Deconstructing the Computer: Report of a Symposium

Dale W. Jorgenson and Charles W. Wessner, Editors, Committee on Deconstructing the Computer, Committee on Measuring and Sustaining the New Economy, National Research Council

ISBN: 0-309-53320-1, 182 pages, 6x9, (2005)

This free PDF was downloaded from:
http://www.nap.edu/catalog/11457.html

Visit the National Academies Press online, the authoritative source for all books from the National Academy of Sciences, the National Academy of Engineering, the Institute of Medicine, and the National Research Council:

- Download hundreds of free books in PDF
- Read thousands of books online, free
- Sign up to be notified when new books are published
- Purchase printed books
- Purchase PDFs
- Explore with our innovative research tools

Thank you for downloading this free PDF. If you have comments, questions or just want more information about the books published by the National Academies Press, you may contact our customer service department toll-free at 888-624-8373, visit us online, or send an email to comments@nap.edu.

This free book plus thousands more books are available at http://www.nap.edu.

Copyright © National Academy of Sciences. Permission is granted for this material to be shared for noncommercial, educational purposes, provided that this notice appears on the reproduced materials, the Web address of the online, full authoritative version is retained, and copies are not altered. To disseminate otherwise or to republish requires written permission from the National Academies Press.

THE NATIONAL ACADEMIES
Advisers to the Nation on Science, Engineering, and Medicine

MEASURING AND SUSTAINING THE NEW ECONOMY

DECONSTRUCTING THE COMPUTER

Report of a Symposium

Dale W. Jorgenson and Charles W. Wessner, Editors

Committee on Deconstructing the Computer

Committee on Measuring and Sustaining the New Economy

Board on Science, Technology, and Economic Policy

Policy and Global Affairs

NATIONAL RESEARCH COUNCIL
OF THE NATIONAL ACADEMIES

THE NATIONAL ACADEMIES PRESS
Washington, D.C.
www.nap.edu

THE NATIONAL ACADEMIES PRESS 500 Fifth Street, N.W. Washington, DC 20001

NOTICE: The project that is the subject of this report was approved by the Governing Board of the National Research Council, whose members are drawn from the councils of the National Academy of Sciences, the National Academy of Engineering, and the Institute of Medicine. The members of the committee responsible for the report were chosen for their special competences and with regard for appropriate balance.

This study was supported by Contract/Grant No. NASW-99037, Task Order 103, between the National Academy of Sciences and the National Aeronautics and Space Administration; Contract/Grant No. OFED-13416 between the National Academy of Sciences and Sandia National Laboratories; Contract/Grant No. CMRC-50SBNB9C1080 between the National Academy of Sciences and the U.S. Department of Commerce; Contract/Grant No. NSF-EIA-0119063 between the National Academy of Sciences and the National Science Foundation; and Contract/Grant No. DOE-DE-FG02-01ER30315 between the National Academy of Sciences and the U.S. Department of Energy. Additional support was provided by Intel Corporation. Any opinions, findings, conclusions, or recommendations expressed in this publication are those of the author(s) and do not necessarily reflect the views of the organizations or agencies that provided support for the project.

International Standard Book Number 0-309-09254-X

Limited copies are available from the Board on Science, Technology, and Economic Policy, National Research Council, 500 5th Street, NW, W547, Washington, DC 20001; 202-334-2200.

Additional copies of this report are available from the National Academies Press, 500 Fifth Street, N.W., Lockbox 285, Washington, DC 20055; (800) 624-6242 or (202) 334-3313 (in the Washington metropolitan area); Internet, http://www.nap.edu

Printed in the United States of America

THE NATIONAL ACADEMIES

Advisers to the Nation on Science, Engineering, and Medicine

The **National Academy of Sciences** is a private, nonprofit, self-perpetuating society of distinguished scholars engaged in scientific and engineering research, dedicated to the furtherance of science and technology and to their use for the general welfare. Upon the authority of the charter granted to it by the Congress in 1863, the Academy has a mandate that requires it to advise the federal government on scientific and technical matters. Dr. Ralph J. Cicerone is president of the National Academy of Sciences.

The **National Academy of Engineering** was established in 1964, under the charter of the National Academy of Sciences, as a parallel organization of outstanding engineers. It is autonomous in its administration and in the selection of its members, sharing with the National Academy of Sciences the responsibility for advising the federal government. The National Academy of Engineering also sponsors engineering programs aimed at meeting national needs, encourages education and research, and recognizes the superior achievements of engineers. Dr. Wm. A. Wulf is president of the National Academy of Engineering.

The **Institute of Medicine** was established in 1970 by the National Academy of Sciences to secure the services of eminent members of appropriate professions in the examination of policy matters pertaining to the health of the public. The Institute acts under the responsibility given to the National Academy of Sciences by its congressional charter to be an adviser to the federal government and, upon its own initiative, to identify issues of medical care, research, and education. Dr. Harvey V. Fineberg is president of the Institute of Medicine.

The **National Research Council** was organized by the National Academy of Sciences in 1916 to associate the broad community of science and technology with the Academy's purposes of furthering knowledge and advising the federal government. Functioning in accordance with general policies determined by the Academy, the Council has become the principal operating agency of both the National Academy of Sciences and the National Academy of Engineering in providing services to the government, the public, and the scientific and engineering communities. The Council is administered jointly by both Academies and the Institute of Medicine. Dr. Ralph J. Cicerone and Dr. Wm. A. Wulf are chair and vice chair, respectively, of the National Research Council.

www.national-academies.org

Committee on Deconstructing the Computer*

Dale Jorgenson, *Chair*
Samuel W. Morris University Professor
Harvard University

M. Kathy Behrens
Managing Director of Medical Technology
Robertson Stephens Investment Management

Michael Borrus
Managing Director
The Petkevich Group, LLC

Ellen Dulberger
Director, AMS Strategy
IBM Global Services

Kenneth Flamm
Dean Rusk Chair in International Affairs
LBJ School of Public Affairs
University of Texas at Austin

William J. Spencer
Chairman Emeritus, *retired*
International SEMATECH

Bronwyn Hall
Professor of Economics
University of California at Berkeley

Mark B. Myers
Visiting Executive Professor of Management
The Wharton School
University of Pennsylvania

William J. Raduchel

Stephen L. Squires
Vice President and Chief Science Officer
Hewlett-Packard Company

---

*As of February 2003.

Committee on Measuring and Sustaining the New Economy*

Dale Jorgenson, *Chair*
Samuel W. Morris University Professor
Harvard University

M. Kathy Behrens
Managing Director of Medical Technology
Robertson Stephens Investment Management

Kenneth Flamm
Dean Rusk Chair in International Affairs
LBJ School of Public Affairs
University of Texas at Austin

Bronwyn Hall
Professor of Economics
University of California at Berkeley

James Heckman
Henry Schultz Distinguished Service Professor of Economics
University of Chicago

Ralph Landau
Consulting Professor of Economics
Stanford University

Richard Levin
President
Yale University

William J. Spencer, *Vice-Chair*
Chairman Emeritus, *retired*
International SEMATECH

David T. Morgenthaler
Founding Partner
Morgenthaler Ventures

Mark B. Myers
Visiting Executive Professor of Management
The Wharton School
University of Pennsylvania

Roger Noll
Morris M. Doyle Centennial Professor of Economics
Stanford University

Edward E. Penhoet
Chief Program Officer
Science and Higher Education
Gordon and Betty Moore Foundation

William J. Raduchel

Alan Wm. Wolff
Managing Partner
Dewey Ballantine

---

*As of February 2003.

Project Staff*

Charles W. Wessner
Study Director

Sujai J. Shivakumar
Program Officer

Kenneth Jacobson
Consultant

David E. Dierksheide
Program Associate

Christopher S. Hayter
Program Associate

Tabitha M. Benney
Program Associate

McAlister T. Clabaugh
Program Associate

---

*As of February 2003.

For the National Research Council (NRC), this project was overseen by the Board on Science, Technology and Economic Policy (STEP), a standing board of the NRC established by the National Academies of Sciences and Engineering and the Institute of Medicine in 1991. The mandate of the STEP Board is to integrate understanding of scientific, technological, and economic elements in the formulation of national policies to promote the economic well-being of the United States. A distinctive characteristic of STEP's approach is its frequent interactions with public- and private-sector decision makers. STEP bridges the disciplines of business management, engineering, economics, and the social sciences to bring diverse expertise to bear on pressing public policy questions. The members of the STEP Board* and the NRC staff are listed below.

