make decisions on a wide variety of things—choice of personnel, whether to answer a particular question or not from the press, which alternative to choose in a situation, whether to speak up about something in a meeting or let it pass and save one’s ammunition for a more important issue, and on and on.

There are undoubtedly other important factors that I have overlooked in this stream of consciousness. And there are certainly long litanies of do’s, don’ts, and suggestions that I will expect to share with the new AD—but those can wait and are, hopefully, simply instantiations of the points made above. NSF is a great place to work, in an exciting and culturally rich city, and the position of AD/CISE affords one an unparalleled opportunity to serve our field and our country.

TO: My successor
FROM: Peter A. Freeman
RE: Issues, Lessons Learned, and Hints for the Future
DATE: December 26, 2006

Preface: This is written when you have not yet been named. My objective is to inform you and help you to stand on my shoulders, learn from my mistakes, and in general to do the very best job you can for our community. I fully recognize and expect that you will have different specific objectives than I had when I came in and when I left. That is good. On the other hand, I assume—and our community expects—that you will always strive to do what is best for the entire community. To the extent that some good things have been continued, started up, or at least plans laid during my tenure, I sincerely hope that you will see fit to improve and build on them.

I trust that you have read my memo of April 27, 2006 (“Reflections on the Position of CISE AD”). I hope that you will reread it—often! There are many lessons to be learned there that I will try not to repeat here.

Communities: I have consistently taken the position that computer science is the primary community we serve. I believe strongly in a broad definition of CS that includes the majority of what CISE supports and is responsible for. At the same time, it is important to recognize that we also support computer engineering, computational science, information theory, communication theory, DSP, information systems, social informatics, and probably several other important, but smaller, communities not mentioned here.

As a result, I often talk about the “computing community.” At the same time, I take every opportunity I can find—especially with our AD colleagues and OD—to remind them that CS is the “mother lode” and “intellectual core” for all of the computing disciplines.

I have found that it is also very important to frequently remind people that CS is a science in its own right, and that CS \neq computational science \neq cyberinfrastructure \neq supercomputing.

As a young, and mostly applied, discipline, we don’t always have a very distinguished image within NSF or in the larger scientific community—I’m sure that you are as aware of that as I am. I point this out, however, because you have an opportunity and a responsibility on behalf of all of us to work to improve that image. I

have tried and hope that you do as well to represent the deeper parts of our field consistently and frequently to my AD colleagues and others within NSF.

At the same time, we need to do more to build the scientific reputation of computer science within academe and the broader society. I will be continuing to work toward that goal as an ex-AD, and I hope that you can provide the leadership needed from the position of AD.

The ADs in my time at NSF have largely formed a collegial and cooperative team, working with each other for the good of science and NSF, often thinking of those objectives first and the objectives of their discipline only secondarily. Our predecessors have not always been known as team players who could be trusted and that fact, coupled with the often prevailing view of CS as just a service discipline, has not served CISE nor the field well. I set out to reverse that view and think I have made good progress. I encourage you to pay careful attention to the community of which you are now a part.

The last community that I direct your attention to is the CISE Advisory Committee. Used properly, it can be an effective advocate for CISE and a valuable source of information and feedback for you personally and for CISE. Make them feel like a community, listen to them, and let them know that you've heard them and value their inputs even if you don't agree with all of them. Choose a broad range of members who truly represent our community with people who can provide good advice that you respect. Keep the discussions with them at a strategic level, which you can easily do by posing the right questions and steering the discussions as needed. Your best ally in this will be the AC Chair, whom you choose. He or she is your partner. Al Aho has been great, and I recommend you keep him for at least your first year.

External issues: You were chosen because you represent the community broadly and bring a perspective that will be useful to NSF. Thus, there is little point in my spending much time explicating my own views since, after almost five years, they are embedded in the programs and initiatives at CISE you are inheriting. I will, however, touch on a few points.

(Let me note parenthetically that I hope you are planning on staying the maximum amount of time allowed. Less than three or four years will not give you the time to have much of a positive impact.)

There is a meta-issue that I believe is *the single most important issue* for our field: Helping the field mature.

The efforts that have been started under my guidance that, I hope, will be the most relevant to addressing this issue are the CCC, GENI, and, at least as far as the CISE part of it, the Computational Discovery and Innovation initiative announced in the 08 budget. I hope that you will continue those, while improving them as needed.

A very important, long-range, lower-profile activity that I encourage you to work on at every possible turn is the development of new leaders in our field. Compared to most scientific disciplines, CS has very few leaders at all levels in all aspects who can serve us well in Washington and in other important venues. The AC is one place to do this. Recommending people from our field for other positions, committees, task forces, and so on is another. You can address this in the speeches you give and the articles you write. You have a pulpit that can be used effectively.

A related issue is the development of thought leaders in science and broader fields as well who are also computer scientists. Be on the lookout constantly as you travel around for those who can be encouraged in this area and use your position to promote them.

A similar meta-issue is turning CISE into the organization that is doing the most for advancing our field in exciting ways (this also involves internal issues). With the decline in influence (and funding) of DARPA, we have not only the opportunity but also the responsibility to do this. The ITR Initiative laid the groundwork and showed that our community can work together on (modestly) large projects of their own initiative. GENI is clearly aimed in this direction. The CCC was chartered precisely to enable the community to envision other such projects. The large-projects program announced in the 08 budget is intended to provide the kind of support that DARPA used to provide—but done in the NSF style.

There is a continuum of funding styles ranging (in caricature) from a pure NSF model to a pure DARPA model. Neither caricature is what you actually find in practice, of course. I have tried to push/pull CISE more toward what I call a “modified DARPA approach” in which we listen to the community and always utilize peer review to inform our decisions, but at the same time articulate clear directions (e.g., GENI) and initiate new programs and seek innovative proposals to help move things along.

One of the lessons I learned (more slowly than I would have liked) is that one has quite a lot of freedom of action at NSF—a lot more than is usually recognized. Another lesson learned is that unsolicited proposals are always welcome. This permits those with creative ideas to be encouraged and supported outside the

framework of specific programs. Planning grants and SGERs (Small Grants for Exploratory Research) also provide the mechanisms to help realize new directions.

I don't need to tell you that software is important and that our ability to design, implement, measure (and predict performance), and modify software—to say nothing of making it secure, robust, etc.—still faces serious challenges. My view is that we are nowhere near having a robust, underlying body of scientifically validated knowledge on which to design software (the structures) and create it (the software engineering processes). In many ways, I believe that is still the biggest, most pervasive issue facing our field. The Science of Design (SOD) program was my (feeble) attempt to start to lay the groundwork for such a body of knowledge. It's too early to tell just how successful it will be (my own estimate was that it would take 5 years minimum to start to see any results—on the other hand, we've been saying for 40+ years that we need something like that and aren't much closer today). The encouragement we gave to Alan Kay is another attempt to try something other than just more of the same. This will continue to be a major issue and I encourage you to try your hand at doing something significant about software.

Chances are that “social informatics” is not a topic that you are terribly familiar with (few of our colleagues are). I happened to be before I came to NSF because of my personal proclivities and the fact that I spent the first part of my career at Irvine (home of much of the early social informatics work). While the scientific study of the personal, social, and organizational impacts of computing (= social informatics, roughly) is generally considered “soft” and is not pursued by many serious computer scientists, in the larger picture of helping our field mature, gain respect as something more than PC and supercomputer programmers, and get the resources that are needed for our research and education programs, it is extremely important. The returns are perhaps the highest per dollar spent compared to any other program. Suzi Iacono is a respected researcher and author in the field and can provide you with any needed background.

Education and workforce (EWF) are topics for which your leadership will be needed. This is a unique responsibility of NSF and one that collectively (considering CISE support for thousands of graduate students) has greater impact than any single research program. It cannot be an add-on to other things. As AD you have the opportunity to set the tone of importance not only for PDs internally, but for the field. CPATH is a start, but it is in its first year (as I write this) and must be continued with appropriate modifications for several years—perhaps five at least. BPC is going well, but as with most programs, evaluation of its true impact is several years away. As long as it appears to be having good, positive impact, it should be continued.

A related issue is that it is very important to have visionary PDs for these programs (perhaps more than for most programs). Jan Cuny has done a great job with BPC, but will be limited by the IPA 4-year rule. We are presently searching for a similar person to put real life into CPATH.

Debbie Crawford, because of her broad experience within NSF, has real insight into what needs to be done programmatically in EWF.

A final note regarding external issues: You will, of course, want to choose what you want to be known for in your stewardship of CISE. If I were continuing and looking to the future, I think I would place a large bet on what is loosely called “educational technology.” I believe the time may be ripe to utilize what we have recently learned about how the brain works, the ever-increasing power of computing hardware, etc. to do something significant in this area.

Internal issues: Don’t ignore them!

One of our predecessors reputedly announced at an All Hands Meeting when they came on board, “I am a scientist with no interest in administration, so please don’t talk with me about such matters.” The result was an accelerated decline in the organizational quality of CISE that still has impact. This, in turn, has direct impact on the quality of the people one can attract to CISE as PDs and admin staff, which in turn has a very serious impact on our community.

Debbie Crawford and I have tried to rebuild the organizational and administrative structure of CISE and turn it into the best-run part of NSF instead of the worst. We’re not there yet, but have made good progress. It is essential that that be continued. Debbie is a master at such things and I will leave it to her to fill in the details, but it will require your support and active involvement.

Let me mention three issues of particular importance: staffing, clusters, and long-range budgeting.

As I write this, the staffing levels of CISE are in crisis. This has been an issue since I came and one that Debbie and I have worked on consistently, yet there has been little relative improvement. We have been putting an extra push on the issue for the past month or two, and it will be the single issue that I am focusing on in my exit interviews with OD. This may have some positive effect, but in our most optimistic moments we do not believe this is going to be solved any time soon. Debbie has developed in-depth analyses and can give you the whole history. In the immediate future, this is perhaps the single most important internal issue for you to work on.

Clusters—of research areas, programmatic elements, and staff—are a mechanism that has been in use in some parts of NSF for years. We introduced them to

CISE in 2003 as part of the overall reorganization. They are a very effective way of gaining more flexibility and breaking down the programmatic stovepipes that otherwise grow up. Some clusters have clicked and worked well. Others haven't or are somewhere in between. This is a function of the leadership that the DDs provide, of the content of specific clusters, and of the individuals involved in the clusters. It is an issue that you need to provide leadership on, but indirectly through the DDs and interactions one on one with leaders among the CISE scientific staff.

Budgeting is something that need not consume a large amount of your time, especially regarding the details, but it is the single most important way of affecting the future of our field. Thus, it is extremely important that the AD pay close attention to the strategic directions that NSF is taking and that you will have an opportunity to help influence. While your first priority should be meeting national needs and the health of science overall, you need to always ask yourself how a proposed direction will impact CISE and our field and then subtly try to steer things in a way that will simultaneously advance science, meet national needs, and help us.

Both for positioning our field in the long term and for making sure that the detailed, yearly allocations within CISE are done in a coherent, constructive manner, it is essential that you take a 3- to 5-year horizon for budget planning. Equally important is having principles and specific objectives that justify and explain your budget planning. As with all things internal, Debbie can be your guide on this.

As with any leader of an organization, one of your most important tasks is the recruitment of good DDs and PDs. That has always been, still is, and always will be true. Don't neglect it.

Getting them here is only the start, however. You must personally guide the DDs and mentor them. And you should interact with PDs as much as you can, to stay close to the science and programs. At the same time, you have to be careful not to go around the DDs. I wish I had done a lot more managing by walking around, attending Division meetings, panels, etc.

One of the most important lessons I have learned is that in spite of appearances (all the scientific staff have Ph.D.s, lots of latitude of action, are as smart as you are, etc. etc.), NSF is NOT a university faculty. There is a delicate balance that must be maintained, but at the end of the day you—and by extension the Deputy and DDs—are ultimately responsible for the important decisions.

While there are definitely things that need to be improved and help that is needed from OD, I believe that overall CISE is running pretty well and I would recommend against any wholesale changes.

There are undoubtedly other important factors that I have overlooked in this stream of consciousness. And there are certainly long litanies of do's, don'ts, and suggestions that I will be glad to share with you. I will be living in Washington for the foreseeable future and while I will do my very best not to meddle in any way (as I have successfully done when leaving previous leadership positions), I would welcome the opportunity to meet with you privately to share what information, knowledge, and insight I have accumulated.

NSF is a great place to work, in an exciting and culturally rich city, and the position of AD/CISE affords one an unparalleled opportunity to serve our field and our country.

Good luck!

These two memos need no further explanation.

At this point in time (2019) and without a full understanding of the current situation in NSF, I would only modify two things in the above memos: First and foremost, to emphasize that the AD/CISE must be a leader. Second, to stress that that leadership must apply to scientific as well as organizational matters.

No one person, of course, can be an expert in all areas of computing and capable of seeing where each area should be headed; but they can seek and be open to the best ideas of those that are experts. The implications of this should permeate their thinking and actions. This prescription applies to other members of the CISE leadership team as well—Deputy ADs, Division Directors, Senior Scientific Advisors, and others as applicable. All must take as their mandate to harvest the best ideas that percolate up, choose among competing ones where applicable, and then take actions to steer their area of responsibility in directions that will help take the field forward.

Notes

1. J. W. Gardner. 1993. *On Leadership*. Simon and Schuster. This reference was not given in my original memo.

PART

**SUMMARY AND
CONCLUSIONS**

13

Summary and Conclusions

Peter A. Freeman, W. Richards Adrion

We have covered a period of almost 70 years, 1950–2016, of the history of computing and NSF. It is worth noting that this period not only covers almost the entire history of NSF, but also most of the history of modern digital computation. The preface to this book provided an overview of its succeeding twelve chapters: the first five chapters provided a chronological narrative of the history and the next seven provided focused case studies of major programmatic initiatives and other topics relating to NSF’s support for computing. Here we have pulled out a few significant events and results to summarize those twelve chapters, identified some themes, and provided a few concluding remarks.

13.1 Summary

One of NSF’s earliest activities, mandated by its “Organic Act,” was to collect and manage scientific information. The Office of Scientific Information (OSI) was one of its earliest organizational units in 1951. OSI supported scientific publication, Soviet-focused projects (machine translation), studies of information processes, indices of publications, and linguistics research related to machine translation, as well as coordinating federal agencies to gather and disseminate scientific information. It also supported work on tools to help in the information-provision process and fundamental research on information retrieval and databases. This was the start of NSF support for fundamental and applied research in computing.

NSF first addressed its responsibility for science education in computing with a workshop in 1954 to train mathematicians on scientific applications of computers. In the 1950s and 1960s, NSF supported early innovators in computer-based education such as Pat Suppes, Dick Atkinson, Don Bittzer, Seymour Pappert, and others.

These were efforts separate from that of OSI, but they contributed indirectly to the educational mission.

Access to computers for scientific research by civilians was severely limited in the 1950s. Thus, some of the earliest funding requests to NSF, as early as 1953, were for assistance in purchasing a computer, and NSF helped fund computing facilities at a few universities. An ad hoc NSF Advisory Panel, led by the polymath John von Neumann, recommended in 1955 that NSF establish a limited program to provide computers and staff for scientific research as well as fund research on the design of advanced computers.

By the late 1960s NSF was reducing its science information that was originally focused on documentation, translation, indexing, retrieval, management, and policy and transferring the activities into the non-research part of the Foundation. At the same time, two parallel, largely independent themes emerged: formal support *for* computing infrastructure to support scientific research, and the growing field of computer science that focused *on* computing systems and theory.

The ad hoc, piecemeal support for computing infrastructure was replaced with institutional grants until the demand became too great in the early 1970s and the program was shut down following a brief attempt to create regional facilities. Provision of computing facilities reappeared in the 1980s with the creation of supercomputer centers and backbone network; we will pick up those developments below.

The computing centers, funded by NSF via institutional grants in the 1950s and 1960s, became *computing research centers* “below-the-radar” as well. Research conducted in these centers laid the basis for computer science as a research discipline, boosted the start of some early academic departments (e.g., Stanford and Purdue), and gave a start to some of the first, future CS doctoral programs. The centers were also home to training programs, which included developing the first courses and curricula in computing and computer science.

In the mid-1960s these trends were being noticed and resulted in the appearance of two influential reports: a study report by J. Barkley Rosser, which highlighted the rapidly expanding use of computers in scientific research and education, and the Pierce Report from the National Academy, which recommended expanding the study of computing by undergraduates.

NSF responded in 1967 by creating the Office of Computing Activities (OCA), providing the first dedicated NSF “home” for computing. The earliest support for computer science research began in the Mathematics Section, so OCA was initially staffed by people moving over from there, led by Milt Rose. Artificial intelligence (AI), pattern recognition, symbolic logic, operating systems, computing theory,

computing security and privacy, human-computer interaction (HCI), and computer graphics, areas originally supported in Mathematics, found a home in OCA and were joined by efforts to support computing facilities, computers in education, and training. These investments, and others such as the use of computers in education, advanced our understanding of many of the basics of today's computing environment. All these efforts were carried out by a few far-seeing program directors and managers at NSF.

Some of these same individuals—for example, Kent Curtis—were later instrumental in forming CISE. John Pasta, a respected senior researcher with a background in mathematics, physics, and computation, came to NSF to head OCA in 1970; he saw it through some turbulent organizational changes and focused it on computing research reflecting the changing nature of computer science. Thus, by 1974 NSF organizationally recognized computer science as a field of *research* rather than as simply a handmaiden to computing facilities and applications of computing, when it created the Division of Computer Research (DCR) as a replacement for OCA.

In the larger picture, computer science and computer engineering emerged on campuses as serious academic subjects worthy of doctoral study at many leading universities, and the broader science community fully realized that computers were essential tools for research. These developments brought pressure to bear on NSF management, who took some consequential actions. An expansion of the NSF science and engineering directorates created a structure fairly close to what exists today, although DCR was “demoted” to a section in a Mathematical and Computer Sciences Division. It was fortuitous, however, that computing was recognized as an area deserving of basic research support prior to the general reorganization that eventually created CISE.

The downgrade to a section and potential loss of some funding was noted and widely lamented by leaders in the computing field. Notable among them was Gordon Bell, a senior leader in industry, who argued in early 1974 that basic research in all areas of computing should be housed in universities, not industry; and that mission agencies such as ARPA (the primary supporter of computing research at that time) could not provide broad enough support to maintain and grow the entire field.

The “Feldman Report” written by a group of distinguished academic and industrial computer scientists led by Jerome Feldman and William Robert “Bert” Sutherland, reporting on a 1978 workshop recommended to NSF that a major infusion of funds be allocated to create 25 well-equipped university laboratories to support computer research. This led to the Coordinated Experimental Research (CER) program in 1980, which over the next ten years greatly expanded the number

of universities capable of experimental CS research at the Ph.D. level. A series of articles to the community, including a famous “eating our seed corn letter” by the president of ACM, capped this broad expression of alarm over the lack of funding for computing research.

At the same time, computers were exploding in power and imploding in cost, prompting NSF to again change strategy, providing workstations and specialized equipment for specific projects in all areas of scientific research. CER funded new computing facilities for computer scientists; this helped to advance research on new architectures, distributed and parallel computing, multi-computers, programming languages, graphics and visualization, computational linguistics, artificial intelligence, and robotics.

As university computing programs began to multiply and grow, a steady stream of research results supported by various funding programs began to appear. In 1979, NSF funded “Theornet” to connect CS theory researchers. This initiative quickly led to a broader CSNET to connect all computer science researchers, which in turn prompted the creation of NSFNET to connect all scientific researchers. It was the popularity and expansion of NSFNET that led directly to the public Internet. These developments were led by smart funding decisions and direct action by a few far-sighted NSF staff, eagerly received by networking researchers.

At the same time, the scientific community was concerned that future advances in their fields would be impeded by the lack of advanced computers. An interagency study, the Lax Report, advocated supercomputing systems for scientific research along a number of dimensions, including networking. Nobel prize-winning physicist Kenneth Wilson coined the phrase “grand challenges” to describe important science problems that could not be attacked with current computing equipment. An internal NSF study on computational science, known as the Bardou-Curtis Report, laid the groundwork for more attention to computational science and supercomputing.

Into this environment came a new Director, Erich Bloch, in 1984. A true technical computer pioneer and IBM senior executive, he knew of the well-documented need for advances in high-end computing and realized that seminal NSF computing research programs were broadening and deepening the field but that these activities were scattered around NSF.

Bloch almost immediately took action to provide high-end computers to the scientific community by buying time on existing resources. Then, in 1986, NSF created five new national supercomputer centers, joining the National Center for Atmospheric Research (NCAR), to address the need for large-scale, dedicated computers for advanced scientific research. The Supercomputer Centers program was

designed to support researchers on dozens of campuses, eventually growing to over one hundred universities affiliated with one or more of the centers. The program also advanced computing technology by creating a new market for supercomputers.

Bloch then turned to the issue of funding for basic research in computing. The result was the creation of CISE, a move that had increasingly been considered internally. It was probably the single most important action that Bloch took to advance the relationship between NSF and computing research and education. Director Bloch and his appointee as CISE AD, Gordon Bell, also from the computer industry, were a potent team to move computing forward inside the Foundation. Chuck Brownstein, Rick Adrion, the Acting DDs, and Jerry Daen helped them to realize their vision.

The structure for CISE lasted largely intact for over a decade and included the supercomputer centers as well as various computing research divisions and a networking research and deployment division. For the first time since computing research had robustly expanded, essentially all of the people at NSF concerned with computing were in one directorate and their AD sat in the high-level meetings concerning funding allocations and other NSF-wide policies regarding NSF fellowships, cost-sharing, and human resources.

Neither Gordon Bell nor his successors during this period—Bill Wulf, Nico Habermann, Paul Young, and Juris Hartmanis—remained in office for more than two years. Experienced CISE staff maintained the momentum and directions envisioned by the ADs throughout this period and beyond. This first decade of CISE was characterized by strengthening CISE's organization, building its reputation within NSF, and creating several significant and far-reaching research and infrastructure programs.

Bell spent his first year as AD on organizational matters and on hiring new program and division directors. He separated the networking infrastructure and research programs from the Supercomputer Centers program to ensure that the center backbone network and burgeoning NSFNET program focused on building a national networking infrastructure, then spent time guiding the research programs on the coming importance of parallel computation, robotics, and VLSI. When he left at the end of 1987, Chuck Brownstein as Acting AD was responsible for laying the groundwork for the expansion of NSFNET and developing Congressional support.

Bill Wulf, who became the second AD/CISE in 1988, was an active researcher and educator—as have been all successive CISE ADs. He made critical decisions regarding the supercomputer centers and supported the nascent Human Genome Project at NIH (e.g., by funding the effort that developed an efficient implementation of the

BLAST algorithm). He coined the term *collaboratory* to represent a “center without walls, in which the nation’s researchers can perform their research without regard to physical location. . . .” He was AD when the first computing-related Science and Technology Centers were funded in 1988–1991, capitalizing on previous CISE initiatives.

When Wulf left in mid-1990, there was again a long period before Nico Habermann arrived to be the third AD/CISE. Brownstein again stepped in as Acting AD before moving into another NSF position outside of CISE. Mel Ciment, a mathematician who had joined CISE in 1988, was named Deputy AD.

Nico Habermann became the third AD/CISE in October 1991. His tenure was not even a full two years since he passed away unexpectedly in August 1993. He was just beginning to make his presence felt at NSF, but even in that short time he burnished CISE’s internal reputation with his experience as a senior researcher and academic leader, providing a face for computing to the rest of NSF. He oversaw the continuing emergence of NSFNET and served as the executive director of a national committee to implement the 1991 High-Performance Computing and Communications (HPCC) Act by Congress. Among other outcomes of NSF funding during this time was the release of the Mosaic browser in January 1993, the first easily usable Web browser and the model for many later browsers.

Paul Young became the fourth AD/CISE in July 1994, serving for exactly two years. He coordinated a program with other agencies to explore multimedia digital libraries, including a large grant to Stanford. Two graduate students, while working on this project, created the technology that became the Google search engine. The final step (15 years in the making) in the conversion of NSFNET to a self-governing entity—known now as the Internet—occurred smoothly in April 1995, toward the end of Young’s tenure. The plan for developing national inclusive partnerships as successors to the supercomputer centers also began at this time. Young enlisted his Advisory Committee and others to study the appropriate programmatic future for different areas of CISE. Their report was delivered just as he was leaving, but his successor built on it to conduct a reorganization of CISE.

Juris Hartmanis, winner of the 1993 Turing Award and a member of the National Academy of Engineering, immediately followed Young as the fifth AD/CISE. He brought significant experience, recognition, and additional respect to CISE. He raised the bar on quality in all of the funding decisions CISE made through his careful, sustained, personal focus on what each of the programs was actually achieving, not just funding. He restructured CISE in the most consequential way since Bell, de-emphasizing the more engineering-oriented programs. Hartmanis stepped down near the end of 1998, leaving a solid foundation for the changes already appear-

ing on the horizon—momentous changes in the practical world of computing that were starting to impinge on CISE and efforts underway to dramatically increase Congressional appropriations for computing research.

The appearance of the Internet, web browsers, email, and other applications in the last half of the 1990s were primarily based on research funded by NSF and built by leaders whose education and professional training had benefited from NSF funding. Given the attractions of industry, universities again were struggling to retain professors and students. Aggravating the problem, it became clear that CISE funding and programs were not keeping pace.

This was the situation when microbiologist Rita Colwell, a strong supporter of computing, became the eleventh NSF Director. A few months later Ruzena Bajcsy, a senior computer scientist with a background in academic and intellectual leadership, became the sixth AD/CISE, serving from 1998 to 2001. As with Bloch and Bell a decade earlier, Colwell and Bajcsy formed a mutually supportive pair that was effective in creating the next step function change in the size and importance of CISE.

The Information Technology Research (ITR) Initiative had been under development and funding for it was appropriated as Colwell and Bajcsy were just starting at NSF; it was turned into a concrete program under their direction. ITR was the first large, multi-year initiative in computer science aside from the order of magnitude smaller CER program. Over a five-year period, the CISE base budget doubled, opening up opportunities to many more researchers. It had the immediate impact of broadening the computing field by funding research in software systems, IT education and workforce, human-computer interfaces, information management, advanced computational science, scalable information infrastructure, and social and economic implications of computing and communications. Through its focus on group and large-scale interdisciplinary projects, ITR increased investments in sensor networks, bio-inspired and quantum computing, “big data,” sustainability, and human-computer networks. The rapid rise of social networks and online communities at this time stimulated new lines of research, while older aspects of computing research were taking new looks at old subjects such as software systems. There was also a new emphasis on group and large-scale projects, changing the modality of computing research.

As with Hartmanis and Habermann, Bajcsy’s personal research record and demeanor served CISE well in the eyes of her peer ADs and other science leaders, as well as in Congress and the research community. She persuaded the CS community to transform the older order, and she explained the importance of the super-computer program even though it often did not directly benefit most computing

researchers. She used the term *cyberinfrastructure* to refer to the practical infrastructure of computers (especially supercomputers), software, networking, sensors, effectors, other devices, and humans. She chaired an NSF Blue Ribbon panel to study what was needed for scientific research and to make recommendations. The Atkins Report was submitted shortly after she stepped down in August 2001 after 33 months as AD; by then the first few ITR solicitations had been issued, additional funding increases were on the way, CISE's scientific reputation was further enhanced, and the directorate was beginning to respond to the new reality of computing.

Peter Freeman was named the seventh AD/CISE at the end of January 2002 and assumed the position in early May. He also benefited from supportive management (Rita Colwell, followed by Arden Bement as Director, with Joseph Bordogna as Deputy Director) and had a knowledgeable and very capable Deputy AD, Deborah Crawford. He had an agenda to address some of the problems of the last half of 1990s and quickly understood the need for a reorganization of CISE, which occurred in 2003.

Managing the support of cyberinfrastructure, a part of CISE in the early 2000s, has never been easy despite the best efforts to address the broader issues involved (e.g., network provisioning and massive data storage)—efforts that have never been completely successful. At the same time, computer scientists often believe that high-performance computing drains funds that could go to their research (a myth), while many scientific users want greater control over this indispensable tool. As a result, the management location of cyberinfrastructure for scientists moved around in NSF. Freeman attempted to tighten the management of the HPC resources for the mutual benefit of users, NSF, and computer science research, but he had limited and interrupted results. As the more broadly construed cyberinfrastructure became increasingly important, the demands for closer control by the science directorates eventually held sway, and control was moved to an office reporting to the NSF Director.

Two of Freeman's objectives were to strengthen networking and cybersecurity research. While these two are intertwined, they each benefited from heightened national concern over the other, so that CISE often had to race to meet the expectations of Congress in both areas. In the aftermath of the 2001 terrorist acts, cybersecurity research was increased to the extent that funding permitted. By 2005, a modest cybersecurity centers program resulted in a small number of multi-year, multi-investigator awards in addition to a growing number of traditional research awards. These efforts succeeded in attracting more researchers to focus on issues of security and privacy, and were later expanded more robustly.

Networking research was initially steered toward new platforms and testbeds before the GENI Project was initiated in 2004. GENI changed the nature of networking research by focusing on a scientific basis for the engineering of entirely new network structures, starting with a national-scale testbed for network innovations. Gurudatta Parulkar had been a professor and entrepreneur and came to CISE as a program director to give back to his adopted country. As an active member of the networking research community, his personal ties and depth permitted him to pull together and enhance several strands of networking research to form the initial GENI concept.

Freeman's final objective was to expand CISE's outreach to women and under-represented minorities. As with any complex, multi-year activity, it is difficult to judge success, but based on the number of people in the community involved, their persistence over time, and the emulation of CISE's efforts widely within NSF and beyond, it seems fair to say the activities were successful. The flagship effort was the Broadening Participation in Computing (BPC) initiative conceived of and led by Jan Cuny; it has continued for more than a decade and served as the basis for several government-wide initiatives.

The creation of the Computing Community Consortium (CCC) with CISE support in late 2006 provided a means for the community to identify fundamental research needs and articulate ways in which computer science could be applied to fundamental societal needs. One of the results was the National Robotics Initiative. Other major results that originated in the budget request for FY 2008 prepared by Freeman and his team were the Expeditions in Computing program, which intended to define the future of computing, and the Foundation-wide Cyber-enabled Discovery and Innovation (CDI), whose goal was to advance science and engineering through the use of computing concepts.

By 2007, major changes in the size, budget, management, daily operations, and programs across CISE were in operation or budgeted. That placed CISE in an enhanced position of leadership within NSF, across the government, and in the computing community. When Freeman stepped down at the end of January 2007, Deborah Crawford was appointed Acting AD/CISE.

Jeannette Wing became AD/CISE in July 2007 and served three years, continuing the CCC and GENI programs and overseeing the implementation of the CDI and Expeditions concepts and appropriations. For several years, CISE had been providing almost 85% of all federal research support for basic computer science research with far too few staff. Wing streamlined proposal processing by limiting the number of proposals per year a PI could submit, thus reducing NSF staff workload. She also re-balanced funding levels across programs to ensure core CS areas received sufficient

support and modified or redirected some programs, including cybersecurity. She reached out to other directorates and agencies to form cross-cutting initiatives, including the National Robotics Initiative. She also initiated a data-intensive computing program in cooperation with Google and IBM, and a data science initiative as a follow-on to CDI. When she stepped down, Peter Arzberger, a division director in Biology and former executive director of the San Diego Supercomputer Center and the National Partnership for Advanced Computing Infrastructure, served for eight months as the Acting AD.

Farnham Jahanian became the ninth AD/CISE in March 2011 and served slightly more than three years. The administrative senior staff and the division directors once again stepped in to provide continuity. Being organizationally experienced, Jahanian recognized the need for a structural solution to continuity and created the post of deputy division directors to be filled by permanent employees. He quickly formed a bond with the newly appointed NSF Director, Subra Suresh, which developed into yet another productive pairing.

Jahanian welcomed the return of the Office of Cyberinfrastructure (OCI) to CISE and defended the move against predictable criticism. He took special interest in the outreach efforts of Jan Cuny and promoted her work across NSF and beyond. He was open to new ideas and encouraged everyone in CISE to innovate. He promoted the CCC's work and made sure CISE was participating in NSF-wide initiatives. He worked closely with OSTP, especially with Tom Kalil and others in the Obama Administration, helping fashion a variety of initiatives including the National Robotics Initiative, the Federal Big Data Research and Development Initiative, the U.S. Ignite program, and the CS for All initiative. He also encouraged several international efforts. When he stepped down in July 2014, Suzi Iacono, Senior Science Advisor in CISE, was appointed Acting AD.

Jim Kurose, the last AD/CISE covered in our study, took office in January 2015 and ended his term in September 2019 (including a brief stint serving in OSTP in 2018). He has increased cross-directorate and international cooperation, started during Jahanian's tenure. At OSTP he led efforts to address the expanding interest in AI. Given that his service started not far from the cutoff point for our study (2016), we have not described activities during his tenure except to note that Chapter 5 describes a number of issues he faced in his early years in CISE.

13.2 Concluding Remarks

Over its first two-thirds of a century, NSF has played an essential role in supporting American computing, a fact that is not widely known. Its support was essential to the spread of academic computer science beyond the handful of universities

supported in the 1960s by ARPA. By the 1990s, it was the leading supporter of basic computing research, and today it provides almost 85% of federal funding for it. The other major federal supporter for computing research was (D)ARPA, which was active in funding the computer science community from the 1960s through today, but with its greatest impact on basic computer science research being in 1960s to early 1990s.

Briefly comparing the funding activities of the two organizations reveals some important facts about NSF's work. While DARPA was primarily focused on serving its defense mission (sometimes in a broad way, but increasingly in a narrow way after 1990), NSF's goal was to support the general needs of the computational science and computer science communities as well as serving the nation overall (not just its defense, but also, for example, its economic and social needs). While DARPA primarily made large grants to researchers at a small number of elite institutions, NSF intentionally spread its funding more widely, to many different researchers at many different kinds of institutions. While DARPA sought to take advantage of the health of the computing community, NSF's goal was to nurture that health. NSF's mission is broader than DARPA's. NSF also takes responsibility for educating a scientific computing workforce and for ensuring that that workforce provides opportunities for a broad range of people, inclusive and equitable toward gender, race, disability, and other diversities.

NSF is best known in the computing field for its support of research. However, some of NSF's early work in computing had to do with facilities; and these efforts have a direct line to computational science. For many years, there has been a tension between computer science and computational science, with the participants in each of these communities skeptical that NSF's computing organizations could provide appropriately balanced management and funding while not favoring the other community. These tensions are behind the organizational movement of computing facilities programs, including supercomputer centers and cyberinfrastructure, in and out of CISE at various times. Although it is an ongoing development and thus hard to fully evaluate at this time, the situation is changing. Increasingly, computing is seen as *central* to the activities of essentially all science and engineering disciplines, and this has resulted in closer ties between directorates and a major increase in cross-directorate programs. It is also becoming clearer that there are mutual intellectual benefits to both computational science and computer science from being close neighbors. Perhaps, over time, tensions between the two will be effaced.

The United States has a large and robust computer science community—the envy of the world. This is due in no small part to the NSF, which has provided the funding for facilities and computing research in the academy, the main seat of

fundamental new ideas in the nation. It is also due to NSF's attention to computer education and the production of an adequately sized and trained workforce in which every American has an increasing opportunity to participate.

We end with three themes, without elaboration, that we believe can be seen in the history we have presented:

- The overall history presented here is one of slow recognition by NSF of computing as a scientific discipline, building gradually to today, when computing is considered one of the most important and basic disciplines for NSF to support.
- Like many organizations, NSF has sometimes struggled to keep its activities ahead of rapidly advancing computer technology, with one important difference—it is usually several years ahead of others in facing the issue of keeping up.
- There has long been a creative tension between the practical and research aspects of NSF computing activities, just as there is in the computing research it supports.

There are no doubt additional themes to be found in this story; we invite the reader to find them!

APPENDIXES

APPENDIX

Computing Organizations at NSF

W. Richards Adrion

As part of the CISE History project, we collected organizational charts covering almost every year for all of the programs, offices, and divisions that would be incorporated into the Directorate for Computer and Information Science and Engineering (CISE) in 1986. We collected CISE organizational charts from 1986 to the present as well as overall NSF organizational structure charts for each year. These charts were collected from a number of sources: NSF annual reports, professional society newsletters, on-line and print NSF staff directories, the “Federal Yellow Book,”¹ and others. All of our organizational charts will be archived with the CISE History collection at the Charles Babbage Institute. In this appendix, we will select organizational charts we believe illustrate important stages in the development of computing programs at NSF.

When it began in 1950 with Alan T. Waterman as Director, NSF was given a specific role for supporting the interchange of scientific information and the authority to arrange for the publication of scientific and technical information. Director Waterman created an Office of Science Information (OSI), initially reporting to the Assistant Director for Administration, and later directly to him. Four individuals led OSI: Robert Tumbelson (1952–1955); Alberto Thompson (1956–1957); Tomas Jones (Acting 1957); and Burton Adkinson (1957–1958).

In 1955, President Eisenhower signed Executive Order 10807, giving NSF a greatly increased responsibility of making scientific information more easily available to scientists. Under the National Defense Education Act (NDEA), NSF was

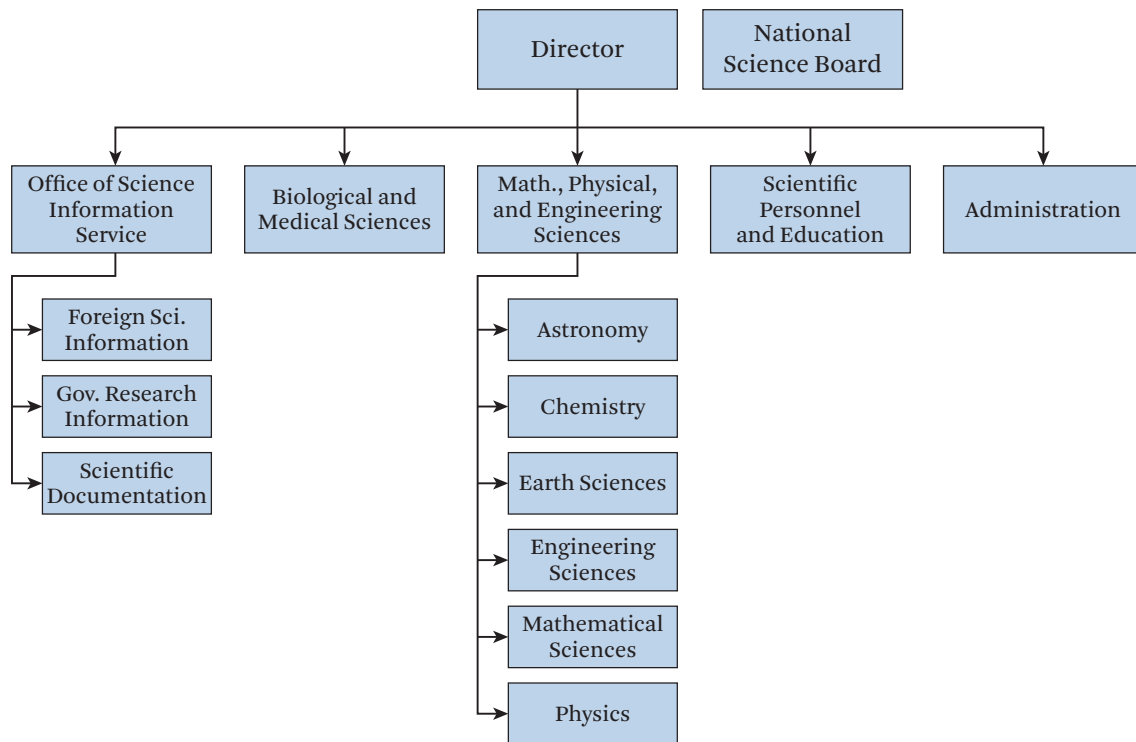


Figure A.1 NSF and the Office of Science Information, FY 1955.

authorized to establish a Science Information Service to address indexing, abstracting, translating, and other services leading to a more effective dissemination of scientific information; and undertake programs to develop new or improved methods for making scientific information available. Burton Adkinson became head of the new office and remained in this role until 1971. Adkinson was succeeded by Melvin Day (1971–1972) and Lee Burchinal (1973–1974).