Dale Jorgenson, *Chair*
Samuel W. Morris University Professor
Harvard University

M. Kathy Behrens
Managing Director of Medical Technology
Robertson Stephens Investment Management

Bronwyn Hall
Professor of Economics
University of California at Berkeley

James Heckman
Henry Schultz Distinguished Service Professor of Economics
University of Chicago

Ralph Landau
Consulting Professor of Economics
Stanford University

Richard Levin
President
Yale University

William J. Spencer, *Vice-Chair*
Chairman Emeritus, *retired*
International SEMATECH

David T. Morgenthaler
Founding Partner
Morgenthaler

Mark B. Myers
Visiting Executive Professor of Management
The Wharton School
University of Pennsylvania

Roger Noll
Morris M. Doyle Centennial Professor of Economics
Stanford University

Edward E. Penhoet
Chief Program Officer
Science and Higher Education
Gordon and Betty Moore Foundation

William J. Raduchel

Alan Wm. Wolff
Managing Partner
Dewey Ballantine

---

*As of February 2003.

STEP Staff*

Stephen A. Merrill
Executive Director

Russell Moy
Senior Program Officer

Craig M. Schultz
Research Associate

Christopher S. Hayter
Program Associate

Sujai J. Shivakumar
Program Officer

Charles W. Wessner
Program Director

Tabitha M. Benney
Program Associate

David E. Dierksheide
Program Associate

McAlister T. Clabaugh
Program Associate

---

*As of February 2003.

Contents

II. RESEARCH PAPER

III. APPENDIXES

Preface

This report of a workshop is the second in a series designed to improve our understanding of the technological and economic trends underlying the growth and productivity increases that have created what many refer to as the New Economy. The "New Economy" refers to a fundamental transformation in the U.S. economy as businesses and individuals capitalize on new technologies, new opportunities, and new national investments in computing, information, and communications technologies. Use of this term reflects a growing conviction that substantial change has occurred in the structure of the U.S. economy, and this change is permanent.[1] The goal of this analytical effort, led by the National Research Council's Board on Science, Technology, and Economic Policy (STEP), is to improve our understanding of the sources of gains in growth and productivity and our understanding of the policies required to sustain the benefits of this New Economy for the nation.

Even the casual observer is aware of the ongoing revolution in communications, computing, and information management.[2] In the mid-1990s, this technological revolution contributed to a distinct rise in the long-term growth trajectory

---

[1] In the context of this analysis, the New Economy does not refer to the boom economy of the late 1990s.

[2] This is especially so for the computer hardware sector and perhaps for the Internet as well, although there is insufficient empirical evidence on the degree to which the Internet may be responsible. For a discussion of the impact of the Internet on economic growth, see "A Thinker's Guide," *The Economist*, March 30, 2000. For a broad study of investment in technology-capital and its use in various sectors, see McKinsey Global Institute, *U.S. Productivity Growth 1995–2000: Understanding the Contribution of Information Technology Relative to Other Factors*, Washington, D.C.: McKinsey & Co., October 2001.

of the United States.[3] The term New Economy captures this new reality and has now become widely accepted by leading economists as a long-term productivity shift of major significance.[4] The concept of a New Economy gained currency in the mid-1990s when new data began to reveal an acceleration of growth accompanying a transformation of economic activity.[5] This shift in the rate of growth had coincided with a sudden, substantial, and rapid decline in the prices of computers—from 15 to about 28 percent annually after 1995—hand-in-hand with significant increases in computing power and function.[6] It also coincided with a shift in the rate of decline in price for memory and logic devices—from 40 percent per year until the mid-1990s to about 60 percent thereafter. What is less widely appreciated is that much of this progress is derived from the significant and sustained increases in semiconductor productivity, predicted over 30 years ago by Gordon Moore and known as Moore's Law.

In approaching a phenomenon as complex as the New Economy, it is important to understand—and sort out—diverse elements of technological innovation, structural change, and the impact of public policy as well as issues of measurement.

- Technological innovation—more accurately, the rapid rate of technological innovation in information technology (including semiconductors, computers, software, and telecommunications) and the rapid growth of the Internet—are seen

---

[3] See Dale Jorgenson and Kevin Stiroh, "Raising the Speed Limit: U.S. Economic Growth in the Information Age," in National Research Council, *Measuring and Sustaining the New Economy*, Dale W. Jorgenson and Charles W. Wessner, eds., Washington, D.C.: National Academy Press, 2002.

[4] The introduction of advanced productivity-enhancing technologies obviously does not eliminate the business cycle. See Organisation for Economic Co-operation and Development, *Is There a New Economy? A First Report on the OECD Growth Project*, Paris: Organisation for Economic Co-operation and Development, June 2000, p. 17. For an early discussion, see also M. N. Baily and R. Z. Lawrence, "Do We Have an E-conomy?" NBER Working Paper 8243, April 23, 2001, at <*http://www.nber.org/papers/w8243*>.

[5] "Despite differences in methodology and data sources, a consensus is building that the remarkable behavior of IT prices provides the key to the surge in economic growth." See Dale W. Jorgenson, "Information Technology and the U.S. Economy," Presidential Address to the American Economic Association, New Orleans, LA, January 6, 2001. This consensus continues to expand. See *The Economist*, "The new 'new economy'," September 11, 2003.

[6] Jorgenson and Stiroh describe this acceleration as a "point of inflection," where the price decline abruptly rose from 15 percent annually to 28 percent. In response to this rapid price decline, investment in computer technology exploded, and its contribution to growth rose more than five-fold. In the latter half of the 1990s, Jorgenson and Stiroh find that computers contributed 0.46 percentage points per year to economic growth. Software and communications equipment contributed an additional 0.30 percentage points per year from 1995 to 1998. Preliminary estimates through 1999 revealed further increases for all three categories (computers, communications equipment, and software). Jorgenson's and Stiroh's analysis builds the case for "raising the speed limit," that is, for revising upward the intermediate-term projections of growth for the U.S. economy and supports the notion that the economy is on a higher productivity path, similar to that experienced from the early 1950s through the early 1970s. See Dale W. Jorgenson and Kevin J. Stiroh, "Raising the Speed Limit: U.S. Economic Growth in the Information Age," *op. cit.*

as the sources of the productivity gains that characterize the New Economy. These productivity gains derive first from the exponential growth in semiconductor performance at ever lower cost.[7] In addition, the use of information technologies in the production of computers has greatly increased the productivity of this industry while having substantial positive effects (albeit with a lag) on the productivity of other important sectors of the economy such as banking, retail, and transportation.[8] Many therefore believe that the productivity gains of the New Economy are closely linked to this unprecedented rate of technological innovation.[9]

- Structural changes arise from a reconfiguration of knowledge networks and business patterns made possible by innovations in information technology. Phenomena, such as business-to-business e-commerce and Internet retailing, are altering how firms and individuals interact, enabling greater efficiency in purchases, production processes, and inventory management.[10] These structural changes are still emerging as the use and applications of the Internet continue to evolve.
- Public policy plays a major role at several levels. This includes the government's role in fostering rules of interaction within the Internet[11] and its discretion in setting and enforcing the rules by which technology firms, among others, compete.[12] More familiarly, public policy concerns particular fiscal and regulatory choices that can affect the rate and focus of investments in sectors such as telecommunications. The government also plays a critical role within the inno-

---

[7] Price declines, for higher performance, have remained on the order of 17–20 percent per annum. See the presentation by Kenneth Flamm in this volume.

[8] See, for example, Stephen Oliner and Daniel Sichel, "The Resurgence of Growth in the late 1990s: Is Information Technology the Story?" *Journal of Economic Perspectives* 14(4), Fall 2000. Oliner and Sichel estimate that improvements in the computer industry's own productive processes account for about a quarter of the overall productivity increase. They also note that the use of information technology by all sorts of companies accounts for nearly half the rise in productivity.

[9] See Alan Greenspan's remarks before the White House Conference on the New Economy, Washington, D.C., April 5, 2000. <*www.federalreserve.gov/BOARDDOCS/SPEECHES/2000/20000405.HTM*>. For a historical perspective, see the Proceedings section of this volume. Ken Flamm compares favorably the economic impact of semiconductors today with the impact of railroads in the nineteenth century.

[10] See, for example, Brookes Martin and Zaki Wahhaj, "The Shocking Economic Impact of B2B," *Global Economic Paper*, 37, Goldman Sachs, February 3, 2000.

[11] Dr. Vint Cerf notes that the ability of individuals to interact in potentially useful ways within the infrastructure of the still-expanding Internet rests on its basic rule architecture: "The reason it can function is that all the networks use the same set of protocols. An important point is these networks are run by different administrations, which must collaborate both technically and economically on a global scale." See comments by Dr. Cerf in National Research Council, *Measuring and Sustaining the New Economy*, Dale W. Jorgenson and Charles W. Wessner, eds., Washington, D.C.: National Academy Press, 2002. Also in the same volume, see the presentation by Dr. Shane Greenstein on the evolution of the Internet from academic and government-related applications to the commercial world.