Other units at the level of the Office of Science Information Service (OSIS) in Figure A.1 were divisions led by assistant directors and other offices (not shown).

Below the divisions and OSIS are programs, so as Fiscal Year 1955 began there were six programs in Mathematical, Physical, and Engineering Sciences, each led by a single program director. The Scientific Documentation program in OSIS, led by Helen Brownson, was the primary funding source for OSIS research and development. By 1964, the two science divisions shown in Figure A.1 were joined by engineering and social sciences divisions and incorporated a Research Directorate, parallel to non-science directorates, and the science programs had evolved into sections. Waterman’s replacement as Director, Leland Haworth, temporarily moved

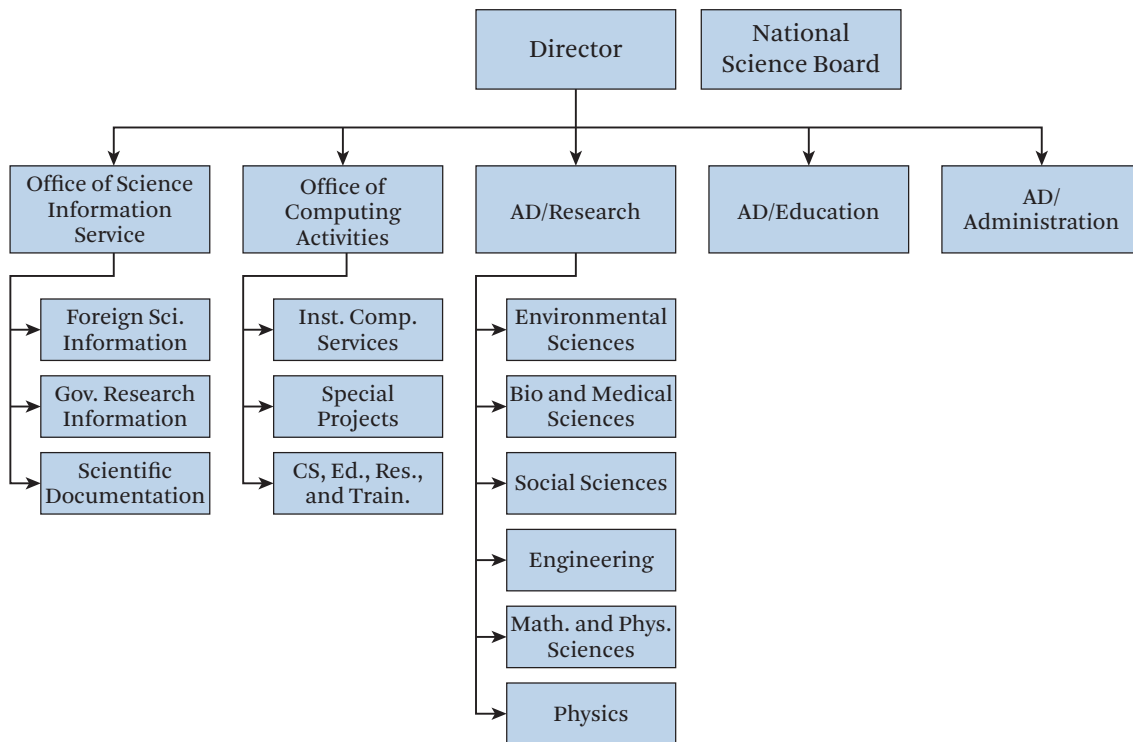


Figure A.2 Office of Science Information Services and the Office of Computing Activities, FY 1968.

OSIS to report to the Assistant Director for Scientific Personnel and Education from 1964 to 1966.

As described in Chapter 1, the National Science Board approved the start of a computing facilities program in 1955, which was approved by Congress in 1959. Until the Office of Computing Activities was established following publication of the Rosser and Pierce reports, computing facilities grants were made through the Mathematics Section (earlier a program), which was housed in the Research Directorate. The Mathematics Section was also home to a nascent computer science program that had begun to make grants for computing research. As noted in Chapter 1, while computing centers were primarily for scientific research, they became computing research centers under the radar. The centers were also home to training programs, which included developing the first courses and curricula in computing and computer science.

As shown in Figure A.2, NSF established the Office of Computing Activities (OCA) in July 1967 to provide federal leadership in the use of computers for research and education. Later, the directive was added as a statutory requirement to the NSF

charter. Concurrently, Director Haworth moved OSIS back to reporting directly to his office with Burton Adkinson remaining as head. The Office of Computing Activities was initially staffed by people moving over from Mathematics. OCA head Milton Rose is credited with recruiting a strong group of program and section managers from outside NSF. For the first time, NSF’s major information and computing programs were at a similar level, but that would be short lived. By 1971, John Pasta had replaced Milton Rose and Melvyn Day had replaced Burton Adkinson. Melvyn Day was later replaced by Lee Burchinal (1973–1974).

In 1972, John Pasta reorganized OCA into three sections: Computer Science and Engineering, Computer Applications in Research, and Computer Innovations in Education, placing a stronger emphasis on research and reflecting the changing nature of computer science and of OCA’s role within NSF. By 1974, NSF had ended both the OCA facilities and training programs, moved the computers in education program into the Education Directorate, and transferred the Office of Computing Activities into the Research Directorate. In 1976, OCA was renamed the Division of Computing Research. While this was the first time NSF organizationally recognized computer science as a field of *research*, challenging times lay ahead for both OCA and OSIS.

In 1976, the NSF would undertake a major foundation-wide reorganization, creating the six directorates shown in Figure A.3 (the administrative directorate and

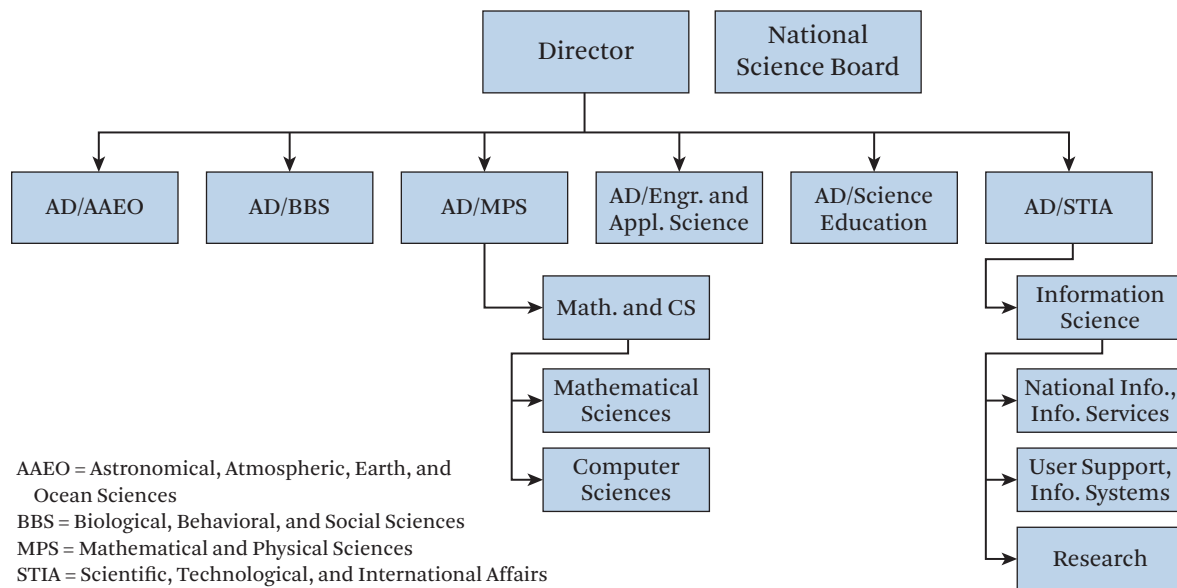


Figure A.3 Computer Science Section and the Division of Information Science, FY 1978.

other offices are not shown). The recently formed Division of Computer Research would again be reorganized as FY 1977 began to become a section in the Mathematical and Computer Sciences Division with John Pasta as Division Director. OSIS would be restructured as the Division of Science Information and moved to the directorate in charge of international affairs. Lee Burchinal remained Division Director, but soon would leave.

Around the same time, the engineering community was beginning to argue strongly to Congress and the Executive Branch for a “National Engineering Foundation.” The NSF response was to move engineering out of MPE and combine it with the Research Applications Directorate, eventually becoming the Engineering Directorate.

In the late 1970s and early 1980s, the information sciences programs were struggling in the Scientific, Technological, and International Affairs (STIA) Directorate, which had cut funding for information science research and centers drastically. The Computer Science Section was struggling to gain prominence and strengthen its reputation. Both units chartered reviews. Following a review of the information sciences programs, DSI became the Division of Information Science and Technology (DIST) under founding director and mathematician Howard Resnikoff. As described in Chapter 2, the series of reports pointing to a crisis in experimental computer science led to larger investments in computer science and eventually to a new Division of Computer Research within MPS in 1984. Resnikoff left, Edward Weiss became Division Director, and DIST was moved to BBS. Just before John Slaughter became Director in 1982, the new Engineering Directorate had created a Division of Electrical, Computer, and Systems Engineering (ECSE). The final piece of the puzzle that Erich Bloch and Gordon Bell would assemble into CISE was the creation in 1984 of the Office of Advanced Scientific Computing (Figure A.4), as a response to a growing demand from the scientific community for access to high-performance computing.

Erich Bloch replaced Edward Knapp as NSF Director in 1985 and in March 1986, Bloch announced his intention to hire C. Gordon Bell as the initial Assistant Director of a new computing directorate. Bloch assigned Charles Brownstein, Rick Adrion, and Gerald Daen to assist Bell in creating it. The details are described in Chapters 2 and 3, but the reorganization went through three phases as shown in Figure A.5.

As the new directorate was developed, Kent Curtis moved from Division Director of Computer Research in MPS to Division Director for Computer and Computation Research (CCR) in CISE. As described in Chapter 3, Peter Freeman soon replaced Curtis when Curtis replaced Adrion as CISE Chief Scientist. John Connolly moved from Office Director to Division Director for Advanced Scientific

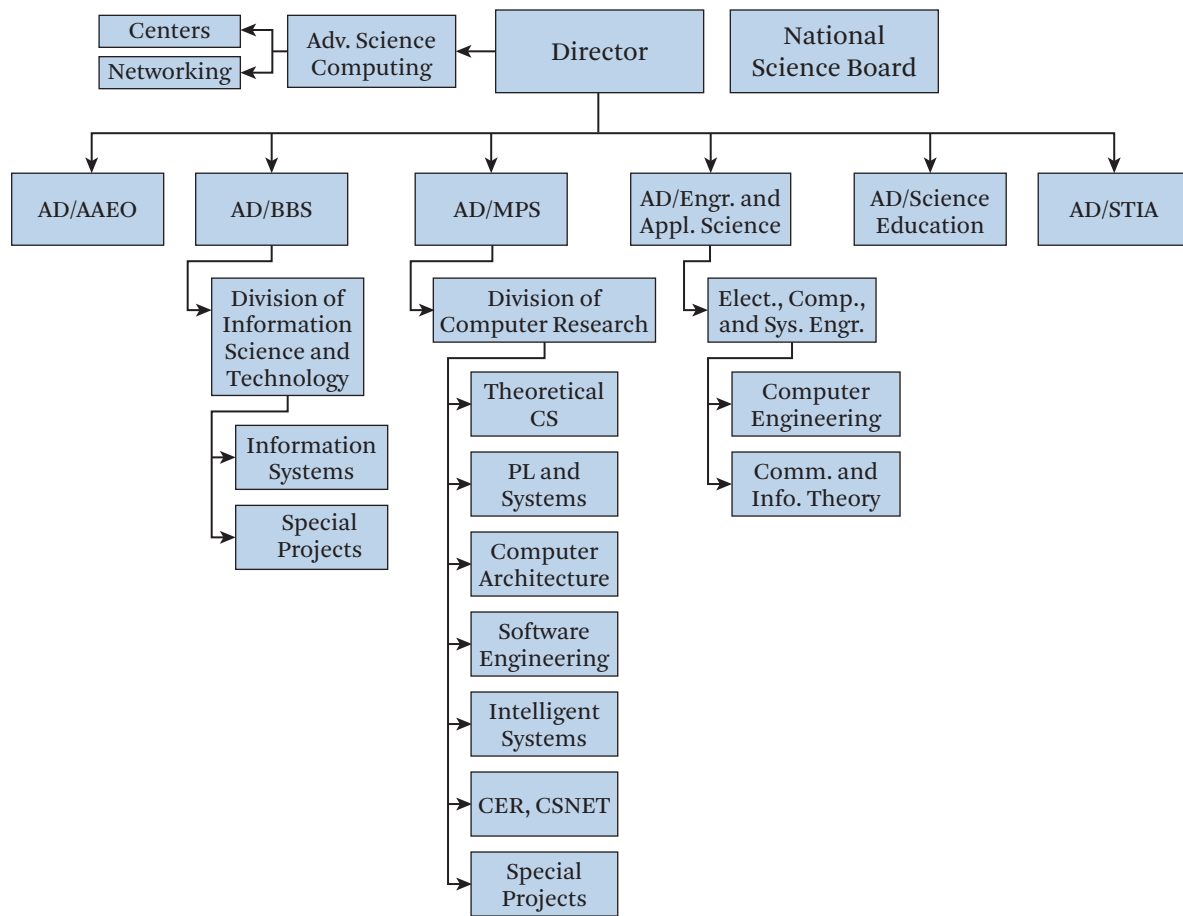


Figure A.4 Computing units across NSF, FY 1984.

Computing (ASC) and was replaced in 1987 by Melvin Ciment. Bernard Chern moved from Engineering to become Division Director for Microelectronic Information Processing Systems (MIPS). Y. T. Chien became acting and, later, permanent Division Director for Information, Robotics, and Intelligent Systems (IRIS). Stephen Wolff had replaced Dennis Jennings as NSFNET program director and became Division Director for Networking and Communications Research and Infrastructure (NCRI) when it was split off from DASC.

The next important change in CISE was the creation of the Office of Cross-Disciplinary Activities (CDA) in 1995. When CISE had been created, the education and infrastructure programs (CER, CISE Equipment, Special Projects) in DCR/MPS

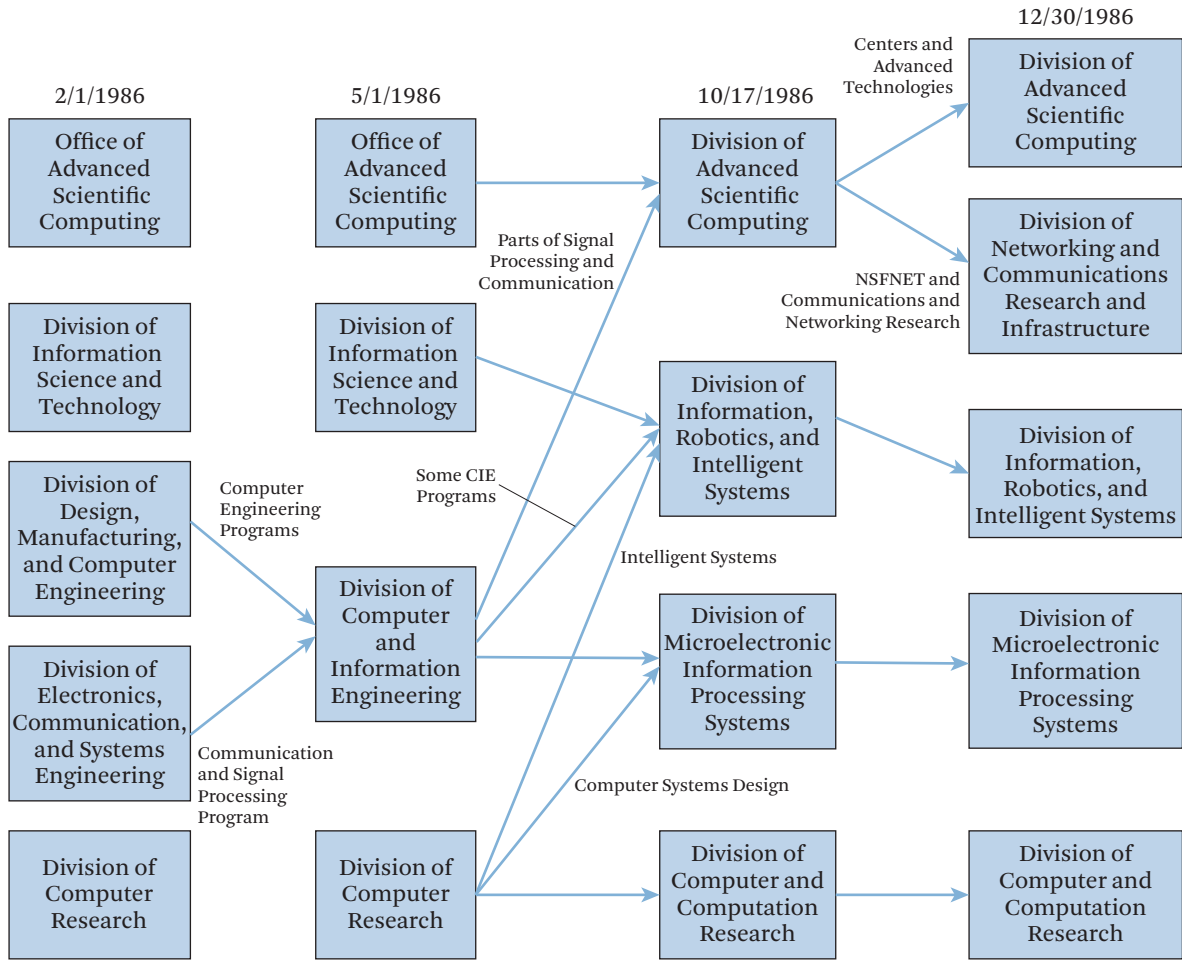


Figure A.5 Creating CISE, 1986.

and DIST/IIS were transferred to CCR/CISE. As an office reporting to the AD/CISE, CDA could coordinate these activities across the directorate.

There were any number of changes in unit names across the Foundation and particularly changes in CISE staff at all levels in the period following the creation of CISE, but the directorate structure remained relatively unchanged until the late 1990s. When Paul Young became AD/CISE, he created three panels to review research directions and a CISE Organizational Review Committee (CORC) to recommend organizational changes. The Committee did so, but Young left before instituting any significant changes. After becoming AD, Juris Hartmanis looked

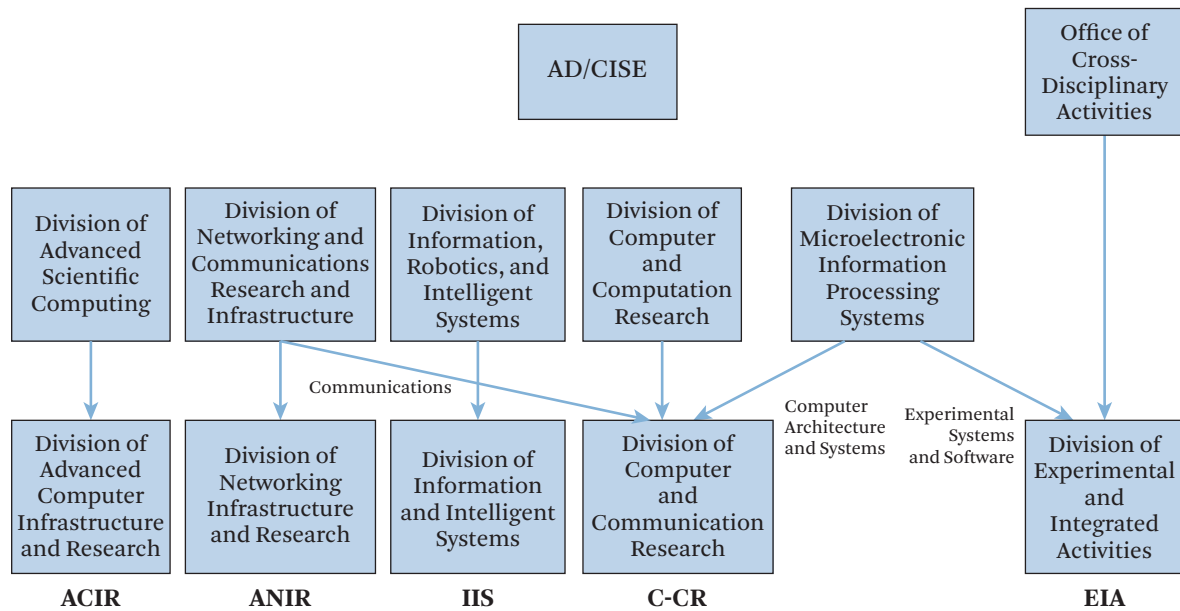


Figure A.6 Reorganizing CISE, 1997.

carefully at the CORC and panel reports, interviewed every program director, and consulted with others in the field. Hartmanis proposed a reorganization and created teams of researchers and CISE staff to review his decision. The result (Figure A.6) placed a greater emphasis on core computer science and de-emphasized some of the strongly engineering aspects (e.g., VLSI and automation) of Bell's design.

The next and perhaps the most significant reorganization occurred under Peter Freeman in 2003. The new organization saw changes in structure, management, and personnel at all levels. The divisional budget and administrative staff were re-assigned in stronger configurations, programs were brought together into clusters with multiple programs and an administrative support team. The networking and high-performance computing divisions lost their research programs and had their operational programs combined into a new Division of Shared Cyberinfrastructure. The networking and communication research programs were moved into a new Division of Computing and Networking Systems (CNS). EIA, which during ITR had grown its biological, new technologies, experimental systems, workforce, and education, saw all of its programs except education and workforce divided among IIS and C-CR (renamed Computing and Communications Foundations). Freeman

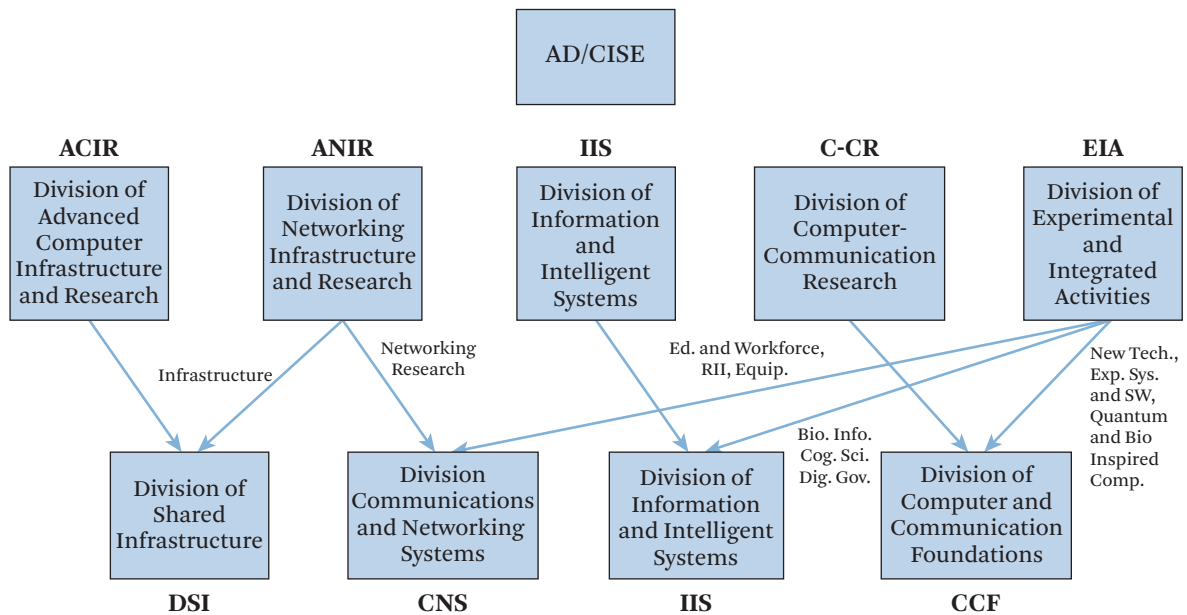


Figure A.7 Reorganizing CISE, 2004.

recruited Jan Cuny to come in to lead a new program in Broadening Participation in Computing (BPC) along with education and workforce. While Cuny's programs were assigned to CNS, her portfolio was directorate-wide.

With the significant exception of the cyberinfrastructure programs, the organization of CISE has remained remarkably stable through today. The constant pressure from computational scientists to pull the cyberinfrastructure programs out of CISE and place them under broader NSF management, coupled with the continuing suspicion by computer science researchers that these programs take funds away from computing research, resulted in the programs moving to the Office of Cyberinfrastructure (OCI) in 2006. In 2013, OCI moved back into CISE as the Division of Advanced Cyberinfrastructure (ACI). Today the office reports both to the current CISE AD and to the Director.

Notes

1. The "Federal Yellow Book," published by Leadership Directories Inc., includes contact information for over 45,000 U.S. federal decision-makers located within the Washington, DC metropolitan area. See: <https://www.leadershipconnect.io/products/print-leadership-directories/>.

APPENDIX

CISE Oral Histories List

The following people were interviewed as part of this project at some time in the period from 2017 to 2018. Other oral histories exist that pertain to our story, and while we have used them in our research, they do not appear in this list. The scholar who wants to read additional interviews about NSF should look first at the general oral history collections held by the Charles Babbage Institute (CBI), the IEEE Center for the History of Electrical Engineering, and the Computer History Museum archives. Initials by the interviewee's name in our list identify the interviewer: RA for Rick Adrion, WA for William Aspray, and PF for Peter Freeman. These interviews were transcribed and lightly edited by the project staff, and signed off by the interviewees. Only the edited and approved transcripts, not the audiotapes or the raw transcripts, will be available. These interviews are being deposited at the CBI, and interested readers should consult their website regarding how to gain access.

Kamal Abdali (RA)

W. Richards Adrion (WA)

Alfred Aho (WA)

Peter W. Arzberger (WA)

Ruzena K. Bajcsy (WA)

C. Gordon Bell (WA)

Arden L. Bement (WA)

Jim Bottum (RA)

Charles N. Brownstein (WA)

John Cherniavsky (RA)

Melvyn Ciment (WA)

Rita R. Colwell (WA)

John Cozzens (RA)

Dennis Jennings (RA)

Deborah Crawford (WA)

Janice E. Cuny (PF)

David J. Farber (RA)

Darleen Fisher (WA)

Michael Foster (WA)

Peter A. Freeman (WA)

Susanne Hambruch (RA)

William C. Harris (PF)

Juris Hartmanis (RA)

Richard Hirsh (WA)

John Hopcroft (RA)

Susan Iacono (PF)

Farnam Jahanian (WA)

Frank Rhodes (PF)

Anita Jones (WA)	Edwina Rissland (RA)
Tom Kalil (RA)	Rita Rodriguez (WA)
John King (RA)	Alfred Spector (WA)
Rao Korasaju (PF)	George Strawn (WA)
James Kurose (WA)	Gary Strong (RA)
Larry Landweber (WA)	Al Thaler (RA)
Carl Landwehr (WA)	Howard Wactlar (PF)
Neal Lane (PF)	Rick Weingarten (PF)
Anita LaSalle (PF)	John White (PF)
Irene Lombardo (RA)	Jeannette Wing (RA)
Keith Marzullo (WA)	Steve Wolff (RA)
Gracie Narcho (PF)	William A. Wulf (WA)
Michael Pazanni (RA)	

APPENDIX

Short Biographies

Following are brief sketches of 25 people named in this book; information on them and others can usually be found easily on the Web. Many more names were mentioned, of course, but the role(s) they played did not figure prominently in our narrative.

Peter Arzberger, a bio-mathematician, was Acting AD/CISE for 8 months beginning July 2010. He also had served at NSF as a Program Officer in Mathematical and later Biological Sciences, and the Division Director of Biological Infrastructure. He has been Executive Director of the San Diego Supercomputer Center (SDSC) and the National Partnership for Advanced Computational Infrastructure (NPACI).

Ruzena Bajcsy, a computer scientist, was the 6th AD/CISE (1998–2001). She holds Ph.D.s in EE from Slovak Technical University (1967) and CS from Stanford (1972) under John McCarthy. She joined the faculty of the Computer and Information Science Department at the University of Pennsylvania in 1972 and was the founder of a lab for robotics and related subjects. She was department chair before joining NSF. As AD/CISE she further raised the stature of CS within NSF and more broadly in the USG, structured the ITR Program to broaden and deepen CS, and successfully managed the substantial increase in funds. After leaving NSF, she joined the faculty at UC Berkeley where she was Director of the CITRIS S&T Center. Today she is Director Emeritus of CITRIS and a Professor at Berkeley. She is a member of the NAE and the IOM.

C. Gordon Bell, an engineer, computer architect, and manager, was the 1st AD/CISE (1986–1987). Trained at MIT, he was a very early employee of the Digital Equipment Corporation, designed the groundbreaking PDP line of computers, and guided the development of the VAX computers while VP of Engineering. He also held a faculty appointment at Carnegie Mellon University beginning in 1967, where he authored a comprehensive book on computer architecture with Allen Newell. Bell played a crucial role in developing the structure of CISE. After NSF he was involved in several start-ups and then joined Microsoft Research to work on the concept of lifelogging.

He is a member of the NAE and was awarded the National Medal of Technology in 1991. He is best known in today's computing field as the founder of the Gordon Bell Prize.

Arden Bement was the 12th director of the National Science Foundation (2004–2010). He has a Ph.D. in metallurgical engineering from the University of Michigan. In his long career, Bement worked in industry, government, and academia. From 1976 to 2001 he served various roles in the USG, including at DARPA; just prior to coming to NSF he was director of NIST. As director of NSF from 2004 to 2010, Bement oversaw many changes and developments including the establishment of the Office of Cyberinfrastructure and the broadening of international collaboration. After NSF he joined Purdue University where today he is a Professor Emeritus.

Erich Bloch was an electrical engineer, trained in Switzerland and the U.S. during and shortly after WW2. He was the 8th NSF Director (1984–1990), the only one ever without a Ph.D., a fact of which he was proud. Before NSF Bloch had a long career at IBM, playing key roles in the design of the Stretch supercomputer and the System 360, accomplishments for which he was awarded the first National Medal of Technology. At IBM, he rose to be Vice President for Technical Personnel. At NSF he introduced the use of email and more effective use of computing, created CISE, appointed women to senior posts, and made other significant changes. He was a member of the NAE, among other honors. After NSF he was the first Distinguished Fellow at the Council on Competitiveness and was one of the founders of the Washington Advisory Group.

Charles (Chuck) Brownstein, a political scientist and early user of computers in that field, was a faculty member at Lehigh before joining NSF in 1974 as a program officer in telecommunications policy and impacts. He worked with Erich Bloch and then Gordon Bell and others on the design of CISE. After 1986 he served two periods as Acting AD/CISE, provided leadership on several government-wide committees, and served in other management positions. After leaving NSF in 1994, he spent ten years at the Corporation for National Research Initiatives (CNRI) as Senior Scientist and later served in senior staff positions at the National Research Council and the ANSER Corporation before retiring.

Melvyn “Mel” Ciment, an applied mathematician, holds a Ph.D. from Courant Institute at New York University. He held academic and research positions at NYU, University of Michigan, Tel Aviv University, the National Bureau of Standards (now the National Institute for Standards and Technology), and the Naval Surface Weapons Center before coming to NSF in 1983. He founded the computational mathematics

program in MPS before joining CISE in 1986 in the Advanced Scientific Computing Division. He served as a Deputy AD and Acting AD at various periods until he left NSF in 1999 to join the Potomac Policy Institute. He was awarded the Meritorious Service Award at NSF for management contributions to the HPCC program.

Rita R. Colwell, a microbiologist, was the 11th NSF Director (1998–2004). Prior to NSF, she was president of the University of Maryland Biotechnology Institute and professor of microbiology and biotechnology, a member of the National Science Board (1984–1990), and president of the American Association for the Advancement of Science (AAAS). Dr. Colwell was instrumental in helping CISE grow, promoting international connections, responding to 9/11, and serving as co-chair of the National Science and Technology Council. She is now a Distinguished University Professor at the University of Maryland and Johns Hopkins University Bloomberg School of Public Health. She is a member of the National Academy of Science (NAS) and received the National Medal of Science in 2006.

Deborah Crawford, an electrical engineer, received her Ph.D. in information systems engineering from the University of Bradford, England. She worked at Bell Labs, UC Santa Barbara, and NASA's Jet Propulsion Lab (JPL) before joining NSF in 1993 as a program director in Engineering. She was Senior Staff Associate for Policy and Strategic Planning in the Office of the Director before joining CISE in 2002 as Deputy AD. While in CISE she was also Acting Director of the just created Office of Cyberinfrastructure for a year and was Acting AD twice. After leaving NSF in 2011, she served as VP of Research at Drexel University, Director of the International Computer Science Institute at UC Berkeley, and currently is Senior Vice President for Research and Economic Development at George Mason University. She was twice recognized while at NSF with a Presidential Rank Award (in 2006 and 2010).

Kent K. Curtis graduated from Yale in 1948, received a master's in physics from Dartmouth in 1950, and studied theoretical physics and music at Berkeley. He led the mathematics and computing group at the Lawrence Berkeley Laboratory (1955–1967), while also serving as a lecturer at UC Berkeley. In 1967 he joined NSF as head of the Computer Science and Engineering Research Section. Through this role, Curtis was highly influential in shaping the development of university and college computing centers as well as growing computer science and engineering research. In 1984, after serving as section head for computer science for 10 years, he became director of the Computer Research Division (DCCR), where he was instrumental in the development of the Coordinated Experimental Research program, CSNET, and

CISE. He was appointed chief scientist for CISE in 1987, but his life was cut short when he died of cancer at the age of 60.

Erwin Gianchandani, a computer scientist, received a Ph.D. in biomedical engineering from the University of Virginia in 2009. He was Executive Director of the Computing Community Consortium until 2012, when he entered NSF as Deputy Director for the Division of Computer Networks and Systems; in 2015 he was appointed Deputy AD/CISE, where he continues to serve today. Gianchandani has published extensively and continues to participate in various scientific meetings.

Nico Habermann, a computer scientist, was the 3rd AD/CISE (1991–1993). He received a doctorate in applied mathematics from Technological University, Eindhoven, Netherlands, in 1967. He first worked on language design and implementation for a number of early computer languages, including Algol 60; then later on process coordination, operating systems, and software engineering. He joined the faculty of the Department of Computer Science at Carnegie Mellon University (CMU) in 1968, becoming the first dean of the School of Computer Science in 1988. He co-founded the CMU Software Engineering Institute in 1985, serving as its first director. He passed away unexpectedly in 1993 while AD/CISE.

Juris Hartmanis, a computer scientist and computational theorist, was the 5th AD/CISE (1996–1998). He received a Ph.D. in mathematics from Cal Tech in 1955 and joined Cornell University in 1965. He helped found the CS department and served as its first chair. In 1993 he was co-winner of the Turing Award. In 1992, he co-edited a National Academy study, “Computing the Future: A Broader Agenda for Computer Science and Engineering,” which has been influential in the development of computer science. After his service at NSF, Hartmanis returned to the faculty at Cornell where he remains active as an Emeritus Professor. He is a member of the National Academy of Engineering and the National Academy of Science, a rare double honor.

Suzi Iacono, a scholar of social informatics, received her Ph.D. in information systems from the University of Arizona. Prior to coming to NSF, she held a faculty position at Boston University School of Management. Since joining NSF in 1998 as a program director, she has served in many leadership roles, including Deputy AD, Acting AD, and Senior Science Advisor—all in CISE. Today she is head of the NSF Office of Integrative Activities (OIA). She has published extensively in social informatics.

Farnam Jahanian, a computer scientist and entrepreneur, was the 9th AD/CISE (2011–2014). After receiving a Ph.D. in computer science from the University of

Texas at Austin in 1989, he worked at the IBM T. J. Watson Research Center until 1983, when he joined the faculty of computer science and engineering at the University of Michigan. His research centered on computer networks, and he was director of the Software Systems Laboratory. He was department chair from 2007 to 2011. After finishing his service at NSF in 2014, he joined the faculty at CMU and served as VP of Research, then provost, before becoming the tenth president of CMU in 2018, a position in which he is still serving.

James Kurose, a computer scientist, was the 10th AD/CISE (2015–2019). He holds a Ph.D. in computer science from Columbia University and has published extensively in computer networks as a faculty member since 1984 at the University of Massachusetts Amherst. He has been a member of the Board of Directors of the CRA and a visiting scientist at several top schools and labs in the U.S. and abroad. Kurose has received many awards for his research, teaching, and service, including the IEEE Infocom Award, the ACM SIGCOMM Lifetime Achievement Award, and the Taylor Booth Award of the IEEE for his educational activities. As AD he has been very active in cross-directorate, interagency, and international activities.

Neal Lane served as the 10th NSF Director (1993–1998). A physicist and strong supporter of improving communication between the scientific community and the general public, he was an NSF rotator and served on several NSF advisory study panels. Lane joined the Rice University faculty in 1966 and was provost (1986–1993) just prior to becoming Director. In 1998, he was appointed director of the White House Office of Science and Technology Policy. After leaving government service in 2001, he returned to the faculty at Rice, where today he is a Professor Emeritus. He has won numerous awards, including the NASA Distinguished Service Award, the National Academy of Sciences Public Welfare Medal, and the National Science Board's Vannevar Bush Award.

John Pasta served as head of the Office of Computing Activities at NSF from 1970 to 1976 and later as Division Director of Mathematical and Computer Sciences from 1976 to 1981. Pasta received his doctorate in physics in 1951 from New York University and began working at Los Alamos National Laboratory, where he collaborated with the senior physicists Enrico Fermi and Stanislaw Ulam on one of the earliest “computer experiments” and co-authored several papers with them. In the early 1960s Pasta joined the faculty at the University of Illinois until 1970 when he came to NSF. He was instrumental in expanding computer science within the Foundation. His service was recognized in 1979 when he received the NSF Distinguished Service Award. His contributions were cut short in 1981 when he died of cancer.