[12] The relevance of competition policy to the New Economy is manifested by the intensity of interest in the antitrust case, *United States versus Microsoft*, and associated policy issues.

vation system.[13] It supports national research capacities, providing incentives (or disincentives) to promote education and training in key disciplines, and funds most of the nation's basic research.[14] The government also plays a major role in stimulating innovation. It does this most broadly through the patent system.[15] In addition, government procurement and innovation awards have played key roles in the development of new technologies to fulfill national missions in defense, agriculture, health, and the environment.[16]

Collectively, these public policies play a central role in the continued development of the New Economy. Sustaining the contributions of this New Economy to growth, productivity, and employment will require public policy to remain relevant to the rapid technological and structural changes that characterize it and responsive to new trends and policy challenges in the global marketplace.

THE CONTEXT OF THIS REPORT

STEP's earlier analysis, *U.S. Industry in 2000*, assessed the determinants of competitive performance in a wide range of manufacturing and service industries, including those relating to information technology.[17] The STEP Board also undertook a major study, chaired by Gordon Moore of Intel, on how government-industry partnerships can support growth-enhancing technologies.[18] Reflecting a growing recognition of the importance of the surge in productivity since 1995, the Board launched a multifaceted assessment, exploring the sources of growth, measurement challenges, and the policy framework required to sustain the New Economy. The first exploratory volume was published in 2002.[19] Subsequent workshops and ensuing reports in this series include *Productivity and Cyclicality in the Semiconductor Industry* and *Deconstructing the Computer*—the present

---

[13] See Richard Nelson, ed., *National Innovation Systems*, New York: Oxford University Press, 1993.

[14] National Research Council, *Trends in Federal Support of Research in Graduate Education*, Washington, D.C.: National Academy Press, 2001.

[15] In addition to government-funded research, intellectual property protection plays an essential role in the continued development of the biotechnology industry. See Wesley M. Cohen and John Walsh, "Public Research, Patents and Implications for Industrial R&D in the Drug, Biotechnology, Semiconductor and Computer Industries," in National Research Council, *Capitalizing on New Needs and New Opportunities: Government-Industry Partnerships in Biotechnology and Information Technologies*, Washington, D.C.: National Academy Press, 2002.

[16] For example, government support played a critical role in the early development of computers. See Kenneth Flamm, *Creating the Computer*, Washington, D.C.: The Brookings Institution, 1988.

[17] National Research Council, *U.S. Industry in 2000: Studies in Competitive Performance*, David C. Mowery, ed., Washington, D.C.: National Academy Press, 1999.

[18] For a summary of this multivolume study, see National Research Council, *Government-Industry Partnerships for the Development of New Technologies: Summary Report*, Charles W. Wessner, ed., Washington, D.C.: The National Academies Press, 2003.

[19] National Research Council, *Measuring and Sustaining the New Economy, Report of a Workshop, op. cit.*

report, which looks at the component industries of the computer. Future reports in the series will address the software sector, as well as the policies required to sustain the New Economy.

SYMPOSIUM

Believing that increased productivity in the semiconductor and computer component industries plays a key role in sustaining the New Economy, the Committee on Measuring and Sustaining the New Economy, under the auspices of the STEP Board, convened a symposium on *Deconstructing the Computer* on February 28, 2003, at the National Academies in Washington, D.C. The symposium focused on metrics currently used in measuring computer performance and the sources of productivity growth in computers, examining current trends in hardware, components, and peripherals. The symposium included presentations and remarks from leading academics and innovators in the information technology sector (Appendix B lists these individuals). These presentations highlighted the need to develop more appropriate quality-adjusted price measures and better data on the performance of computers and computer components.

The "Proceedings" chapter of this volume contains summaries of their workshop presentations and discussions. The speakers have each approved the summary of their statements. Also included in this volume is a paper by Jack Triplett, "Performance Measures for Computers," which was presented at the symposium. Given the quality and the number of presentations, summarizing the workshop proceedings has been a challenge. We have made every effort to capture the main points made during the presentations and the ensuing discussions. We apologize for any inadvertent errors or omissions in our summary of the proceedings. The lessons from this symposium and others in this series will contribute to the Committee's final consensus report on *Measuring and Sustaining the New Economy*.

ACKNOWLEDGMENTS

In order to better understand the technological drivers and appropriate regulatory framework for the New Economy, as well as obtain a better grasp of its operation, a number of agencies have played a role in the creation and development of the New Economy. We are grateful for the participation and the contributions of the National Aeronautics and Space Administration, the Department of Energy, the National Institute of Standards and Technology, the National Science Foundation, and Sandia National Laboratories.

We are indebted to Ken Jacobson for his preparation of the meeting summary. Several members of the STEP staff, including Sujai Shivakumar, also deserve recognition for their contributions to the preparation of this report. We are also indebted to Tabitha Benney and Christopher Hayter, who prepared the con-

ference and, with David Dierksheide and McAlister Clabaugh, helped prepare this report for publication.

NATIONAL RESEARCH COUNCIL REVIEW

This report has been reviewed in draft form by individuals chosen for their diverse perspectives and technical expertise, in accordance with procedures approved by the NRC's Report Review Committee. The purpose of this independent review is to provide candid and critical comments that will assist the institution in making its published report as sound as possible and to ensure that the report meets institutional standards for quality and objectivity. The review comments and draft manuscript remain confidential to protect the integrity of the process.

We wish to thank the following individuals for their review of this report: Ana Aizcorbe, Bureau of Economic Analysis; Robert Keyes, IBM Thomas Watson Research Center; William Knaus, University of Virginia; and Juri Matisoo, Semiconductor Industry Association.

Although the reviewers listed above have provided many constructive comments and suggestions, they were not asked to endorse the content of the report, nor did they see the final draft before its release. Responsibility for the final content of this report rests entirely with the authoring committee and the institution.

STRUCTURE

This report has three parts: a summary of the proceedings of the February 28, 2003, symposium; a research paper by Dr. Jack Triplett; and finally, a bibliography that provides additional references.

This report represents an important step in a major research effort by the Board on Science, Technology, and Economic Policy to advance our understanding of the factors shaping the New Economy, the metrics necessary to understand it better, and the policies best suited to sustaining the greater productivity and prosperity that it promises.

Dale W. Jorgenson Charles W. Wessner

I

PROCEEDINGS

Introductory Remarks

Dale W. Jorgenson
Harvard University

Dr. Jorgenson, Chairman of the Board on Science, Technology, and Economic Policy (STEP) of the National Research Council, welcomed those assembled for the day's symposium, "Deconstructing the Computer," the second in a series of STEP Board-sponsored meetings devoted to "Measuring and Sustaining the New Economy." This project focuses on the role of information technology (IT) in the U.S. economy.

Posing the question "Why information technology?," Dr. Jorgenson offered three answers:

- Developments in IT have been the main source of productivity growth in the U.S. economy for three decades.
- This relatively small industrial segment, about 5 percent of the economy in terms of value added, by itself produced the great resurgence of the latter half of the 1990s that put the U.S. economy back on something approximating its historic track of growth.
- IT had been, in addition, the major force behind a remarkable development of productivity in the years 2001–2002, during which the economy was in a recession.

For these reasons, Dr. Jorgenson explained, the STEP Board had chosen information technology as a launching pad for examining the relationships among science, technology, and economic policy. Semiconductors, the bedrock of IT's development, had served as the subject of the series' first symposium, "Productivity and Cyclicality in Semiconductors: Trends, Implications, and Questions."

Dr. Jorgenson proposed summarizing the major findings of that initial conference as a way of providing a context for the day's meeting. To highlight the speed of technological change prevailing in the semiconductor industry, he cited an analogy formulated in 1998 by Gordon Moore, Chairman Emeritus of Intel Corporation: "If the automobile industry advanced as rapidly as the semiconductor industry, a Rolls-Royce would get half a million miles per gallon, and it would be cheaper to throw it away than to park it." Dr. Jorgenson specified that Moore's intent was not to disparage the automobile industry but rather to illustrate how much faster the semiconductor industry had advanced than the rest of the economy as it progressed from the invention of the transistor through the development of the integrated circuit and on to the creation of memory and logic chips. As a second expression of the speed of development of semiconductor technology, he invoked another, more familiar, statement by Dr. Moore commonly referred to as Moore's Law: "The number of transistors on a chip doubles every 18–24 months." This he illustrated with the aid of a graph tracing the growth that took place between 1970 and 2000 in the number of transistors per chip for both processing and storage devices (see Figure 1). "What's remarkable about Moore's Law," he declared, "is not just the rapid development that Moore predicted in 1965, but the fact that this has tracked the history of the industry for the subsequent three decades."

Turning to the impact of the semiconductor industry on the economy, Dr. Jorgenson explained that the economist's basic task in this case—a very difficult one, he admitted, owing to the rate at which semiconductor technology has developed—is to hold the performance of the device constant while tracking its price. He displayed a chart that plots in logarithmic scale indexes for the prices of computers, memory chips, and logic chips, all of them assigned a value of 1 for 1996 (see Figure 2).