George Strawn received his Ph.D. in mathematics from Iowa State University (ISU) in 1969 and was a member of the CS faculty there, eventually heading the ISU Computation Center and becoming chair of the CS Department. He joined NSF in 1995 as division director of networking. Strawn served as executive officer for Ruzena Bajcsy and Acting AD/CISE when she left. Starting in 2002, he was NSF CIO for six years and then headed the interagency NITRD Coordination Office until 2015. He is currently the director of the Board on Research Data and Information for the National Academies. He was elected a Fellow of AAAS in 2012.

Subra Suresh, a materials engineer, was the 13th director of NSF (2010–2013). Suresh holds a Doctor of Science from MIT; he then did post-doctoral work at Berkeley. He was on the faculty at Brown University for ten years before moving to MIT in 1993. Just prior to joining NSF, Suresh was dean of the MIT School of Engineering and the Vannevar Bush Professor of Engineering. In 2013 he left NSF to become the ninth president of CMU. In 2018 Suresh became president of Singapore's National Technological University.

Alan T. Waterman was the first director of the National Science Foundation (1951–1963). He received his Ph.D. in physics from Princeton in 1916; after graduation, he taught physics at the University of Cincinnati. Over the next 30 years, Waterman continued his academic career, but he also began a career with the government, holding positions such as deputy chief and chief scientist of the Office of Naval Research. It is likely this experience that made Waterman the perfect fit as NSF's first director. Over his two terms, Waterman truly shaped NSF into the institution it would become.

Frederick (Rick) Weingarten received a Ph.D. in mathematics from Oregon State University in 1966, spent a short time at NSF, and after post-doc work at Livermore National Laboratory took a permanent position at NSF as Assistant Program Director and later Program Director within the Office of Computing Activities, the Division of Computing Research and the Computer Sciences Section. After over a decade overseeing research programs in networking and the social impacts of information technology, he joined the Office of Technology Assessment, where he continued his support via several important studies. Later he became executive director of the Computer Research Association, putting it on the path to becoming an important, professional association. After CRA, he was Director of the Office for Information Technology Policy of the American Library Association. He is now retired.

Jeanette Wing, a computer scientist, was the 8th AD/CISE (2007–2010). She earned her Ph.D. in CS from MIT in 1983. She then joined the faculty at the University of Southern California briefly before joining the faculty at CMU. She was head of the Department of Computer Science within the School of Computer Science from 2004 to 2007: a position she stepped down from to come to NSF. After NSF she joined Microsoft Research, quickly becoming corporate VP of Research. She is currently the Avaneessians Director of the Data Sciences Institute at Columbia University, as well as a professor of computer science there.

William A. (Bill) Wulf, a computer scientist, was the 2nd AD/CISE (1988–1990) and was a faculty member at Carnegie Mellon University from 1966 to 1981. He is known for his work on operating systems, systems software, and computer architecture. In 1981 he co-founded Tartan Labs where he remained until 1988, when he became AT&T Professor of Computer Science at the University of Virginia. After NSF Wulf briefly returned to his UVA position. He was then president of the National Academy of Engineering (NAE) for 11 years. Wulf has been awarded several prestigious awards, including the Karl V. Karlstrom Outstanding Educator Award and the ACM Policy Award.

Paul Young, a computer scientist, was the 4th AD/CISE (1994–1996). He received his Ph.D. from MIT in 1963. He began his career at Purdue, then moved to the University of New Mexico and later to the University of Washington, where he was a professor and chair from 1983 to 1988 and then Associate Dean of Engineering. He chaired the Board of the CRA from 1989 to 1991. When he left NSF, he returned to the University of Washington as a faculty member and later moved to the University of Wisconsin–Madison. He is now retired.

APPENDIX

CISE History Archive (CHA)

As part of the process of this project, the team collected approximately 5,000 documents in paper or electronic form. The process by which these materials were collected is described in the preface. The project team has screened the materials to remove those that NSF would not want to see in the public domain. In particular, the team removed “privileged” NSF documents that were produced after 2006, as well as any documents that divulge the names of reviewers, candidates for positions, respondents to NSF solicitations, personnel actions, and other personal information. In the case of longer documents that have only a small amount of private information, the team has occasionally redacted the private information so that the majority of the document may be made public.

These materials were donated in mid-2019 to the Charles Babbage Institute Archives (CBIA) (<http://cbi.umn.edu>), an archival unit of the University of Minnesota. The vast majority of the collection is in digital form, either because it was created digitally or because it was scanned for the project. The materials have been only roughly organized by the project team. Further organization, if any, will be done by CBIA archivists; finding aids, if any, will be also prepared by them. CBIA archivists will determine the access regulations for the collection.

Materials were selected for this collection because they are relevant to the history of CISE and its NSF predecessors. They include the following categories:

- Publications (formal, bound, public)
- Near-print materials (often unbound, often hard to access)
- Unpublished documents (announcements, Dear Colleague letters, personal collections, others)
- Internal documents (non-sensitive, prior to 2007)
- Internal documents (non-sensitive, 2007–2016, cleared)

- Notes, emails, drafts (non-sensitive, cleared or old)
- Documents/publications from non-NSF sources (non-sensitive)
- Items generated by our work that are relevant (e.g., biographies)
- Oral histories done for this project, which have been completed and have been signed off by the interviewees and do not reveal private information
- Other sources regarding NSF, CISE, and related activities (e.g., website for National Academies Press publications)

APPENDIX

Abbreviations and Acronyms

NOTE: This list may not be complete and contains some abbreviations not used in this book, but used in cited references. It is important to note that some abbreviations (e.g., CCR) are used in various forms (e.g., DCR, CCF, C-CF) over time to designate a similar set of programs.

AAAI	American Association for Artificial Intelligence
AAAS	American Association for Advancement of Science
AAEO	Astronomical, Atmospheric, Earth, and Ocean Sciences (Directorate of)
AB	Advisory Board
ACI	Advanced Computing Infrastructure
ACM	Association for Computing Machinery
AD	Assistant Director
AFOSR	Air Force Office of Scientific Research
ARO	U.S. Army Research Office
ARPA	Advanced Research Projects Agency
ARRA	American Recovery and Reinvestment Act
ASC	Advanced Scientific Computing
BAA	Broad Agency Announcement
BBS	Biological, Behavioral and Social Sciences (Directorate of)
BD2K	Big Data to Knowledge initiative (NIH)
BIO	Biological Sciences (Directorate of)
BPC	Broadening Participation in Computing program
CAI	computer-assisted instruction

CCC	Computing Community Consortium
CCF	Computer and Communications Foundations Division
C-CR	Computer-Communications Research Division
CCR	Computer and Computation Research Division
CDA	Cross Disciplinary Activities (Office of)
CDI	Cyber-enabled Discovery & Innovation program
CE21	Computing Education for the 21st Century program
CENS	Center for Embedded Networking Systems
CER	Coordinated Experimental Research program
CISE	Computer and Information Science and Engineering (Directorate of)
CMS	Civil & Mechanical Systems Division
CNRI	Corporation for National Research Initiatives
CNS	Computer and Network Systems Division
Co-PI	Co-Principal Investigator
COSATI	Committee on Scientific and Technical Information
CPATH	CISE Pathways to Revitalized Undergraduate Computing Education program
CPS	Cyber-Physical Systems program
CRA	Computing Research Association
CREN	Corporation for Research and Educational Networking
CRII	CISE Research Initiation Initiative
CS	Computer Science
CSIA	Cyber Security and Information Assurance
CSNET	Computer Science Network
CSR	Computer Systems Research program
CSS	Computer Science Section
CSTB	Computer Science and Telecommunications Board
DAD	Deputy Assistant Director
DARPA	Defense Advanced Research Projects Agency
DASC	Division of Advanced Scientific Computing
DCIE	Division of Computer and Information Engineering
DCR	Division of Computing Research

DD	Division Director
DDD	Deputy Division Director
DGA	Division of Grants and Agreements
DIST	Division of Information Science and Technology
DMCE	Division of Design, Manufacturing, and Computer Engineering
DMCS	Division of Mathematical and Computer Sciences
DMS	Division of Mathematical Sciences
DoD	Department of Defense
DoE	Department of Energy
DSI	Division of Science Information
ECCS	Electrical, Communications and Cyber Systems Division
ECSE	Electrical, Computer and Systems Engineering Division
E2CDA	Energy-Efficient Computing: from Devices to Architectures program
EFT	Extensible Terascale Facility
EHR	Education and Human Resources (Directorate of)
EIA	Experimental and Integrative Activities Division
ENG	Directorate for Engineering
EPSCoR	Established Program to Stimulate Competitive Research
ERC	Engineering Research Center
ERE	Environmental Research and Education (Directorate of)
FCC	Federal Communications Commission
FCCSET	Federal Coordinating Council for Science, Engineering, and Technology
FCSTC	Federal Coordinating Scientific Technical Committee
FEVS	Federal Employment Viewpoints Surveys
FFRDC	Federally Funded Research & Development Center
FOCS	Foundations of Computer Science (conference)
FRICC	Federal Research Internet Coordinating Committee
GENI	Global Environment for Network Innovations
GEO	Geosciences (Directorate of)
GLEON	Global Lake Ecological Observatory Network
GRASP	General Robotics, Automation, Sensing & Perception (Lab)

GRPA	Government Performance and Results Act
HPC	High Performance Computing
HPCCI	High Performance Computing & Communications Initiative
HSST	House of Representatives' Science, Space, and Technology Committee
IC	Intelligence Community
I-CORPS	Innovation Corps program
ICST	Institute for Computer Science and Technology (NBS/NIST)
IEEE	Institute for Electrical and Electronic Engineers
IEEE-CS	IEEE Computer Society
IETF	Internet Engineering Task Force
IIP	Industrial Innovation and Partnerships
IIS	Information and Intelligent Systems Division
ILLIAC	Illinois Automatic Computer
IOM	Institute of Medicine (now National Academy of Medicine)
IPA	Intergovernmental Personnel Act Mobility Program assignment
IR&D	Independent Research & Development
IRIS	Information, Robotics, and Intelligent Systems Division
IT ²	Information Technology for the 21st Century
ITAR	International Traffic in Arms Regulations
ITR	Information Technology Research program
ITWF	Information Technology Workforce program
IUCRC	Industry-University Cooperative Research Centers
JCL	Job Control Language
KDI	Knowledge Distributed Intelligence (initiative)
MCC	Microelectronics and Computer Technology Corporation
MEDEA	Measurements of Earth Data for Environmental Analysis
MIPS	Microelectronic Information Processing Systems Division
MPE	Math, Physical Sciences, & Engineering (Directorate of)
MPS	Mathematical and Physical Sciences (Directorate of)
MREFC	Major Equipment and Facilities Construction
MRI	Major Research Instrumentation
NAE	National Academy of Engineering

NAI	Network Access Identifier
NBCR	National Biomedical Computation Resource
NBS	National Bureau of Standards (see NIST)
NCAR	National Center for Atmospheric Research
NCO	National Coordination Office
NCRI	Networking and Communications Research and Infrastructure Division
NCSA	National Center for Supercomputing Applications
NDEA	National Defense Education Act
NEON	National Ecological Observatory Network
NeTS	Networking Technology and Systems program
NIST	National Institutes of Standards and Technology (see NBS)
NITRD	Networking & Information Technology Research Development (Office of)
NPACI	National Partnership for Advanced Computational Infrastructure
NRC	National Research Council
NREN	National Research and Education Network
NRI	National Robotics Initiative
NSA	National Security Agency
NSB	National Science Board
NSF	National Science Foundation
NSFNET	National Science Foundation Network
NTSC	National Science and Technology Council
OAC	Office of Advanced Cyberinfrastructure
OAD	Office of the Assistant Director
OASC	Office of Advanced Scientific Computing
OCA	Office of Computing Activities
OCI	Office of Cyberinfrastructure
OCR	Optical Character Recognition
O/D	Office of the Director
OGC	Office of General Counsel
OIA	Office of Integrative Activities

OISE	Office of International Science and Engineering
OLPA	Office of Legislative and Public Affairs
OMB	Office of Management and Budget
ONR	Office of Naval Research
OSI	Office of Scientific Information
OSIS	Office of Science and Information Services
OSRD	Office of Scientific Research and Development
OSTP	Office of Science and Technology Policy
OTA	Office of Technology Assessment
PACI	Partnerships for Advanced Computational Infrastructure (program)
PAWR	Platforms for Advanced Wireless Research program
PC3	Pervasive Computing and Communications Collaboration program
PCAST	President's Council of Advisors on Science and Technology
PD	Program Director
PI	Principal Investigator
PITAC	President's Information Technology Advisory Committee
PSAC	President's Science Advisory Committee (under Kennedy)
RFI	Request for Information
RI	Research Infrastructure program
RIA	Research Initiation Awards
R&RA	Research and Related Activities (NSF budget category)
S&AS	Smart and Autonomous Systems program
SBE	Social, Behavioral, & Economic Sciences (Directorate of)
SBIR	Small Business Innovation Research program
SCORE	Special Cyber Operations Research and Engineering
SIAM	Society for Industrial and Applied Mathematics
SIGACT (ACM)	Special Interest Group on Algorithms and Computation Theory
SIGSOFT (ACM)	Special Interest Group on Software Engineering
SMET	Science, Mathematics, Engineering, and Technology (see STEM)
SMRT	Senior Management Round Table

SOCS	Social and Computational Systems program
SRC	Semiconductor Research Corporation
STC	Science and Technology Centers program
STEM	Science, Technology, Engineering, and Mathematics
STIA	Scientific, Technological, and International Affairs (Directorate of)
TCS	Theoretical Computer Science program
TIGR	The Institute for Genomic Research
TRUST	Team for Research in Ubiquitous Secure Technology
USG	United States Government
vBNS	Very high-speed Backbone Network Service
VLSI	Very-large-scale integrated circuit technology

Bibliography of Secondary Literature on NSF and the History of U.S. Computing

- J. Abbate. 2000. *Inventing the Internet*. Cambridge, MA: MIT Press.
- J. Abbate. 2010. Privatizing the Internet: Competing visions and chaotic events, 1987–1995. *IEEE Annals of the History of Computing*, 32(1). DOI: [10.1109/MAHC.2010.24](https://doi.org/10.1109/MAHC.2010.24).
- J. Abbate. 2012. *Recoding Gender*. Cambridge, MA: MIT Press.
- J. Abbate. October 2013. Is Computer Science “science”? A half-century debate. Keynote talk, 2nd International Conference on History and Philosophy of Computing. Paris. <https://hapoc2013.sciencesconf.org/27047/document>.
- B. W. Adkinson. 1976. National Science Foundation—Science information. In A. Kent, H. Lancour, and J. Daily, eds., *Encyclopedia of Library and Information Science: Volume 19*. CRC Press.
- J. Agar. 2003. *The Government Machine: A Revolutionary History of the Computer*. Cambridge, MA: MIT Press.
- J. Agar. 2006. What difference did computers make? *Social Studies of Science*, vol. 36. DOI: [10.1177/0306312706073450](https://doi.org/10.1177/0306312706073450).
- A. Akera and F. Nebeker. 2002. *From 0 to 1: An Authoritative History of Modern Computing*. New York: Oxford University Press.
- A. Akera. 2007. *Calculating a Natural World: Scientists, Engineers, and Computers During the Rise of U.S. Cold War Research*. Cambridge, MA: MIT Press.
- R. Allan. 2001. *A History of the Personal Computer: The People and the Technology*. London: Allan Publishing.
- E. Altman. 1993. National Science Foundation’s support of information research. *Encyclopedia of Library and Information Science*, 52(15).
- E. Altman and S. Brown. 1991. From retrieval to robotics: The National Science Foundation’s support of information research. *Government Information Quarterly*, 8(4).

- J. A. Anderson. 2017. *After Digital: Computation as Done by Brains and Machines*. New York: Oxford University Press.
- T. Appel. 2000. *Shaping Biology: The National Science Foundation and American Biological Research, 1945–1975*. Baltimore, MD: Johns Hopkins University Press.
- W. Aspray. 2016. *Participation in Computing: The National Science Foundation's Expansionary Programs*. Cham, Switzerland: Springer.
- W. Aspray and B. O. Williams. April 27–29, 1993. Computing in science and engineering education: The programs of the National Science Foundation. *Electro 93 International, Conference Record, Volume 2*, pp. 234–239, Edison, NJ.
- W. Aspray and B. O. Williams. 1994. Arming American scientists: NSF and the provision of scientific computing facilities for universities, 1950–1973. *IEEE Annals of the History of Computing*, 16(4). DOI: [10.1109/85.329758](https://doi.org/10.1109/85.329758).
- W. Aspray, B. O. Williams, and A. Goldstein. 1996. The social and intellectual shaping of a new mathematical discipline: The role of the National Science Foundation in the rise of theoretical computer science and engineering. In R. Calinger, ed., *Vita Mathematica: Historical Research and Integration with Teaching*. Mathematical Association of America Notes Series.
- H. E. Bamford and C. N. Brownstein. 1986. National Science Foundation support for computer and information science and engineering. *Information Processing & Management*, 22(6). DOI: [10.1016/0306-4573\(86\)90095-6](https://doi.org/10.1016/0306-4573(86)90095-6).
- D. O. Belanger. 1998. *Enabling American Innovation: Engineering and the National Science Foundation*. Lafayette, IN: Purdue University Press.
- L. Berlin. 2017. *Troublemakers: Silicon Valley's Coming of Age*. New York: Simon & Schuster.
- C. Bernhardt. 2016. *Turing's Vision: The Birth of Computer Science*. Cambridge, MA: MIT Press.
- M. Blumenthal. May 1998. Federal government initiatives and the foundation of the information technology revolution: Lessons from history. *American Economic Review*, 8(2).
- A. Bousquet. 2008. Cyberneticizing the American war machine: Science and computers in the Cold War. *Cold War History*, 8(1). DOI: [10.1080/14682740701791359](https://doi.org/10.1080/14682740701791359).
- M. Campbell-Kelly. 2003. *From Airline Reservations to Sonic the Hedgehog: A History of the Software Industry*. Cambridge, MA: MIT Press.
- M. Campbell-Kelly. 2014. The history of computing in colour: A picture essay. *International Journal for the History of Engineering and Technology*, 84(2). DOI: [10.1179/1758120614Z.00000000044](https://doi.org/10.1179/1758120614Z.00000000044).
- M. Campbell-Kelly, W. Aspray, N. Ensmenger, and J. Yost. 2013. *Computer: A History of the Information Machine*. Boulder, CO: Westview, 3rd ed.
- P. Ceruzzi. 1986. An unforeseen revolution: Computers and expectations, 1935–1985. In *Imagining Tomorrow: History, Technology, and the American Future*. Cambridge, MA: MIT Press.

- P. Ceruzzi. 1989. Electronics technology and computer science, 1940–1975: A coevolution. *Annals of the History of Computing*, vol. 10. DOI: [10.1109/MAHC.1988.10036](https://doi.org/10.1109/MAHC.1988.10036).
- P. Ceruzzi. 2003. *A History of Modern Computing*. Cambridge, MA: MIT Press, 2nd ed.
- P. Ceruzzi. 2012. *Computing: A Concise History*. Cambridge, MA: MIT Press.
- J. Cohn. 2015. Transitions from analog to digital computing in electrical power systems. *IEEE Annals of the History of Computing*, 37(3). DOI: [10.1109/MAHC.2015.37](https://doi.org/10.1109/MAHC.2015.37).
- G. Con Diaz. 2012. Ownership and the history of American computing. Think Piece. *IEEE Annals of the History of Computing*, 34(2). DOI: [10.1109/MAHC.2012.29](https://doi.org/10.1109/MAHC.2012.29).
- J. W. Cortada. 1993. *The Computer in the United States: From Laboratory to Market, 1930–1960*. Armonk, N.Y.: M.E. Sharpe.
- J. W. Cortada. 2005. *The Digital Hand: How Computers Changed the Work of American Financial, Telecommunications, Media, and Entertainment Industries*. New York: Oxford University Press.
- J. W. Cortada. 2016. *All the Facts: A History of Information in the United States since 1870*. New York: Oxford University Press.
- J. W. Cortada. 2017. A history of information in the United States since 1870. *Information & Culture*, 52(1).
- D. Cox, E. Sandor, and J. Fron, eds. 2018. *New Media Futures: The Rise of Women in the Digital Arts*. Urbana, IL: University of Illinois Press.
- K. K. Curtis, N. C. Metropolis, W. G. Rosen, Y. Shimamoto, and J. N. Snyder. 1983. John R. Pasta, 1918–1981—An unusual path toward computer science. *Annals of the History of Computing*, 5(3). DOI: [10.1109/MAHC.1983.10078](https://doi.org/10.1109/MAHC.1983.10078).
- S. Dasgupta. 2018. *The Second Age of Computer Science: From Algol Genes to Neural Nets*. New York: Oxford University Press.
- G. E. Dawson. 1994. *The Digital Revolution*. Limited ed. Downers Grove, Ill.: GMN Corp.
- P. N. Edwards. 1996. *The Closed World: Computers and the Politics of Discourse in the Cold War*. Cambridge, MA: MIT Press.
- J. M. England. 1982. *A Patron for Pure Science: The National Science Foundation's Formative Years, 1945–1957*. Washington, DC: National Science Foundation. DOI: [10.1086/ahr/89.1.102](https://doi.org/10.1086/ahr/89.1.102).
- N. Ensmenger. 2010. *The Computer Boys Take Over: Computers, Programmers, and the Politics of Technical Expertise*. Cambridge, MA: MIT Press.
- C. L. Evans. 2018. *Broad Band: The Untold Story of the Women Who Made the Internet*. New York: Portfolio.
- P. L. Frana. 2018. Telematics and the early history of international digital information flows. *IEEE Annals of the History of Computing*, 40(2). DOI: [10.1109/MAHC.2018.022921442](https://doi.org/10.1109/MAHC.2018.022921442).
- T. Friedman. 2005. *Electric Dreams: Computers in American Culture*. New York: New York University Press.

- T. F. Gieryn. 1999. The U.S. Congress demarcates natural science and social science (twice). *Cultural Boundaries of Science: Credibility on the Line*. Chicago: University of Chicago Press.
- K. Gotkin. 2014. When computers were amateur. *IEEE Annals of the History of Computing*, 36(2). DOI: [10.1109/MAHC.2014.32](https://doi.org/10.1109/MAHC.2014.32).
- D. S. Greenberg. 1967. *The Politics of Pure Science*. Chicago: University of Chicago Press.
- D. S. Greenberg. 2001. *Science, Money, and Politics*. Chicago: University of Chicago Press.
- S. Greenstein. 2015. *How the Internet Became Commercial: Innovation, Privatization, and the Birth of a New Network*. Princeton, NJ: Princeton University Press.
- M. Grenia. 1997. *History of Computing: An Encyclopedia of the People and Machines that Made Computer History*. Lexicon Services.
- G. K. Gupta. 2007. Computer science curriculum developments in the 1960s. *IEEE Annals of the History of Computing*, 29(2). DOI: [10.1109/MAHC.2007.20](https://doi.org/10.1109/MAHC.2007.20).
- T. Haigh. 2010. Computing the American way: Contextualizing the early US computer industry. *IEEE Annals of the History of Computing*, 32(2). DOI: [10.1109/MAHC.2010.33](https://doi.org/10.1109/MAHC.2010.33).
- T. Haigh. 2019. *Exploring the Early Digital*. Springer. DOI: [10.1007/978-3-030-02152-8](https://doi.org/10.1007/978-3-030-02152-8).
- D. Hemmendinger. 2016. Two early interactive computer network experiments. *IEEE Annals of the History of Computing*, 37(3). DOI: [10.1109/MAHC.2015.44](https://doi.org/10.1109/MAHC.2015.44).
- T. Ivall. 1986. Fifty years of computer science. *Electronics and Wireless World*, vol. 92.
- D. J. Kevles. 1977. The National Science Foundation and the debate over postwar research policy, 1942–1945: A political interpretation of *Science—The Endless Frontier*. *Isis*, 68(5). DOI: [10.1086/351711](https://doi.org/10.1086/351711).
- D. L. Kleinman. 1995. *Politics on the Endless Frontier: Postwar Research Policy in the United States*. Durham, NC: Duke University Press.
- D. L. Kleinman and M. Solovey. 1995. Hot Science/Cold War: The National Science Foundation after World War II. *Radical History Review*, vol. 63. DOI: [10.1215/01636545-1995-63-111](https://doi.org/10.1215/01636545-1995-63-111).
- O. N. Larsen. 1992. *Milestones and Millstones: Social Science at the National Science Foundation, 1945–1991*. New Brunswick: Transaction Publishers.
- B. M. Leiner, et al. 1997. Origins of the Internet. In *A Brief History of the Internet*. Internet Society. DOI: [10.1145/1629607.1629613](https://doi.org/10.1145/1629607.1629613).
- S. W. Leslie. 1993. *The Cold War and American Science*. New York: Columbia University Press.
- M. Lomask. 1976. *A Minor Miracle: An Informal History of the National Science Foundation*. Washington, DC: National Science Foundation.
- S. Lukasik. 2011. Why the Arpanet was built. *IEEE Annals of the History of Computing*, 33(3). DOI: [10.1109/MAHC.2010.11](https://doi.org/10.1109/MAHC.2010.11).
- P. G. Lykos, M. G. Morgan, and F. W. Weingarten. 1973. Computer impact on society perspective on a new NSF initiative. *Proceedings of the ACM Annual Conference (ACM '73)*. New York, NY: ACM. DOI: [10.1145/800192.805738](https://doi.org/10.1145/800192.805738).

- R. F. Maddox. 1979. The politics of World War II science: Senator Harley M. Kilgore and the legislative origins of the National Science Foundation. *West Virginia History*, vol. 41.
- M. S. Mahoney. 2005. The history of computing(s). *Interdisciplinary Science Reviews*, 30(2). DOI: [10.1179/030801805x25927](https://doi.org/10.1179/030801805x25927).
- J. Mailland and K. Driscoll. 2017. *Minitel: Welcome to the Internet*. Cambridge, MA: MIT Press.
- G. T. Mazuzan. 1992. Good science gets funded: The historical evolution of grant making at the National Science Foundation. *Knowledge: Creation, Diffusion, Utilization*, vol. 14. DOI: [10.1177/107554709201400104](https://doi.org/10.1177/107554709201400104).
- G. T. Mazuzan. 1994. *The National Science Foundation: A Brief History*. Washington, DC: National Science Foundation.
- J. Minker. 2007. Forming a computer science center at the University of Maryland. *IEEE Annals of the History of Computing*, 29(1). DOI: [10.1109/MAHC.2007.5](https://doi.org/10.1109/MAHC.2007.5).
- J. Minker. 2007. Developing a computer science department at the University of Maryland. *IEEE Annals of the History of Computing*, 29(4). DOI: [10.1109/MAHC.2007.55](https://doi.org/10.1109/MAHC.2007.55).
- T. Misa. 2007. Understanding “How Computing Has Changed the World.” *IEEE Annals of the History of Computing*, 29(4). DOI: [10.1109/MAHC.2007.68](https://doi.org/10.1109/MAHC.2007.68).
- T. Misa. 2010. *Gender Codes: Why Women Are Leaving Computing*. Hoboken, NJ: Wiley-IEEE Computer Society.
- T. Misa. 2012. Computer science: Digital dawn. *Nature*, vol. 483. DOI: [10.1038/483032a](https://doi.org/10.1038/483032a).
- T. Misa. 2016. *Communities of Computing: Computer Science and Society in the ACM*. California: ACM Books. DOI: [10.1145/2973856](https://doi.org/10.1145/2973856).
- T. Misa. 2016. Computer security discourse at RAND, SDC, and NSA (1958–1970). *IEEE Annals of the History of Computing*, 38(4). DOI: [10.1109/MAHC.2016.48](https://doi.org/10.1109/MAHC.2016.48).
- T. Misa and J. Yost. 2016. *Fastlane: Managing Science in the Internet World*. Baltimore, MD: Johns Hopkins University Press.
- A. L. Norberg and J. W. O’Neill. 2000. *Transforming Computer Technology: Information Processing for the Pentagon, 1962–1986*. Baltimore, MD: Johns Hopkins University Press.
- G. O’Regan. 2012. *A Brief History of Computing*. London: Springer International.
- G. O’Regan. 2016. *Introduction to the History of Computing: A Computing History Primer*. Switzerland: Springer International.
- S. Pollack. 1982. The development of computer science. In S. Pollack, ed., *Studies in Computer Science*. Washington DC: Mathematical Association of America.
- W. E. Pugh and W. Aspray. 1996. Creating the computer industry. *IEEE Annals of the History of Computing*, 18(2). DOI: [10.1109/MAHC.1996.490112](https://doi.org/10.1109/MAHC.1996.490112).
- Z. Qin et al. 2018. *Fundamentals of Software Culture*. Singapore: Springer International.
- K. Redmond. 2000. *From Whirlwind to MITRE: The R&D Story of the Sage Air Defense Computer*. Cambridge, MA: MIT Press.

- M. Rees. 1987. The computing program of the Office of Naval Research. *Communications of the ACM*, vol. 30. DOI: [10.1145/30408.30415](https://doi.org/10.1145/30408.30415).
- E. D. Reilly. 2003. *Milestones in Computer Science and Information Technology*. Westport, CT: Greenwood Press.
- H. L. Resnikoff. 1980. Information science and technology at the National Science Foundation. *Information Processing and Management*, vol. 16. DOI: [10.1016/0306-4573\(80\)90013-8](https://doi.org/10.1016/0306-4573(80)90013-8).
- J. Rohde. 2017. Pax Technologica: Computers, international affairs, and human reason in the Cold War. *Isis*, 108(4). DOI: [10.1086/695679](https://doi.org/10.1086/695679).
- A. Roland and P. Shiman. 2002. *Strategic Computing: DARPA and the Quest for Machine Intelligence, 1983–1993*. Cambridge, MA: MIT Press.
- M. Rothenberg. 2010. Making judgments about grant proposals: A brief history of the merit review criteria at the National Science Foundation. *Technology and Innovation*, vol. 12. DOI: [10.3727/194982410x12895770313952](https://doi.org/10.3727/194982410x12895770313952).
- J. Rudolph. 2002. *Scientists in the Classroom: The Cold War Reconstruction of American Science Education*. New York: Palgrave.
- J. F. Sargent, Jr. and D. A. Shea. August 17, 2017. *Office of Science and Technology Policy (OSTP): History and Overview*. Washington, DC: Congressional Research Service. <https://fas.org/sgp/crs/mis/R43935.pdf>.
- K. Schaffel. 1991. *The Emerging Shield: The Air Force and the Evolution of Continental Air Defense 1945–1960*. Washington, DC: Office of Air Force History.
- M. Solovey. 2013. *Shaky Foundations: The Politics-Patronage-Social Science Nexus in Cold War America*. New Brunswick, NJ: Rutgers University Press.
- G. Stever. 2002. *In War and Peace: My Life in Science and Technology*. Washington, DC: Joseph Henry Press.
- J. Stine. 1992. Scientific instrumentation as an element of U.S. science policy: National Science Foundation support of chemistry instrumentation. In R. Bud and S. E. Cozzens, eds., *Invisible Connections: Instruments, Institutions, and Science*. Bellingham, Washington: SPIE.
- A. Tatnall, ed. 2010. *History of Computing: Learning from the Past*. Berlin: Springer-Verlag Berlin Heidelberg.
- A. Tatnall and B. Davey, eds. 2014. *Reflections on the History of Computers in Education: Early Use of Computers and Teaching about Computing in Schools*. Berlin: Springer International.
- J. E. Tomayko. 2000. *Computers Take Flight: A History of NASA's Pioneering Fly by Wire Project*. Washington, DC: NASA.
- C. Tozzi. 2017. *For Fun and Profit: A History of the Free and Open Source Software Revolution*. Cambridge, MA: MIT Press.

- W. F. Vogel. 2017. 'The spitting image of a woman programmer': Changing portrayals of women in the American computing industry, 1958–1985. *IEEE Annals of the History of Computing*, 39(2). DOI: [10.1109/MAHC.2017.14](https://doi.org/10.1109/MAHC.2017.14).
- J. Wang. 1995. Liberals, the progressive left, and the political economy of postwar America: The National Science Foundation debate revisited. *Historical Studies in the Physical and Biological Sciences*, vol. 26. DOI: [10.2307/27757758](https://doi.org/10.2307/27757758).
- J. Wang. 1999. *American Science in an Age of Anxiety: Scientists, Anticommunism, and the Cold War*. University of North Carolina Press.
- M. Warner. 2015. Notes on the evolution of computer security policy in the U.S. government, 1965–2003. *IEEE Annals of the History of Computing*, 37(2). DOI: [10.1109/MAHC.2015.25](https://doi.org/10.1109/MAHC.2015.25).
- S. Weinberger. 2017. *The Imagineers of War: The Untold History of DARPA, the Pentagon Agency that Changed the World*. New York: Alfred A. Knopf.
- M. R. Williams. 1985. *A History of Computing Technology*. Englewood Cliffs: Prentice-Hall.
- H. Wulforst. 1982. *Breakthrough to the Computer Age*. New York: Scribners.
- J. R. Yost. 2002. *A Bibliographic Guide to Resources in Scientific Computing, 1945–1975*. Westport, CT: Greenwood Press.
- J. R. Yost [Special Issue Guest Editor]. 2016. Computer security, Part II. *IEEE Annals of the History of Computing*, 38(4). DOI: [10.1109/MAHC.2016.44](https://doi.org/10.1109/MAHC.2016.44).
- J. R. Yost. 2017. *Making IT Work: A History of the Computer Services Industry*. Cambridge, MA: MIT Press.
- G. P. Zachary. 1999. *Endless Frontier*. Cambridge, MA: MIT Press.

Author Biographies

W. Richards Adrion



part-time, and visiting positions at a number of not-for-profit organizations, corporate labs, and universities here and abroad. He is a fellow of the ACM and AAAS.

W. Richards “Rick” Adrion is Professor Emeritus of Computer Science at the University of Massachusetts Amherst. He continues to be an active researcher in software engineering and computer science education, leading several large NSF projects during his career. He served as an NSF program director from 1976–1978 and 1980–1986, a division director from 2000–2002, and a CISE senior scientist in 1986 and 2002–2003. He held full-time positions at the University of Texas at Austin, Oregon State, and the National Bureau of Standards, as well as consulting,

William Aspray



a history of computing project at NSF in the early 1990s, and he has published three articles and a book about NSF computing activities (*Participation in Computing: The National Science Foundation’s Expansionary Programs*, Springer, 2016).

William Aspray is Professor of Information Science at the University of Colorado Boulder. He is trained as a historian of science and has spent most of his career studying the histories of computing, information, mathematics, and electrical engineering. He formerly taught at Harvard University, Indiana University Bloomington, University of Pennsylvania, University of Texas at Austin, Virginia Tech, and Williams College, and served in leadership positions at the Charles Babbage Institute, Computing Research Association, and IEEE History Center. He led

Peter A. Freeman



ACM, AAAS and IEEE.

Peter A. Freeman, a computer scientist (Carnegie Mellon University, 1970), was the seventh AD/CISE from 2002–2007 and DD/CCR from 1987–1989. Previously he held faculty positions at the University of California, Irvine and Georgia Tech, where he was Founding Dean of Computing in 1990. After NSF he returned to Georgia Tech as Emeritus Dean and Professor and was an active consultant with the Washington Advisory Group. While in academe he was an active researcher, educator, and participant in professional society activities. He is a fellow of the

Index

- A4RC program, 318
- AAAI (American Association for Artificial Intelligence), 30, 74
- Abbreviations and acronyms, 373–379
- Abstraction, data, 222–224
- Academic experimental computer science, need for, 48–55
- AccessComputing program, 318
- ACE (American Council on Education), 44–45
- Adrion, W. Richards, 27–28
 - academic experimental computer science, 48–55
 - ARPANET expansion, 278–279
 - background, 1974–1986 period, 28–29
 - CISE beginnings, 66–74
 - CISE structure, 88–91
 - CISE vision, 110
 - Computer Science Section, 1979–1984, 46–48
 - computing alliances, 317
 - computing and information research, 1974–1978, 35–39
 - computing research, 341
 - computing research and infrastructure, 33–34
 - cryptography with NSA, 43–46
 - CSNET, 55–60
 - directorate interface, 266
 - facilities, 41–43
 - innovators in computer science, 39–41
 - NSF computing research programs, 29–32
 - NSF departure, 93
 - OASC, 279, 355
 - OASC and NSFNET, 60–65
 - parallelism paper, 73
 - Senior Advisor, 130
 - Theorynet project, 278
- ADVANCE program, 139, 315
- Advanced Research Projects Agency (ARPA)
 - ARPANET. *See* ARPANET
 - computer research funding, 31, 347
 - funding limitations, 339
 - networking research, 133
 - Project MAC, 12
 - university funding, 40
- Advanced Scientific Computing (ASC), 90, 95, 97, 356
- Advisory Committee on Supercomputer Access, 61–62
- Advisory Panel on University Computing Facilities, 182–183
- AI (artificial intelligence)
 - AI lab, 41
 - research, 241–249
- ALERT system, 213
- Aloha channels, 238
- Alphard language, 222–224
- American Association for Artificial Intelligence (AAAI), 30, 74
- American Council on Education (ACE), 44–45

- American Recovery and Reinvestment Act, 150
- Application-specific integrated circuits (ASICs), 212
- AR (automated reasoning), 203
- Argus language, 225–226
- ARPA. *See* Advanced Research Projects Agency (ARPA)
- ARPANET
 - access, 48, 55–56, 64–65
 - CSNET with, 58, 64
 - decommissioning, 281
 - expansion, 278
 - Internet origin, 276–277
 - NSFNET protocols, 279–280
 - opportunities, 62
 - Sciencenet Phase 1, 63, 279
 - security concerns, 59
- Artificial intelligence (AI)
 - AI lab, 41
 - research, 241–249
- Arzberger, Peter
 - biography, 363
 - CISE leadership, 153–155, 346
 - Kurose supported by, 166
- ASC (Advanced Scientific Computing), 90, 95, 97, 356
- ASICs (application-specific integrated circuits), 212
- Atkins Report, 138, 309
- Atkinson, Richard
 - computer-based education, 337
 - computer science support, 46
 - experimental computer science research, 47
 - mathematical models of student learning, 15
- Atomic transactions, 226
- AUTODIN II network, 277
- Automated reasoning (AR), 203
- Automatic theorem proving, 242–246
- Bajcsy, Ruzena
 - accomplishments, 343–344
 - biography, 363
 - CISE leadership, 121–124
 - Information Technology Research, 265–268
- Baker Report, 5
- Bardon-Curtis Report, 33–34, 61
- BASIC (Beginner's All-purpose Symbolic Instruction Code) language, 16
- Basic Linear Algebra Subroutines (BLAS3), 219
- Bell, Gordon
 - biography, 363–364
 - Bloch credited by, 87
 - CISE creation, 341
 - CISE organization, 64–65, 68, 70–73, 88–91, 341
 - CISE vision, 110
 - computer hardware architect, 136
 - cyberinfrastructure efforts, 138
 - DIST advisory group, 9
 - Gordon Bell Prize, 95
 - high-performance computing, 307
 - management skills, 96
 - MOSIS service, 96
 - NCRI Division, 280
 - networking concept, 91, 95
 - networking program, 279
 - NSF role requests, 29, 66–67
 - NSF workshop, 32
 - OASC, 355
 - peer review importance, 96
 - research funding, 31
 - strategic and operational advice, 93
 - supercomputer centers, 96, 98, 138
 - U.S. computing leadership strategy, 94
 - university research, 339
- Bement, Arden
 - biography, 364
 - cyberinfrastructure activities, 148
 - Freeman guided by, 344
 - supercomputer center support, 307
- Bibliography, 381–387
- Big Data Research and Development Initiative, 157