The computer price shows a rate of decline far exceeding the growth of productivity in the U.S. economy: This decline held steady at about 15 percent per year until 1995, at which point it doubled. The rate of decline began accelerating in 1995 for memory and logic devices as well, with the latter showing the most pronounced drop-off of all: At about 40 percent per year until the mid-1990s, it hit about 60 percent per year thereafter. Pointing out that when economists look at price indexes, they usually expect to see a trend going in the opposite direction, Dr. Jorgenson called the price decline of computer prices "a phenomenon" and that of microprocessor prices "stupendous"; the latter, he noted, indicates "very, very rapid" technological change.

Connecting developments in technology to developments in prices is of interest to economists, he said, because these developments are at the heart of the revival of U.S. economic growth and of the trend in productivity that has pre-

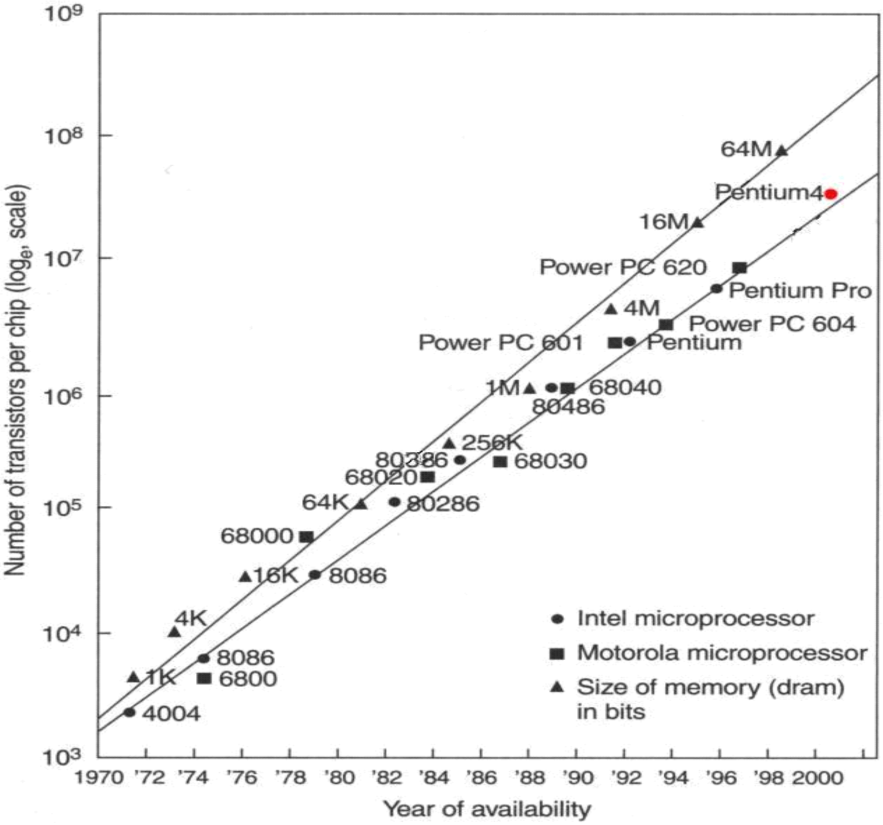

FIGURE 1 Transistor density on microprocessors and memory chips.

vailed for the last three decades. Still, he preferred to focus on the period beginning in the mid-1990s, where, he asserted, “the lodestone” is to be found. “What,” he asked, “happened to generate this acceleration in the rate of [price] decline that showed up in computers as well?” To help in examining this question, Dr. Jorgenson displayed a chart that contrasts industry predictions for the increase in miniaturization of semiconductor feature size with the advances actually realized (see Figure 3).

Even though the feature size that in 1995 had been projected for 1998 was reached early, in 1997, the industry remained cautious: Its consensus in 1997 was that the accelerated rate of feature-size decline would continue through 1999 at the latest. In fact, that rate still holds true, Dr. Jorgenson said, and the current

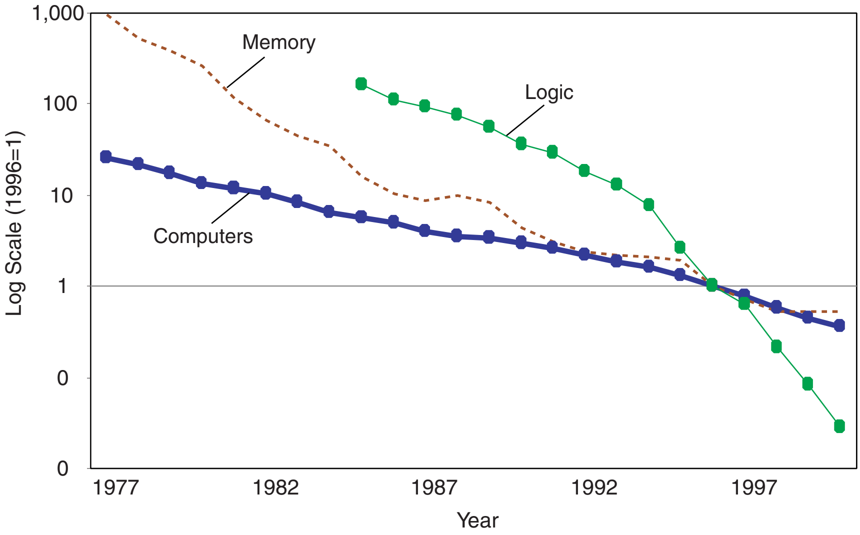

FIGURE 2 Relative prices of computers and semiconductors, 1977–2000.

International SEMATECH roadmap posits that it will continue to hold at least through 2007.[1] "Moore's Law accelerated," he observed, "and that is precisely what produced the sharp increase in the rate of decline of prices."

Stressing the significance of this link, Dr. Jorgenson proposed as a goal for the symposium the generation of a series of similar road maps, each covering a different component of computer prices. This would amount to deconstructing the curve representing the evolution of the computer price into curves capturing the changes in technology underlying it. Also on the agenda, he said, is projecting those changes out for a decade or so in hopes of understanding how much of the acceleration in the rate of progress is permanent and how much has been temporary and can be expected to abate, perhaps reverting to its previous pace. Beyond that, telecom equipment, information about whose price trends is now accumulating, and "the very difficult problem of software" were put forward as subjects of similar study.

---

[1] A Technology Roadmap can help identify, evaluate, and select alternative technology paths to realize the product needs of a company or an organization. It identifies the system requirements, the product and process performance targets, and the technology alternatives and milestones for meeting those targets. Under different circumstances with uncertainty or risk, one or multiple paths can be selected and pursued for achieving those objectives. A benefit of technology roadmapping is that it provides information to make better decisions for technology research investment. See Ronald N. Kostoff and Robert R. Schaller, "Science and Technology Roadmaps," *IEEE Transactions on Engineering Management*, May 2001, 48(2):132–143.

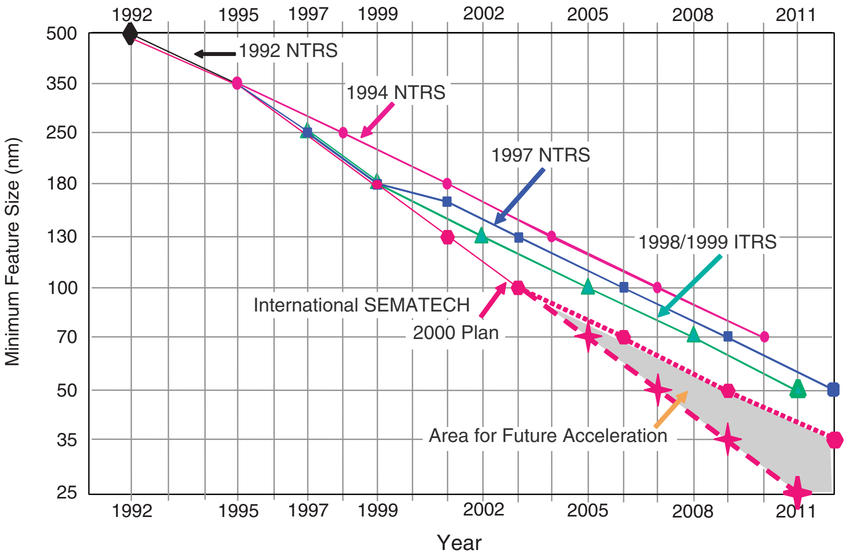

FIGURE 3 Semiconductor roadmap acceleration.

Dr. Jorgenson then introduced the moderator of the day's first panel, Steve Landefeld, the Director of the Bureau of Economic Analysis (BEA) of the U.S. Department of Commerce. Identifying BEA as the source of information on IT trends that has been "the key" to understanding the connection between technology and the U.S. economy, Dr. Jorgenson praised the agency for its pioneering use in 1985 of a price index for computers. This index, he noted, would be discussed in detail by the panel's first speaker, Jack Triplett of the Brookings Institution. He then turned the floor over to Dr. Landefeld.