- Biographies, 363–369
- BITNET
 - CSNET merger, 60
 - expansion strategy, 63, 279
 - gateways and connections, 59
 - TCP/IP services, 65
- BLAS3 (Basic Linear Algebra Subroutines), 219
- Bledsoe, Woodrow, artificial intelligence research by, 241–246
- Bloch, Eric
 - biography, 364
 - CISE funding, 69, 97
 - CISE organizational initiation, 90–91
 - CISE vision, 87
 - computing activities consolidation, 66–68
 - directorate development, 29
 - directorate organization, 70
 - high-performance computing, 34, 306–307, 340–341
 - new leadership initiative, 109–110
 - NSF Director, 47
 - NSF roles and support questions, 73
 - OASC, 355
 - peer review importance, 96
 - program consolidation, 88–89
 - program transfers, 71
 - science and research policy activities, 99–100
 - senior staff lunchroom, 39
 - supercomputer centers, 138
 - U.S. computing leadership strategy, 94
- Borchers, Robert, 266
- Boucher Amendment, 92, 101
- BPC (Broadening Participation in Computing) program, 140, 315–318
- BRAIN (Brain Research through Advancing Innovative Neurotechnologies), 157–158
- Branch-and-bound methods, 212
- Broad computer instruction, 189
- Broadband access, 150–151
- Broadening participation in computing, 313–321
- Broadening Participation in Computing (BPC) program, 140, 315–318
- Brownstein, Charles
 - biography, 364
 - CISE funding, 97
 - CISE long range planning, 73
 - CISE organization, 91
 - CISE vision, 110
 - CSSP transfer, 71
 - Executive Officer, 29, 72
 - Freeman guided by, 93
 - high-performance computing, 99–100, 108
 - information science research, 9
 - NSF computing activities consolidation, 66
 - OASC, 355
 - social and organizational impacts of computing, 89
- Browsers, 129, 285
- CAD (computer-aided design systems) research, 209
- CAHSI program, 318
- CAI (computer-aided instruction) systems in education, 185
 - support for, 15, 181
- CAITE program, 318
- CCC. *See* Computing Community Consortium (CCC)
- CCF (Computing and Communication Foundations) Division, 137, 156, 271
- CCR (Computer and Communications Research) Division, 72, 90, 93, 355
- CDI program. *See* Cyber-enabled Discovery & Innovation (CDI) program
- CE21 (Computing Education for the Twenty-First Century) program, 315
- CEDAR system, 219
- Centers of Excellence program, 49

- CER. *See* Coordinated Experimental Research (CER) program
- CHA (CISE History Archive), 371–372
- Chern, Bernard
 CISE reorganization, 71, 73
 MIPS director, 356
 MOSIS VLSI fabrication facility, 17
- Chien, Y. T.
 Intelligent Systems program, 46
 IRIS director, 356
- Children’s Television Workshop, 14
- CIC (Coordination and Information Center), 58–59
- Ciment, Mel
 ASC director, 356
 biography, 364–365
 CISE implementation plan, 94–95
 CISE leadership, 102
 deputy AD, 342
- CIRI (Clean-slate Internet Re-invention Initiative), 290–293
- CISE. *See* Computer and Information Science and Engineering (CISE) Directorate
- CISE History Archive (CHA), 371–372
- CISE Organizational Review Committee (CORC), 104–105, 357–358
- Civil rights movements, 313–317
- Clean-slate Internet Re-invention Initiative (CIRI), 290–293
- CLU language, 224–225
- Cluster Exploratory (CluE), 152
- Clusters in CLU language, 224–225
- Clusters system in CISE, 131–132
- CMOS (complementary metal-oxide-semiconductor) technology, 212–213
- Collaborative relations, 93–94
- Colwell, Rita
 biography, 365
 CISE funding, 121, 123, 343
 computation importance, 120
 computing research successes, 142
 ITR, 123, 141, 265
 national security needs, 126–127
- Commercial Internet, 282–286
- Committee on the Status of Women in Computing Research, 316
- Compilers, support for, 220–221
- Complementary metal-oxide-semiconductor (CMOS) technology, 212–213
- Computational parallelism, 95
- Computational thinking, 153
- Computer-aided design systems (CAD) research, 209
- Computer-aided instruction (CAI) systems in education, 185
 support for, 15, 181
- Computer and Communications Research (CCR) Division, 72, 90, 93, 355
- Computer and Information Science and Engineering (CISE) Directorate
 analysis, 108–110, 323–333
 Arzberger leadership, 153–155
 Bajcsy leadership, 121–124
 beginnings, 66–74
 computer science new initiatives
 introduction, 119–121
 Crawford leadership, 147–148
 cybersecurity research, 136–137
 education and outreach, 138–141
 Freeman leadership, 124–140
 GENI Project, 292–293
 Habermann leadership, 100–102
 Hartmanis leadership, 105–107
 history archive, 371–372
 Iacono leadership, 161–162
 information science research, 9
 initial actions, 91–92
 initial structure and leadership, 88–89
 ITR program, 265–271
 Jahanian leadership, 155–161
 Kurose leadership, 162–166
 leadership changes, 92–97
 networking research, 133–136
 oral histories list, 361–362
 organizational initiation, 89–91
 participation in computing, 313–321

- reorganization, 128–132
- Research Infrastructure program, 53
- software fundamentals program, 136
- supercomputer centers and cyberinfrastructure, 138
- transition to routine operation, 97–107
- Wing leadership, 148–153
- Wulf leadership, 97–100
- Young leadership, 103–105
- Computer architecture
 - advanced computation, 307–308
 - CER awards, 50
 - facilities programs, 41–42
 - graphics research, 251–252
 - numerical research, 233
 - panel, 34
 - research, 213–220, 340
 - support for, 204–208
 - XPS program, 159
- Computer circuits research, 208–213
- Computer education
 - conclusions, 192–193
 - division of labor, 186–188
 - Great Society support, 184–186
 - national computer science community, 188–190
 - need for computer scientists, 190–192
 - support for, 181
 - support for in 1950s, 182–184
- Computer engineering, theoretical, 234–241
- Computer graphics
 - contributions to, 203
 - research, 249–254
- Computer labor, 189–190
- Computer Science Advisory Committee, 48–49
- Computer Science and Economics program, 151
- Computer Science and Engineering (CS&E) Section, 18–19, 354
- Computer Science and Engineering Research Study (COSERS), 31–32, 35
- Computer Science and Telecommunications Board (CSTB), 33
- Computer science applications, 208
- Computer Science Collaboration Project (CSCP), 320
- Computer science education, 181, 189
- Computer science expansion (1986–1998), 87–88
 - analysis, 108–110
 - CISE initial actions, 91–92
 - CISE initial structure and leadership, 88–89
 - CISE leadership changes, 92–97
 - CISE organizational initiation, 89–91
 - CISE transition to routine operation, 97–107
- Computer science innovators funding, 39–41
- Computer Science Network (CSNET)
 - ARPANET sites for, 62
 - BITNET merger, 29
 - budget for, 50
 - computer research group connections, 279
 - creation, 55–60
 - funding, 49, 55–60, 65
 - management of, 28
 - network of networks approach, 64
 - operational network, 276
 - SRPS impact, 52
 - supercomputers, 63
 - support for, 40, 47, 95
 - Theorynet project, 43, 278
- Computer science new initiatives (1999–2006), 119–121
 - CISE leadership by Bajcsy, 121–124
 - CISE leadership by Freeman, 124–140
 - cybersecurity research, 136–137
 - education and outreach, 138–141
 - networking research, 133–136
 - observations, 141–142
 - software fundamentals program, 136
 - supercomputer centers and cyberinfrastructure, 138

- Computer Science Principles course, 317–318
- Computer science research, 17–20
- Computer Science Research Equipment program, 42
- Computer Science Section (CSS)
 - 1974–1978, 35–39
 - 1979–1984, 46–48
 - careers launched by, 40
- Computer scientists, need for, 190–192
- Computer theory, support for, 207
- Computers in science education, 181
- Computing and Communication Foundations (CCF) Division, 137, 156, 271
- Computing and information research, 1974–1978, 35–39
- Computing centers
 - early, 188–189
 - facilities programs, 176–178
- Computing Community Consortium (CCC)
 - creation, 293
 - function and importance, 156–157
 - funding, 156, 166
 - GENI project, 329
 - mission, 150
 - OSTP initiative, 157
 - research needs identification, 345
 - workshops, 165
 - XPS program, 159
- Computing Education for the Twenty-First Century (CE21) program, 315
- Computing infrastructure
 - demand for, 9–13
 - importance, 33–34
- Computing organization charts, 351–359
- Computing research
 - artificial intelligence, 241–249
 - computer architecture, 213–220
 - computer circuits, 208–213
 - conclusions, 254
 - funding, 38–39
 - graphics, 249–254
 - importance, 33–34
 - numerical analysis, 227–234
 - overview, 203–208
 - programs for, 29–32
 - software, 220–226
 - theoretical computer engineering, 234–241
- CONDUIT sharing network, 14–15, 187
- Coordinated Experimental Research (CER) program
 - academic experimental computer science, 48–55
 - ARPANET opportunities, 62
 - budget for, 49–50
 - computing facilities funding, 340
 - experimental research in colleges and universities, 75
 - grants, 41–42
 - impacts, 52
 - institutions benefited from, 53
 - management of, 28
 - networks, 278
 - overview, 51
 - purpose, 48
 - recommendations for, 32, 339
 - support for, 47
- Coordination and Information Center (CIC), 58–59
- CORC (CISE Organizational Review Committee), 104–105, 357–358
- Cordova, France, 163
- Cornell Theory Center, 306
- COSERS (Computer Science and Engineering Research Study), 31–32, 35
- Course Content Improvement program, 183–184
- Cowell, Wayne, numerical analysis by, 227–232
- CPATH (Pathways to Revitalized Undergraduate Computing Education), 317
- CRA-W/CDC Alliance program, 318
- Crawford, Deborah
 - background, 127

- biography, 365
- changes aided by, 316
- CISE leadership, 147–148
- CISE move, 130
- communication plan, 131
- education and workforce insights, 331
- GENI project, 290–292
- networking research, 135
- Cross-cutting programs, 151
- Cryptography
 - Encryption Wars, 102
 - NSA interactions, 43–46
 - research, 238–241
- CS Theory program, 137
- CS&E (Computer Science and Engineering)
 - Section, 18–19, 354
- CS10K effort, 317, 319
- CSCP (Computer Science Collaboration Project), 320
- CSCT (Continuous-Space-Continuous-Time) machines, 215
- CSDT (Continuous-Space-Discrete-Time) machines, 215
- CSNET. *See* Computer Science Network (CSNET)
- CSS (Computer Science Section)
 - 1974–1978, 35–39
 - 1979–1984, 46–48
 - careers launched by, 40
- CSTB (Computer Science and Telecommunications Board), 33
- Csuri, Charles, graphics research by, 253–254
- Cuny, Jan
 - CISE reorganization, 359
 - CS10K program, 317
 - diversity and inclusion, 160, 316
 - participation in computing, 140, 315, 331, 345–346
 - women in computing, 139
- Curtis, Kent
 - academic computer science, 32
 - ARPANET expansion, 278
 - ARPANET opportunities, 62
 - Bardon-Curtis Report, 34, 61, 340
 - biography, 365–366
 - CCR director, 355
 - CISE formation, 339
 - CISE long range planning, 73
 - CISE structure, 88, 90
 - Computer Science Section, 35
 - computing program organization, 66
 - Coordinated Experimental Research program, 55
 - COSERS, 31
 - cryptographic research, 46
 - experimental computer science research, 47
 - Hopcroft-Kennedy Report, 33
 - Institutional Computing Services Section, 18
 - SRPS budget, 49
 - strategic and operational advice, 93
 - submission encouragement, 36
 - supercomputer services, 305
 - weekly meetings, 38
- Cyber-enabled Discovery & Innovation (CDI)
 - program
 - computational thinking, 153
 - data science initiative, 152, 346
 - funding, 141
 - goals, 345
 - importance, 149
- Cyber-Physical Systems program, 151
- Cyberinfrastructure
 - CISE program, 162–163
 - high-performance computing, 122–123
 - and supercomputer centers, 138
 - support for, 309–310
- Cyberinfrastructure Report, 128
- Cybersecurity R&D Act, 134, 136–137
- Cybersecurity research, 136–137
- CyberSEES (Cyber-Innovation for Sustainability Science and Engineering), 158–159
- DARPA. *See* Defense Advanced Research Projects Agency (DARPA)

- Data abstraction, 223–224
- Data types, 222–224
- DCR. *See* Division of Computer Research (DCR)
- Defense Advanced Research Projects Agency (DARPA)
 - academic experimental computer science, 48
 - ARPANET, 62, 64–65
 - artificial intelligence, 241
 - BioComputational Systems program, 266
 - BRAIN initiative, 158
 - broad agency announcement, 264
 - CER program, 50–52
 - CISE research budget, 71–72
 - computer graphics research, 249
 - computer network technology, 133–134
 - computer science support, 149
 - computing research support, 347
 - CSNET, 56–57
 - Cyber Trust proposals, 271
 - data science initiative, 152
 - Federal Big Data Research and Development Initiative, 157
 - high-performance computing, 153
 - Illinois Center for Supercomputing Research, 220
 - Liskov research funding, 226
 - long-distance networking, 286
 - MOSIS VLSI fabrication facility, 17, 69, 96
 - research areas, 203
 - SRI funding, 248
 - Strategic Computing Initiative, 100
 - Test Net, 284
- Defense Department
 - programs, 52
 - research divestment, 204
- Dependence graphs, 217
- Digital Computer Needs in Universities and Colleges report, 12
- Digital Encryption Standard (DES), 43, 239–240
- Digitalizers, 215
- Directorates
 - organizational charts, 351–359
 - overview, 8–9
- Distributed computer systems, 225
- Distributed Terascale Facility (DTF), 308–309
- Diversity and inclusion. *See* Participation in computing
- Division of Computer Research (DCR)
 - Adrion directorship, 29
 - CCR responsibilities, 72
 - CISE consolidation, 70
 - computer science support, 31
 - computing and information research, 35
 - creation, 47
 - Mathematical and Computer Sciences Division, 339
 - from OCA, 8, 20, 27, 31, 186
 - Pasta directorship, 11
 - reorganization, 355
 - Special Projects program, 43
 - transfer proposal, 71
- Division of Information Science and Technology (DIST), 8–9, 355
- DSCT (Discrete-Space-Continuous-Time) machines, 215
- DSDT (Discrete-Space-Discrete-Time) machines, 215–216
- DTF (Distributed Terascale Facility), 308–309
- Eden Project, 50–51
- Education
 - computer. *See* Computer education
 - computers in, 14–17
 - issues, 154
 - strengthening, 138–141
- Education and Human Resources (EHR) Directorate, 158, 161, 275
- EES (Escrowed Encryption Standard), 102
- EIA (Experimental and Integrative Activities), 29, 130, 266
- EISPACK package, 230

- ELLPACK package, 233
- Encryption
 - Encryption Wars, 102
 - NSA interactions, 43–46
 - research, 238–241
- Energy simulation, 251
- Engineering (ENG) Directorate
 - CISE organization, 70–71, 90
 - description, 37
 - funding, 123
- Escrowed Encryption Initiative, 102
- Escrowed Encryption Standard (EES), 102
- ETF (Extensible Terascale Facility), 309
- Executive Order 10807, 351
- Experimental and Integrative Activities (EIA), 29, 130, 266
- Exploring Parallelism and Scalability (XPS) program, 159
- Extensible Terascale Facility (ETF), 309
- Facilities programs
 - 1959 to 1986, 175–176
 - 1970s and 1980s, 180–181
 - instructional use of computers, 178–180
 - overview, 41–43
 - university computing centers, 176–178
- Federal Coordinating Committee for Science, Engineering, and Technology (FCCSET), 93, 307
- Federal government
 - academic research funding, 3–4
 - budget cycle, 126
 - computing infrastructure support, 9–13
 - university computing center funding, 177
- Federal High-Performance Computing program, 33, 94–95
- Federal support for computer science, 149
- Fellowships and traineeships in computer science, 183
- Fermi-Pasta-Ulam (FPU) problem, 11
- FORTTRAN language, 222, 229–231
- Freeman, Peter
 - AD/CISE memo, 323–333
 - CCR director, 355
 - CISE Implementation Plan, 95
 - CISE leadership, 124–140
 - CISE reorganization, 128–132, 270, 358–359
 - cyberinfrastructure programs, 148
 - cybersecurity research, 136–137
 - diversity in computing, 316
 - education and outreach, 138–141
 - GENI project, 289–293
 - high-performance computing, 344
 - networking research, 133–136, 288–289
 - outreach to women and underrepresented minorities, 345
 - research emphasis, 315
 - software fundamentals program, 136
 - supercomputer centers and cyberinfrastructure, 138
- Funding a Revolution* report, 3
- FUNPACK package, 230–231
- Fuzzy sets, 236–237
- General Problem Solver, 246
- GENI project. *See* Global Environment for Networking Innovations (GENI)
- Georgia Computes program, 318
- Geosciences Directorate (GEO), 266–267
- Gianchandani, Erwin
 - biography, 366
 - Kurose reliance on, 166
- Gigabit Network Testbed Initiative, 286
- Global Environment for Networking Innovations (GENI)
 - conceptual design (2004–2006), 291–293
 - implementation (2007–2016), 293–294
 - initiation (2004–2006), 289–290
 - origin (2004), 290–291
 - purpose, 149–150
- Google search engine, 103, 285
- Gordon Bell Prize, 95
- Government Performance and Results Act (GPRA), 268
- Grand challenges, 34

- Graphics research, 249–254
- Graphs, 217
- Great Society education initiatives, 181, 184–186
- Greenberg, Donald, graphics research by, 249–253

- Habermann, A. Nico
 - biography, 366
 - CISE leadership, 100–102, 342
- Hartmanis, Juris
 - biography, 366
 - CISE leadership, 105–107, 342–343
 - CISE reorganization, 129, 134
 - computer research, 207
 - CORC, 357–358
 - Reorganization Working Group, 29
- Haworth, Lee, 352–354
- Hayes Report, 308
- Health IT initiative, 151
- Hedges, Harry, 28, 50
- Hellman, Martin, 237–241
- High-performance computation (HPC),
 - 303–304
 - 1950–1954, 304
 - 1955–1983, 304–306
 - 1984–1991, 306–308
 - 1992–2000, 308
 - 2000–2004, 308–309
 - 2005-present, 309–310
 - computational parallelism, 95
- High Performance Computing Act (HPCA), 263, 307
- High Performance Computing and Communications (HPCC) Act, 95, 100–101, 264
- High Performance Computing and Communications Initiative (HPCCI), 263
- Hopcroft-Kennedy Report, 33
- HPC. *See* High-performance computation (HPC)
- HPCA (High Performance Computing Act), 263, 307

- HPCC (High Performance Computing and Communications) Act, 95, 100–101, 264
- HPCCI (High Performance Computing and Communications Initiative), 263
- Human-Centered Systems, 104
- Human simulation type theorem proving systems, 243
- Human systems, 236
- Hybrid architecture, 216

- Iacono, Suzi
 - background, 330
 - biography, 366
 - CISE leadership, 161–162, 346
 - Jahanian helped by, 156, 159
 - NITRD activities, 152
 - strategic plan for international activities, 155
- IC (Intelligence community), 126–127
- IEEE, 44, 190
- IGERT (Integrated Graduate Education and Research Training) program, 139
- II-LS (Large-Scale) grants, 53
- II-SS (Small-Scale) grants, 53
- ILLIAC IV Supercomputer, 211, 217
- ILLIP (Illinois Integer Programming) program, 211–212
- IMP (Interface Message Processor), 276
- Implication method, 244
- IMPLY theorem prover, 244–245
- INCLUDES (Inclusion across the Nation of Communities of Learners) initiative, 164
- Independent Research and Development (IR&D), 28, 47
- INFEWS (Innovations at the Nexus of Food, Energy and Water Systems), 164
- Information Technology for the Twenty-First Century (IT2) federal initiative, 263–265
- Information Technology Research (ITR) program
 - CDI program, 149

- community cooperation, 329
- computer science initiative, 343
- cyberinfrastructure, 130
- funding, 123–124, 128–129, 131, 141
- high-performance computing, 308
- importance, 149
- ITR Report, 308–309
- ITWF research, 316
- overview, 263–272
- Information Technology Workforce (ITWF) program, 315–316
- Information theory, 238
- Innovations at the Nexus of Food, Energy and Water Systems (INFEWS), 164
- Innovators in computer science funding, 39–41
- Institutes for Science, Mathematics and Engineering Teachers program, 183
- Institutional grants program, 11–12
- Instructional use of computers facilities programs, 178–180
- Integrated Digital Library Project, 285
- Integrated Graduate Education and Research Training (IGERT) program, 139
- Intelligence community (IC), 126–127
- Intelligent Systems program, 36–37
- Interactive learning systems, 187
- Interface Message Processor (IMP), 276
- Internal issues in AD/CISE, 331–332
- International emphasis in CISE, 155
- International Traffic in Arms Regulations (ITAR) Act, 44
- Internet, 129
 - ARPANET, 276
 - browsers, 129
 - commercial, 282–286
 - CSNET, 57–60
 - management, 101–102
 - network support, 96
 - NSFNET, 100–101, 109, 280–281
 - protocols, 63
- Internet Protocol (IP), 277
- Invention Secrecy Act, 44
- IR&D (Independent Research and Development), 28, 47
- ITAR (International Traffic in Arms Regulations) Act, 44
- ITPACK packages, 233
- ITR program. *See* Information Technology Research (ITR) program
- ITWF (Information Technology Workforce) program, 315–316
- Jahanian, Farnam
 - biography, 366–367
 - CISE leadership, 154–161
 - CRII program, 166
 - federal investment, 163
- John von Neumann Center (JvNC), 63–64, 306
- Josephson junctions, 208–209
- Karplus, Walter, computer architecture research by, 213–217
- Keys in cryptography, 239–240
- Kilgore Committee, 3
- Knapp, Edward, 47, 61
- Knapsack problem, 241
- Kuck, David, computer architecture research by, 217–220
- Kurose, James
 - biography, 367
 - CISE leadership, 162–166, 346
- Lane, Neal
 - Advisory Committee on Supercomputer Access, 61
 - biography, 367
 - blue-ribbon panel, 101
- Languages, support for, 220–222
- Large-Scale (II-LS) grants, 53
- Lax Report, 34, 60–61, 305
- Lehmann, John, 17–18
- Lesk, Michael, 266
- LINPACK package, 231–232
- Liskov, Barbara, software research by, 220–226

- LISP language, 247
- Logic circuitry, 210–213
- Logo language and project, 16, 185
- Machine learning, 41
- Major Research Equipment and Facilities Construction (MREFC) program
 - funding, 150
 - NSB rules, 163
- Mansfield Amendment, 204
- Massey, Walter E., 99
- Mathematical software, 229
- Mathematics and Physical Sciences (MPS) Directorate
 - advisory report, 191
 - CISE reorganization, 68, 70, 90
 - computer science research, 175
 - computing research, 27
 - DCR in, 29
 - emphasis changes, 266
 - leadership, 47, 89
 - optical and wireless technology developments, 275
 - Probability and Statistics program, 153
 - SaTC program, 158
- McElroy, William, 180
- Memex, 276
- Memory hierarchy management, 219
- MERIT network, 280–282
- Micro world environments, 16
- Military Authorization Act, 204
- Minorities
 - programs, 314
 - regional centers, 13
- Mosaic browser, 101, 285
- MPS. *See* Mathematics and Physical Sciences (MPS) Directorate
- MREFC (Major Research Equipment and Facilities Construction) program
 - funding, 150
 - NSB rules, 163
- Muroga, Saburo, computer circuit research by, 208–213
- National Academy of Engineering (NAE), 99
- National Center for Supercomputer Applications (NCSA), 101
- National Center for Women & Information Technology (NCWIT), 140, 316
- National computer science community, 188–190
- National Defense Education Act (NDEA), 5, 181, 183, 351–352
- National Girls Collaborative Project (NGCP), 320
- National LambdaRail (NLR), 287–288
- National Partnership for Advanced Computational Infrastructure (NPACI), 154, 308
- National Research Council study on the future of computing, 106
- National Robotics Initiative (NRI), 151, 157, 165
- National Science Foundation Network (NSFNET)
 - backbone service, 92, 281, 341
 - CISE reorganization, 68
 - commercialization, 98, 101, 105, 133, 282–286
 - creation, 340
 - development, 91
 - high-speed networks, 95, 100
 - Internet origin, 109
 - leadership, 279–280, 342
 - MERIT network, 179
 - national vs. supercomputer network, 72 and Office of Advanced Scientific Computing, 60–65
 - operational network, 276
 - proposal and the initial ideas, 34
 - regional networks, 15, 92
 - service areas, 90
 - supercomputer centers, 99
 - support for, 97
 - Theorynet project, 40, 43, 278
- National security, 126–127
- National Security Agency (NSA)
 - cryptography, 19, 43–46, 238, 240

- CS Theory program, 102
- NATS project, 208, 230–231
- NATS II project, 230–231
- NCIH (North Carolina Information Highway), 286
- NCSA (National Center for Supercomputer Applications), 101
- NCWIT (National Center for Women & Information Technology), 140, 316
- NDEA (National Defense Education Act), 5, 181, 183, 351–352
- Netscape Navigator browser, 101
- Networking research and deployment, 275–277
 - commercial Internet, 282–286
 - computer science new initiatives, 133–136
 - conclusions, 294
 - facilities programs, 43
 - GENI implementation (2007–2016), 293–294
 - GENI origin (2004), 290–291
 - GENI Project (2004–2006), 289–290
 - major science effort support, 92
 - NSF public networking (1980–1995), 278–282
 - NSF-supported research on networks (1950–1980), 277–278
 - research (2002–2004), 288–289
 - research after NSFNET (1996–2001), 286–288
- New Directions in Cryptography report, 44
- Next-Generation Internet (NGI) initiative, 264, 287
- NGCP (National Girls Collaborative Project), 320
- NLR (National LambdaRail), 287–288
- North Carolina Information Highway (NCIH), 286
- NPACI (National Partnership for Advanced Computational Infrastructure), 154, 308
- NRI (National Robotics Initiative), 151, 157, 165
- NSA (National Security Agency)
 - cryptography, 19, 43–46, 238, 240
 - CS Theory program, 102
- NSF programs. *See* National Science Foundation (NSF) programs
- NSFNET. *See* National Science Foundation Network (NSFNET)
- Numerical analysis, 203, 227–234
- NYSERNet, 280–281
- OAC (Office of Advanced Computing), 132
- OASC (Office of Advanced Scientific Computing)
 - creation, 355
 - and NSFNET, 60–65
- Object-oriented languages, 225–226
- OCA. *See* Office of Computing Activities (OCA)
- OCI (Office of Cyberinfrastructure), 271, 346, 359
- OD (Office of the Director), 130–132, 327, 332
- Office of Advanced Computing (OAC), 132
- Office of Advanced Scientific Computing (OASC)
 - creation, 355
 - and NSFNET, 60–65
- Office of Computing Activities (OCA)
 - advisory committee, 184
 - Computer Innovations in Education Section, 16
 - computer research, 204
 - computer science programs, 191
 - computer science research, 17–20
 - DCR replacement, 339
 - education program, 182, 192
 - establishment, 12, 14, 181–182, 184, 353–354
 - funding by, 185–186
 - instructional use of computers, 178–180
 - national and regional research networks, 277
 - need for computer scientists, 190–192
 - regional centers, 13, 178–179

- Office of Computing Activities (OCA)
 - (continued)*
 - reorganization, 354
 - support by, 305, 338
- Office of Cyberinfrastructure (OCI), 271, 346, 359
- Office of Science and Technology Policy (OSTP)
 - high-performance computing, 94
 - initiatives, 157, 346
 - Program Plan, 307
- Office of Science Information Service (OSIS)
 - grants, 5–7
- Office of Scientific Information (OSI), 4–5
- Office of Scientific Research and Development (OSRD), 3
- Office of the Director (OD), 130–132, 327, 332
- Omega networks, 218–219
- Operating systems, support for, 220–221
- Organizational charts, 351–359
- OSI (Office of Scientific Information), 4–5
- OSIS (Office of Science Information Service)
 - grants, 5–7
- OSRD (Office of Scientific Research and Development), 3
- OSTP (Office of Science and Technology Policy)
 - high-performance computing, 94
 - initiatives, 157, 346
 - Program Plan, 307
- Outreach, strengthening, 138–141
- PACI (Partnerships for Advanced Computational Infrastructure), 308
- Paging, 219
- Panel on Computers in Higher Education, 184
- Parallelism in computing, 217–220
- Participation in computing
 - BPC program, 140, 315–318, 345
 - broadening, 313–321
 - CS10K effort, 319
 - introductory computer science courses, 318–319
- NGCP, 320
- programs, 314–315
- Supercomputer Centers program, 314
- virtual workshop, 315–316
- women’s rights and civil rights movements, 313–314
- Partnerships for Advanced Computational Infrastructure (PACI), 308
- Pascal language, 222
- Pasta, John
 - academic computer science, 32
 - background, 11
 - biography, 367
 - computer science research, 21
 - COSERS, 31
 - DCR leadership, 186
 - MCS leadership, 35, 355
 - OCA budget, 180
 - OCA leadership, 339
 - OCA reorganization, 19–20, 354
 - SRPS budget, 49
 - submission encouragement, 36
 - weekly meetings, 38
- Pathways to Revitalized Undergraduate Computing Education (CPATH), 317
- Peer review system, 127
- Phonetnet system, 56–57
- Pierce Report, 14, 181, 338
- PITAC (President’s Information Technology Advisory Committee), 101, 263–265, 315
- Pittsburgh Supercomputing Center (PSC), 306
- PlanetLab, 287
- PLATO (Program Logic for Automatic Teaching Operations) project, 14–16, 185
- President’s Information Technology Advisory Committee (PITAC), 101, 263–265, 315
- President’s Science Advisory Committee (PSAC), 184, 305

- Program analyzers, 218
- Program for Women and Girls, 314
- Program Logic for Automatic Teaching Operations (PLATO) project, 14–16, 185
- Project HARP, 226
- Project Mercury, 226
- Project Thor, 226
- Prolog language, 248
- PSAC (President's Science Advisory Committee), 184, 305
- PSC (Pittsburgh Supercomputing Center), 306
- Public key cryptography, 44, 239–240
- Public Law 81–507, 4

- RANN (Research Applied to National Needs) program, 180, 208
- Raphael, Bertram, 246–249
- Rawstorage data types, 223
- Ray-tracing in graphics, 251–252
- Reduced Instruction Set Computer (RISC) processors, 219
- Regional centers, 178–179
- Regional networks, 92, 179–180
- Research Applied to National Needs (RANN) program, 180, 208
- Resolution principle, 243–244
- Review process, 36–38
- RISC (Reduced Instruction Set Computer) processors, 219
- Robotics
 - National Robotics Initiative, 151, 157, 165
 - Shakey project, 247–248
- Role of the NSF in Supporting Cryptological Research report, 45–46
- Rosser Report, 12, 338
- RSA algorithm, 44, 240–241

- S&AS (Smart and Autonomous Systems) program, 165
- S&CC (Smart and Connected Communities) program, 165
- Sabre network, 276
- SAGE systems, 276
- San Diego Supercomputer Center (SDSC), 306, 308
- SaTC (Secure and Trustworthy Cyberspace) program, 150, 158
- SBE (Social, Behavioral, and Economic Sciences) Directorate, 151, 153, 158
- Science and Technology Centers (STC) program, 51–52, 91–92
- Science and Technology Equal Opportunity Act, 313
- Science foundation, 4
- Science information (1950s to 1980s), 4–9
- Science of Design (SOD) program, 330
- Sciencenet, 63, 279
- SDSC (San Diego Supercomputer Center), 306, 308
- Search engines, 284–285
- Secure and Trustworthy Cyberspace (SaTC) program, 150, 158
- Set-Var theorem prover, 245–246
- Shakey project, 247–248
- Shaw, Mary, software research by, 220–226
- Shielding term removal (STR) operation, 245
- Shuffle exchange networks, 219
- Silo-ing, 151
- SKIPJACK algorithm, 102
- Slaughter, John, 46–47, 355
- Small-Scale (II-SS) grants, 53
- Smart and Autonomous Systems (S&AS) program, 165
- Smart and Connected Communities (S&CC) program, 165
- Social and Computational Systems (SOCS) program, 151
- Social, Behavioral, and Economic Sciences (SBE) Directorate, 151, 153, 158
- SOD (Science of Design) program, 330
- Software crisis, 190, 221
- Software Engineering program, 36–37
- Software fundamentals program, 136
- Software research, 220–226

- Software Systems Science program, 36–37
- SOR (successive overrelaxation) method, 232–233
- Southeastern Universities Research Association network (SURAnet), 280
- Special Interest Committee on Numerical Mathematics, 229
- Special Projects program, 36–37
- Specialists vs. scientists, 7
- Sputnik launch, 4, 181, 183
- Stanford Research Institute Problem Solver (STRIPS), 247–248
- STAR TAP provider, 287
- STARS program, 318
- State-space theory of linear systems, 236
- STC (Science and Technology Centers) program, 51–52, 91–92
- Stever, Guy, 46, 197
- STR+VE theorem prover, 245–246
- Strawn, George
 - biography, 368
 - CIO, 125, 135
 - CISE leadership, 124
 - networking research, 134
- STRIPS (Stanford Research Institute Problem Solver), 247–248
- Structured programming, 222
- Successive overrelaxation (SOR) method, 232–233
- Supercomputer centers
 - broadening participation, 313–314 and cyberinfrastructure, 138
 - high-performance computing, 306–307
 - importance, 122
 - support for, 96, 98
- Supercomputers, 33
- SURAnet (Southeastern Universities Research Association network), 280
- Suresh, Subra
 - biography, 368
 - CISE reorganization, 160
 - Jahanian supported by, 156, 346
- Switching theory, 210
- TCP (Transmission Control Protocol), 277
- TCS (Theoretical Computer Science) program, 36–37, 40
- Themes, 131
- Theoretical computer engineering, 234–241
- Theoretical Computer Science (TCS) program, 36–37, 40
- Theorynet project
 - CER initiative, 28
 - funding, 133, 340
 - public networking, 278
 - Telemail participation, 43
- TICCIT (Time-shared Interactive Computer Controlled Information Television), 16–17, 185–186
- Time-shared Interactive Computer Controlled Information Television (TICCIT), 16–17, 185–186
- TOOLPACK package, 231–232
- Transmission Control Protocol (TCP), 277
- Triangle Universities Computation Center, 179
- Turing Award, 39–40
- Two-way interactive televisions and computers, 187
- U.S. computing leadership in computing, 94
- U.S. Ignite program, 157
- Undergraduate Instructional Development program, 187
- University computing centers facilities programs, 176–178
- USERNET, 62
- very high-speed Backbone Network Service (vBNS)
 - origin, 281
 - provide solicitation, 284
 - universities supported by, 286–287
- Very-large-scale integrated circuit (VLSI) technology, 211

- von Neumann, John
 - Advisory Panel on University Computing Facilities, 10, 182, 338
 - Conference on High-Speed Computing in Meteorology and Oceanography, 11
 - internal economy of computers, 228
- von Neumann architectures, 213
- Waterman, Alan
 - biography, 368
 - computer use growth study, 12, 305
 - first director, 351
- Weingarten, Fredrick
 - biography, 368
 - computer networking-related grants, 36
 - Computer Science Education, Research, and Training Section, 18
 - Congressional Office of Technology Assessment, 46
 - cryptography work, 240
 - Special Projects program, 43
 - Theorynet project, 278
- What Can Be Automated?: Computer Science and Engineering Research Study report*, 32
- Wilkinson-Reinsch routines, 229–230
- Wing, Jeanette
 - biography, 369
 - CISE leadership, 148–153, 162, 345
 - computer science curricular projects, 319
- Wolff, Steven
 - background, 65
 - CISE implementation plan, 94–96
 - CISE vision, 110
 - NCRI, 73, 90, 280
 - Networking and Distributed Computing, 72
 - NSFNET benefits, 92
 - NSFNET director, 356
 - NSFNET disappointments, 281
 - NSFNET expansion, 98
 - NSFNET transition to public network, 133, 282–283
- Women in computing
 - BPC program, 316–319
 - education and outreach, 139–140, 345
 - funding, 41
 - NCWIT, 140, 316
 - participation programs, 314–316
 - women’s rights and civil rights movements, 313–317
- Wulf, William A. (Bill)
 - Alphard language, 222
 - biography, 369
 - Boucher Amendment, 101
 - CISE leadership, 97–100
 - collaboratories, 97
 - Human Genome Project, 97
 - supercomputer centers, 307, 341
- XPS (Exploring Parallelism and Scalability) program, 159
- Young, David, numerical analysis by, 232–234
- Young, Paul R.
 - biography, 369
 - CISE leadership, 103–105, 342
 - Committee on CISE Organization, 29
 - CORC, 357
 - NSFNET backbone service, 281
 - supercomputer centers, 307
- Zadeh, Lofti, theoretical computer engineering by, 234–237

Computing and the National Science Foundation, 1950-2016

Building a Foundation for Modern Computing

This organizational history relates the role of the National Science Foundation (NSF) in the development of modern computing. Drawing upon new and existing oral histories, extensive use of NSF documents, and the experience of two of the authors as senior managers, this book describes how NSF's programmatic activities originated and evolved to become the primary source of funding for fundamental research in computing and information technologies.

The book traces how NSF's support has provided facilities and education for computing usage by all scientific disciplines, aided in institution and professional community building, supported fundamental research in computer science and allied disciplines, and led the efforts to broaden participation in computing by all segments of society.

Today, the research and infrastructure facilitated by NSF computing programs are significant economic drivers of American society and industry. For example, NSF supported work that led to the first widely-used web browser, Netscape; sponsored the creation of algorithms at the core of the Google search engine; facilitated the growth of the public Internet; and funded research on the scientific basis for countless other applications and technologies. NSF has advanced the development of human capital and ideas for future advances in computing and its applications.

This account is the first comprehensive coverage of NSF's role in the extraordinary growth and expansion of modern computing and its use. It will appeal to historians of computing, policy makers and leaders in government and academia, and individuals interested in the history and development of computing and the NSF.

ABOUT ACM BOOKS



ACM Books is a series of high-quality books published by ACM for the computer science community. ACM Books publications are widely distributed in print and digital formats by major booksellers and are available to libraries and library consortia. Individual ACM members may access ACM Books publications via separate annual subscription.