Panel I

Performance Measurement and Current Trends

MEASURES OF PERFORMANCE USED IN MEASURING COMPUTERS

Jack E. Triplett
The Brookings Institution

The aim of this initial presentation, Dr. Triplett said, was to build a bridge between technologists and economists—professions that "don't talk very much to each other" but had been brought together for the symposium because both are interested in performance measures for computers and their components. The opening of his presentation, directed in particular to technologists, would describe in more detail than had Dr. Jorgenson what economists do with computer performance measures and what kinds of computer performance measures they really want.

Starting with what he called a "really simple, stylized example," Dr. Triplett posited the existence in the economy of a single computer, UNIVAC. Referring to figures obtained from a book by fellow speaker Kenneth Flamm, he cited UNIVAC's price at $400,000 in 1952 and said that three units were made that year, so that "current-price output" for computers in 1952 was $1.2 million. For the sake of example he put forward the assumption that in 1955 a fictitious "New Computer" came along and supplanted UNIVAC as the only computer in the

TABLE 1 Example 1: UNIVAC and the "New Computer"

	UNIVAC (1953)	"New Computer" (1955)
Number produced	3	10
Price, each	\$400K	\$600K
Current price output	\$1.2 million	\$6.0 million
Performance index	1.0	1.5
Computer inflation (with estimated performance premium = 1 + 0.7(0.5))	1.00	1.11
Computer "constant price" output index	1.00	4.60

economy, that 10 were produced, and that they sold for \$600,000 apiece. In the 1955 economy, therefore, the revenue to the producers of computers—or the total amount of computer output, or the total amount of investment in computers, all of these being the same thing—would have been \$6.0 million (see Table 1). Having observed such an increase in unit price and growth in output, he said, "The first thing economists will ask is: 'How much inflation is there in the computer market?'"

Inflation cannot be determined from the above data alone, however: A performance index is necessary in addition so that increase in price can be adjusted for change in performance. Dr. Triplett proposed setting the performance index at 1.0 for UNIVAC and assigning New Computer 1.5 performance units; had the two computers been produced at the same time, he explained, it would have been clear that part of the difference in price between them would have represented a performance premium. "Similarly," he said, "we don't want to show '\$1.2 million to \$6 million' as the change in output if, in fact, the computer in the second period is in some sense more computer than the one in the first period." To make these adjustments, it is necessary to put a value on the increase in performance from UNIVAC to New Computer.

Better for this purpose than having just UNIVAC and New Computer in production at the same time would be having a number of computers available simultaneously. In that case, Dr. Triplett said, a value could be mathematically ascertained and used to adjust not only the price of New Computer compared to UNIVAC in order to get a measure of inflation, but also the change in the output measure in order to get the change in real output. This could be done with a Hedonic Function:[2]

$$P_i = a_0 + a_1 (M)_i + e_i \qquad (1)$$

---

[2] Hedonic price indices are a way to correct deflators for quality changes in the goods they are measured for, using product characteristics such as memory capacity or processor speed (in the case

TABLE 2 Example 2: UNIVAC and the "New Computer"

	UNIVAC (1953)	"New Computer" (1955)
Number Produced	3	10
Price, Each	\$400K	\$600K
Current Price Output	\$1.2 million	\$6.0 million
Performance Index	1.0	1.8
Computer Inflation (with Estimated Performance Premium = 1+0.7(0.5))	1.00	0.96
Computer "Constant Price" Output Index	1.00	5.17

a regression in which P is a vector of prices of computers, models are indexed by the letter i, M is the associated performance measure for each computer, and e_i is the regression error term. The value of the coefficient a_1 derived from the regression would be used to make the adjustments.

In the example of UNIVAC and New Computer, if a_1 is 0.7 and the increase in performance is 0.5, computer-price inflation comes out at 11 percent—not the 50 percent that would be obtained by simply comparing the machines' prices. Thus, taking out of the inflation measure a performance premium estimated from a cross-section of computers leaves the true computer-price change. Similarly, the standard way that the Bureau of Economic Analysis (BEA) would calculate computer output would be to deflate the change of \$4.8 million, from \$1.2 million to \$6 million, by the price index (that is: (\$6.0 / \$1.2) / 1.11 = 4.60, which has the interpretation that output increased 4.6 times). In the example, the increase in computer output is a little smaller than the ratio of \$1.2 million and \$6 million because computer inflation is taking place that reduces the current-price output measure to below the constant-price output measure.

Dr. Triplett then demonstrated what happens when the example is altered by raising New Computer's performance units to 1.8 from 1.5: The same regression yields a value of less than 1.0, indicating negative inflation in the computer price (see Table 2). In "this more realistic case," he said, the computer price falls "be-

---

of computers), or speed and space (in the case of automobiles). See Zvi Griliches, (ed.), *Price Indexes and Quality Change*, Harvard University Press, Cambridge, MA, 1971; Zvi Griliches and M. Ohta, "Automobile Prices and Quality: Did the Gasoline Price Increase Change Consumer Tastes in the U.S.?" *Journal of Business and Economic Statistics*, April 1986, 4(2):l87–198. Price changes are "corrected" for changes in product characteristics using a so-called hedonic function that is estimated econometrically. Hedonic price indices are now available for the U.S. computer industry (which is the leading ICT producer in the world), as well as for some other countries. See A. W. Wyckoff, "The Impact of Computer Prices on International Comparisons of Labor Productivity," *Economics of Innovation and New Technology,* 3(3-4):277–294, 1995.

TABLE 3 Private Fixed Investment in Computers and Peripheral Equipment

	1995	1996	1997	1998	1999	2000	2001	2002
Billions of current dollars	64.6	70.9	79.6	84.2	90.4	93.3	74.2	74.4
Billions of chained 1996 dollars	49.2	70.9	102.9	147.7	207.4	246.4	239.9	284.1
Chained price index	131.29	100	77.38	56.99	43.6	37.87	30.91	26.27
Quantity index	69.4	100	145.22	208.39	292.64	347.77	338.61	400.92

SOURCE: Bureau of Economic Analysis, NIPA Tables 5.4, 5.5, and 7.6.

cause, with a bigger increment of performance, I take out a larger amount when I take out the performance premium." Again, using the same index to correct the output measures yields a change in constant-price output—which is an indicator of the rate of growth of output—that is larger than the change in current-price output. (In national accounts, he noted, constant-price output is called "real output" by economists. BEA publishes it in tables under the title "Chained-Type Quantity Index.") Dr. Triplett stressed that the performance measure is of "extraordinary importance," saying the example illustrates that the determination of two critical factors—the level of price inflation and whether real investment growth is outstripping nominal investment growth—depends on the measure of computer performance.

Displaying BEA figures for private fixed investment in computers and peripheral equipment for 1995 through 2002 (see Table 3), Dr. Triplett pointed out that actual investment went up until 2000 but had come down since then as a result of a slump in purchases of high-tech equipment. BEA's "chained price index" shows, however, that computer prices as a measure of the national account fell consistently, and quite rapidly, over the entire period. The change in actual investment, or current-price shipments, divided by the change in the price index yields the deflated output measure for computers—"billions of chained 1996 dollars," in BEA parlance—which he described as growing "much, much, much, much faster" than expenditures on computers.

Thus, although the growth rate for current expenditures on computer equipment for the period 1995–2000 was 7.6 percent, the price index dropped 22 percent over that interval, raising the deflated output measure "that's telling you what's going on in terms of real performance" by 38 percent, he said (see Table 4). Even during the post-2000 slump, as the rate of growth of actual spending dropped 10 percent per year, the price index kept moving downward—with the result that the deflated investment in computers increased. To illustrate just how long this pattern of decrease in the price index has been under way, Dr. Triplett showed a graph of the computer-price index going back to 1953, the date of the

TABLE 4 Fixed Private Investment in Computer and Peripheral Equipment

	Average Annual Growth Rates	
	1995–2000	2000–2002
Billions of current dollars	7.6	–10.7
Chained price index	–22.0	–16.7
Quality index	38.0	7.4

commercial computer's introduction in the United States (see Figure 4). The price index, because it is adjusted for performance, represents the price of computer power, he said, and so what the chart demonstrates is that the price of computer power today is 0.001 of 1 percent of what it was when the computer was introduced. Noting that a dotted line representing the price index for PCs has been falling more rapidly in recent years than the index for mainframes, he stated: "That's why we've got so many PCs." He then reminded the audience that all the foregoing numbers "depend crucially on having a measure of computer performance."

Next, Dr. Triplett took up the question of which computer performance measures economists have actually used. Early research, focusing on the mainframe,

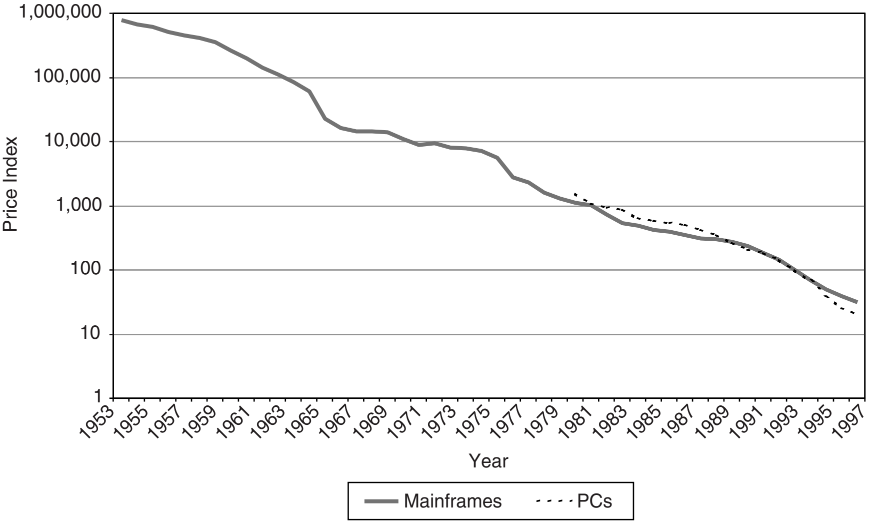

FIGURE 4 Price indexes for mainframes and PCs.

borrowed measures used by technologists studying the rate of technical progress in computers. Initially, measures of computer performance tended to be simple, "some sort of single instruction speed," a popular example of which was multiplication time. When that began to seem too elementary, what is now called "clock speed," came into vogue among technologists. This measure was picked up, in turn, by economists. During a third stage attention turned to the various tasks a computer can perform, and researchers asked themselves: "Why don't we go out and actually take a sample of jobs, try to measure the speed with which jobs are done, and use that as a performance measure?" That course turned out to be far more complicated than it might seem at first blush because not only do computers do many different things, they have many different users. While opining that "trying to measure the performance of a computer by looking at jobs" is better than looking solely at clock speed, Dr. Triplett cautioned that it "inevitably involves an aggregation over the instructions in the jobs, because computers have different kinds of jobs, and then an aggregation over the users of the computers, because users have different mixes of instructions." The early history of research into computer performance culminated with an innovative project carried out in the 1980s by IBM and BEA; that study's measure was MIPS (millions of instructions per second), one of a number of weighted-instruction mixes current at the time. Bidding farewell to the mainframe period, Dr. Triplett referred to the audience to a table summarizing the variables then used in hedonic functions for computers.[3]

He returned to the regression he had previously discussed along with a second, far more complex hedonic function:

$$P_i = a_0 + a_1 (M)_i + e_i \tag{1}$$

$$\ln P_i = a_0 + a_{\mathrm{i}} \ln (M_1)_i + a_2 \ln (M_2) + \ldots a_k \ln (M_k) + e_i \tag{2}$$

He explained that while the simpler regression, having a single variable, might apply where there is a single characteristic of performance, computer performance is multidimensional and therefore demands a hedonic function with more variables. Economists who study computers, wanting to take many different measures of performance into account, would therefore estimate an equation that looks more like the latter. In it, each characteristic of computer performance is represented by M_1 through M_k, giving coefficients a_1 through a_k covering the set of computer performance characteristics. The coefficients in this second, "more realistic" hedonic function are those employed by the Bureau of Labor Statistics (BLS), which now produces the price indexes that go into the BEA national accounts, to adjust for changes in computer performance.

---

[3] See Annex A, "Comparison of Variables in Hedonic Functions for Computers," in Dr. Triplett's paper, "Performance Measures for Computers," in this volume.

Dr. Triplett then referred the audience to a table summarizing the variables used in computer hedonic functions found mainly in economists' recent studies of personal computers.[4] The data in the extreme left-hand column of the table come, however, not from a study of PCs but rather from the IBM study of mainframes used by BEA in 1985 to derive the first computer-price indexes for its national accounts. Of the five products covered in that study—the mainframe computer, disk drive, tape drive, printer, and display—the three that are relevant to PCs are found in the table: the processor, hard-disk drive, and monitor. Thus, he observed, "Economists who have done work on PCs have consciously or unconsciously taken over the variables used in the initial IBM study and applied them to PCs—which makes some sense." He said the list of variables from the recent studies, among which he singled out Chwelos (2002) and the BLS study for praise, comprises "the most exhaustive specifications that have been used so far" in modeling computer performance.[5] Nonetheless, there are "problems" with the variables, he cautioned, remarking that "nobody seems to pay much attention to the speed of the hard disk, even though that was an important performance variable in the IBM study." He then referred to a second table enumerating not only additional hardware features but also dummy variables.[6] Because few PCs today lack a sound card, for instance, a sound-card dummy might be entered into the regression as evidence that the researcher had asked whether it was present or not. "The best economists have done so far with many of the features that differentiate a modern PC from the assemblage of equipment that was in the computer center in 1985 is just to ask 'Was it there?'" he said, "not, 'What is its performance?'"

In conclusion, Dr. Triplett conceded one would be justified in observing that this aggregate of performance measures "doesn't seem real exciting as a model of what a computer does" and in asking: "Why haven't we gone further?" The answer, he said, is "data. Where do we get the data to improve the model of the computer?" And he posed a second question that is somewhat subsidiary to the question of data: "How do we figure out what performance measure we're really looking for today?" One of the things economists would like to learn from technologists so that they can increase the sophistication of their model of computer operation—which is, in essence, what a hedonic function represents—is how to look at the performance of some components of modern computers. Once they understand better what they should be looking for, the second thing economists hope to get from technologists is the knowledge of where to find the data to

---

[4] See Annex B, "Variables in Computer Hedonic Functions, Selected Studies [Hardware Components Only], in Dr. Triplett's paper, "Performance Measures for Computers," in this volume.

[5] See Michael Holdway, 2001, "Quality-Adjusting Computer Prices in the Producer Price Index: An Overview," Bureau of Labor Statistics, October 16.

[6] See Annex B, "Computer Hedonic Functions, Other Hardware Features," in Dr. Triplett's paper, "Performance Measures for Computers," in this volume.

improve their measurements. Looking forward to the rest of the symposium, Dr. Triplett stated: "What I'm hoping to hear is: What are the measures of performance in these components? What are the measures of performance used in the industry? And where can data be obtained to put some of these performance measures into better models of the computer by economists?"

OVERVIEW OF THE IBM GLOBAL PRODUCT PLAN

David F. McQueeney
International Business Machines

Dr. McQueeney defined the challenge before the symposium as figuring out how to measure value all the way across information technology's "quite complex food chain of value-added." Information technology insiders and economists who study the industry have both focused on the transistor, which he put near the bottom of that food chain, for a variety of reasons: It was an obvious place to start; it was easy to deal with; and it demonstrated constant, rapid progress that was clearly beneficial. He warned, however, that if greater attention is not paid to the rest of the food chain from now on, there will be repercussions for the industry's production of business value: "We have to be very careful about how we think about the different levels of this food chain, which goes all the way from quantum mechanics and material science at the very bottom to, at the top, a business process that gives you—if you are an IT company—a net economic value-add for your customer's business."[7]

Laying out the path for his talk, Dr. McQueeney signaled that he would react with a "yes, but" to the trend toward "faster, cheaper, better" in technological progress—an exploration of which, he commented, must lead off "any good technology outlook." He would next raise the question of what happens if the IT industry gets "too carried away with deploying technology for technology's sake" and suggest that, by failing to exercise caution, it has in some cases created problems in deploying and provisioning technology, as well as in providing IT services to its customers. Finally, he would address the creation of business value itself by hardware, software, and applications. While acknowledging this as the perennial "Holy Grail for customers," he noted that customers increasingly want to push the responsibility for creation of business value onto their suppliers rather than remaining content to take that responsibility upon themselves.

To illustrate how "faster, cheaper, better" can collide with what he called the "good-enough phenomenon," Dr. McQueeney offered several "thumbnail sketches," the first concerning optical-data transmission bandwidth. Several hun-

---

[7] Of course, there are off-market applications other than those related to business that have value, e.g., military, medical, and scientific activities.

dred colors of light can now be sent over optical fiber, with the result that there are "individual fibers with enough bandwidth to connect every person in North America to every person in Eastern and Western Europe and to allow all to have a phone conversation at the same time." Moreover, during the Internet boom—when such bandwidth was expected to go into commercial use more quickly than it has done in reality—tremendous capacity in optical fiber was installed between various cities and within metropolitan areas of the United States. Dr. McQueeney said he therefore accepts as reasonable "at the raw fiber level" predictions about bandwidths becoming free such as those made by the technology pundit George Gilder. The problem, however, is that "the intelligence needed to light up those fiber-optic networks and make them actually do something useful—the servers, the routers, the switches—is in fact quite expensive, and we're still struggling with a good investment model that will let us build out that control infrastructure to use the fiber capacity that we have." That the quantity of raw glass deployed outstrips the industry's understanding of how to use it in a responsible business fashion constitutes a "good-enough phenomenon" with respect to installed fiber-optic capacity.

Displays and home PCs have also crossed easily understood "good-enough" thresholds. Monitors built within the year previous to the symposium and on sale at the time, he said, feature a resolution in pixels per inch about at the limit of what the human eye can detect in a windowing environment like that of a PC at an 18-inch viewing screen. While there is still progress to be made in such things as medical imaging and preserving art treasures, which call for the world's best digital imagers and displays, displays used for everyday desktop-PC applications have become so good that "further technological improvements aimed at more pixels per inch would not be detectable to the end user." As for overall home-PC performance, under the constraints of current architecture turning up the clock speed does not have a big impact on throughput because the rest of the system is not scaling with it. Similarly, disk capacity has become so large that many casual users never fill the hard drive in the two or three years they keep a computer. "Guys in our research lab can fill up a disk that comes with a standard PC in a busy afternoon," Dr. McQueeney remarked, "and I'm sure that those of you in the room that do economic analysis could do the same thing." For many home users, however, current PC performance is good enough.

So, although research and development do not stop, at a certain point their benefits begin to show up in price reduction and cost-performance reduction rather than in performance measures. "The raw capabilities of the technology have in some cases gotten to a point where either the economics of how you sell them and how you ascribe value to them is changing," he explained, "or you are forced to look elsewhere in the system performance stack to get real improvements."

Dr. McQueeney then displayed a chart comparing the growth in transistor switching speed projected in 1995 with growth achieved (see Figure 5). This chart, he noted, parallels one shown earlier by Dr. Jorgenson (Semiconductor

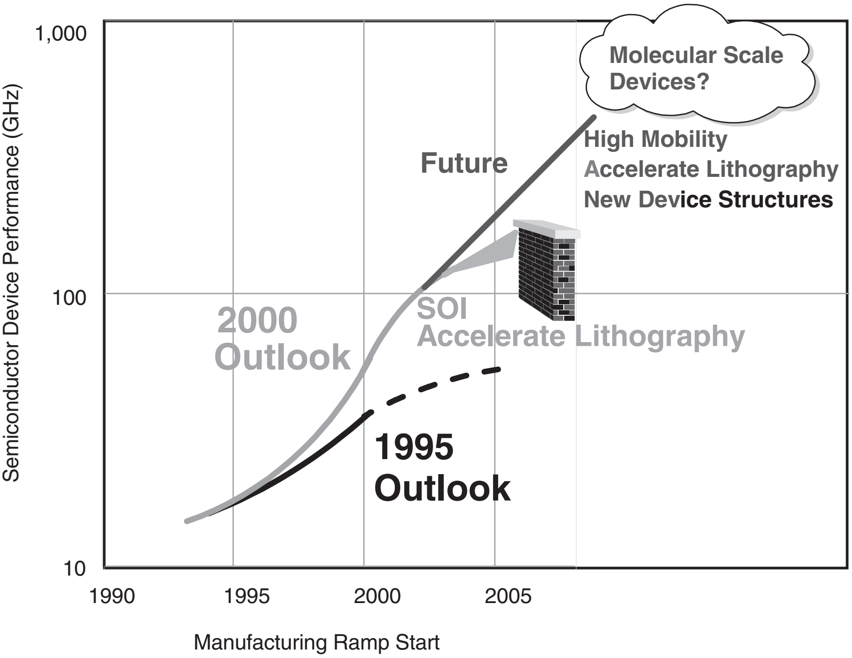

FIGURE 5 Growth in transistor switching speed.
SOURCE: ASCI Roadmap *www.llnl.gov/asci*, IBM. Brain ops/sec: Kurzweil 1999, *The Age of Spiritual Machines,* Moravec 1998, *www.transhumanist.com/volume1/moravec.htm.*

Roadmap Acceleration) that was based on feature size rather than device performance but also clearly showed actual growth far exceeding industry expectations. Observing that researchers have consistently accomplished more than they themselves had predicted, Dr. McQueeney said there is debate inside IBM over whether this discrepancy reflects "sandbagging" by technologists who may fear being fired for missing their own projections. But he also acknowledged the intrusion of a factor he called "technical, or psychological, or both": Anticipating a future encounter with an engineering problem for which no solution exists can incline technologists to look ahead warily. "'I'm not sure that I won't know how to solve it, but I don't know the solution today, so we're going to back off a little bit and say we're not sure if we can keep [advancing] on a straight line,'" is how he characterized their thinking.

He then talked about what is known in the industry as the "brick wall": the point at which semiconductor device developers will "run out of the ability to do clever engineering and run into physics problems." Although the expectation in 2000, as reflected on the chart, was that a brick wall would obstruct progress as of

2003, Dr. McQueeney said that improvement of the semiconductor products available commercially could be anticipated to continue at a rate conforming to Moore's Law for at least 10 more years. But if no brick wall would be standing in the way of shipments for another decade, the barrier it represents was already a real concern in the laboratory, for semiconductor research had reached the point at which the building block for the transistor was the individual atom. To illustrate, he displayed a year-old image of one of the smallest field-effect transistors made to date, on which three layers were visible: at the top, a poly-silicon electrode used to turn the transistor on and off; next, an intermediate layer of silicon dioxide insulator whose thickness, no more than 10 atoms across, is critical to its performance; and at the bottom, the silicon crystalline lattice of the wafer, its individual atoms showing up distinctly (see Figure 6). "What happens," he asked the audience, when the force behind Moore's Law—the ability to make the transistor ever smaller—drives progress to the point that engineers "start tripping over atoms?"

As this era arrives and the era of transistor miniaturization's going hand in hand with improvements in lithography ends, new alternatives for realizing different kinds of nanostructures are making their way to the fore. "Instead of building transistors in the bulk out of atoms and depending on their properties as bulk materials," Dr. McQueeney explained, "we have to start thinking about using atoms themselves to try to do some of the computations or to build devices on an atomic scale." Already, information technology and the life sciences have begun to converge: "The only manufacturing technology we know of in the entire scien-

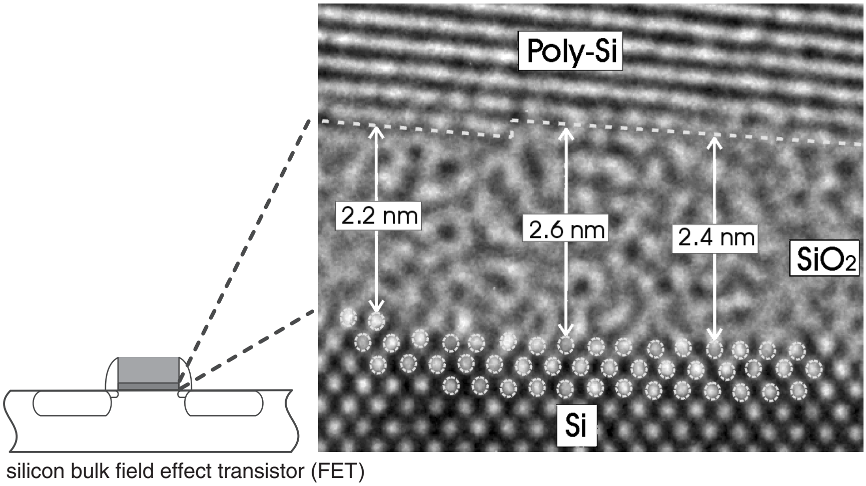

Oxide thickness is approaching a few atomic layers

FIGURE 6 Image of smallest field-effect transistors made to date.

tific field that can manufacture devices in high volume where the atoms are assembled in precise locations on the atomic scale," he noted, "is the replication-of-DNA process that's used in biology." Even while admitting that there is "nothing like that in the world of the technology that we use today," he suggested that the audience take comfort in achievements like that recorded by the image displayed before it.

But the transistor, even if its performance has grown at an impressive compound annual rate of between 16 and 20 percent, is merely one element of the IT picture. "Is there a Moore's Law at the system level?" was the way Dr. McQueeney framed the question before the symposium. To begin answering it, he cited 50 percent as the number that's "kicked around" for the compound annual growth rate (CAGR) of single microprocessors, listing device design, chip design, packaging, the microprocessor core, and the core of the compiler as fertile ground for innovation. Moreover, since computing advancement less typically occurs in single microprocessors than in parallel arrays of them—whether small arrays in the case of PCs or large arrays in the case of Department of Energy supercomputers—there is room for innovation in shared memory system performance, development tools, middleware, and applications as well. This brings the CAGR to 80 percent.

Dr. McQueeney then displayed a graph representing the compound annual growth rate in the performance of supercomputers as measured in Teraflops, or trillions of operations per second (see Figure 7). A series of five machines ordered for nuclear-weapons modeling under the U.S. Department of Defense's Accelerated Strategic Computing Initiative (ASCI) traces a CAGR of 84 percent. Several other machines, which he considers "not fully general purpose" because their hardware is designed for a unique problem, describe a somewhat steeper curve: Deep Blue, the chess computer that beat Gary Kasparov in 1995; a Japan molecular dynamics machine, Riken MDM; and Blue Gene, another special-purpose machine IBM expects to build by 2005. Rounding out the chart was a comparison of animal brain speed with computer speed, which showed the 11-Teraflop ASCI machine built for Lawrence Livermore National Laboratory to be on the performance level of a mouse and the fastest desktop computer on that of a lizard. The human, meanwhile, was positioned a factor of 10 higher than "the fastest machine that we can envision ever possibly being able to build using today's technology projected forward," Dr. McQueeney said. Using these comparisons to characterize technological advance, he posited that every 20 years supercomputers grow faster by the ratio of human intelligence to lizard intelligence, or a factor of about 10^4. "The kind of intelligence that we can apply to business problems," he summed up, "is growing at an alarming rate."

Even as such performance growth has made it possible to "do all kinds of interesting things in optimizing business processes in real time," it has carried along with it what Dr. McQueeney labeled "a complexity problem." He recalled "a big company that processes credit card transactions" that, driven by a need to

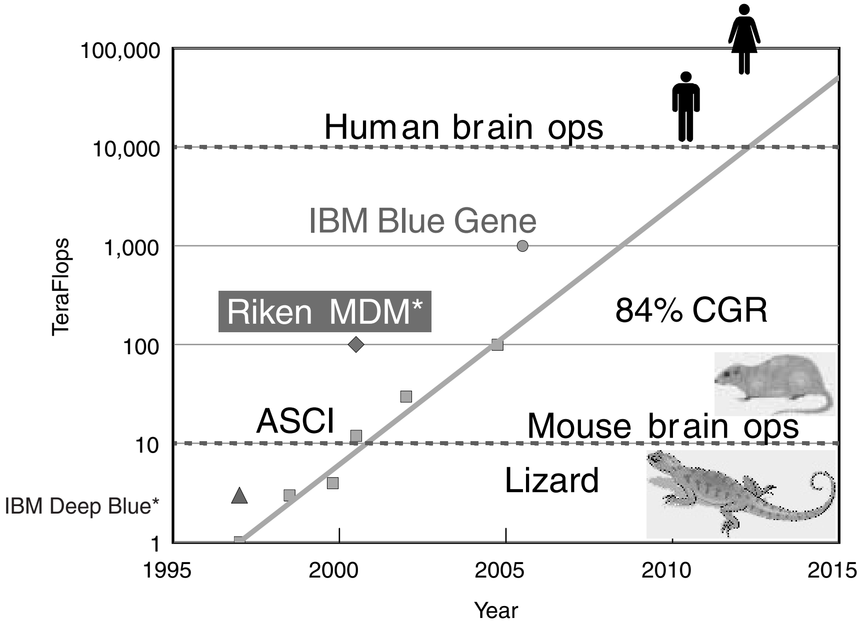

* Special purpose machines, i.e. chess, molecular dynamics, protein folding.

FIGURE 7 Compound annual growth rate in the performance of supercomputers.

eliminate downtime entirely, had pieced together a computer system so complicated that only three of the company's employees worldwide understood it well enough to manage it when it showed signs of breaking down. The firm came to IBM because the last time its system's performance had begun to degrade, two of these three, working on it together, were unable to pinpoint a cause. Although the in-house experts kept the system from going down by "making small changes to its configuration parameters until it started working correctly again," they could not explain afterwards how they had solved the problem. The senior VP in charge of the operation then decided to bring in a firm that could "assemble a lot more expertise over the problem" because, as he put it, "we've built a system so complex that we're having trouble assuring its reliable operation and maintenance."

A second major problem that has arisen with performance growth is what Dr. McQueeney called "efficiency of deployment." Looking at use of computing capacity averaged over a 24-hour day, he offered the following "typical" figures: for a PC, well under 5 percent and perhaps as low as 1 percent; for a UNIX workstation acting as an applications server, between 50 percent at the bottom and 80 percent at the very top; and for a mainframe, somewhere in the 90 percent range.[8] The reigning pattern is thus one of good utilization of central resources, poor utilization of remote resources, and "something intermediate" between. A factor in these numbers is the tendency to overprovision that is inherent in the model on which delivery of information across computer networks is based. For example, to be sure it had Web hosting capacity sufficient to handle the increased interest it expected in reaction to a Super Bowl advertising spot, a company might have to provision the front end of its site to accommodate a spike reaching 50 or even 100 times average demand.

In 1998 IBM began to investigate developing a more efficient delivery model for the middle to low end of IT resources based on aggregating different customers' demand for the front end of Web service, simple computation, and other services that can be handled by generic applications. Checking data its customers furnished against Bell Labs' late-nineteenth-century analysis of the statistical fluctuation of demand on telephone switches, the company found the patterns observed in the two instances to follow essentially the same mathematics. By 2001 IBM had achieved a vision of a new model in which many of the functions that had previously been part of an application written by an end customer would become capabilities that the public infrastructure or a set of middleware provided. This model would allow IBM to "build more value" into applications and code them in a more organized way, while at the same time free customers' applications writers from focusing on the internal details of connections between machines so that they could concentrate more on the value of the business logic to the enterprise. In January 2002 IBM launched a service business supplying computing capability on demand, much in the way public utilities supply power, water, and telecommunications.

Meanwhile, IBM had been on the trail of what turned out to be the new model's "missing piece"—Web Services—through an investigation of how scientists and engineers deal with complexity. "What they typically do," Dr. McQueeney said, "is try to step back and understand how to encapsulate the complexity of lower levels and only expose interfaces above that." IBM also surveyed the history of the PC platform from the days when it comprised only hardware, DOS, a file system, and an application through the advent of the windowing system for presentation and on to the addition of middleware to handle databases and online transactions.[9] The applications area of the mainframe, he noted, grew in parallel, with middleware put in to handle the layer between the applications and the operating system. "It was all about hiding the complexity at layer N underneath layer N–1 and just exposing a simpler interface," he explained.

---

[8] This does not, of course, imply an anomaly. Use of personal computers 5 percent of the time may be comparable to other common household appliances.

[9] Middleware is software that connects two otherwise separate applications, or separate products that serve as the glue between two applications. It is, therefore, distinct from import and export features that may be built into one of the applications. See CREN glossary at *www.cren.net/crenca/glossary/cpglossary.html*.

With the development of Web Services, according to Dr. McQueeney "a way to find, describe, and execute business functions on the Web," new questions arose:

- "Doesn't Web Services provide yet another layer near the top of this stack that lets us further abstract the applications and again have another simplification of the coding of business processes?
- "Is there an analogy [to] the traditional layering of the computer from hardware to operating system to middleware to the applications interface, which we could call a virtual computer?
- "Can we in fact treat the Web as a virtual computer? Does it have computing power? Does it have storage? Does it have I/O? Does it have applications with Web Services?"

The existence of Web Services, he noted, "completes the protocol stack for the Internet and for the Web and lets us think of that, at least from an engineering-design point of view, as if it were a virtual computing platform."

This allowed IBM to rationalize "Good Computing," which Dr. McQueeney described as "all about tying together back-end, high-performance computing resources in a way that then aggregates that image up to a scientific and technical user—and, in the future, a business user." And it made room for another innovation, called "autonomic computing" by analogy to autonomic biological systems, which devotes 20 percent of computing power to self-configuration, self-diagnostics, and self-maintenance. The thinking behind this, as summarized by Dr. McQueeney, was: "Let's not route all the raw performance [accrued thanks to Moore's Law] to the end users, because we have already overwhelmed them; let's turn some of that back inside to make the system easier to manage."

In closing, Dr. McQueeney turned to a change in the business environment for information technology that IBM had observed in the previous year or two: Buyers had begun demanding that the supplier assume a larger share of the responsibility for the return on investment from purchases of IT goods and services. Looking back to the early days of computing, he said customers took it upon themselves to integrate components they bought, to build applications for them, and to add the business value. Then, with increasing technological integration and the emergence of the layered structure he had outlined, fewer orders came for equipment alone: "'I don't want [even] really good components,'" he recalled customers saying. "'I want to buy integrated, end-to-end applications'—meaning, 'I want you to integrate all the IT so that it works end to end, but we'll still be responsible for the bridging of that to business value.'"

Lately, he said, customers have taken yet another step, so that their refrain might be characterized thus: "'I really don't want to get involved with the connection between even the end-to-end IT and the business process. I want to design the business process—and I'm still responsible for my business—but I want more of your risk and your value, Mr. IT Company, to be connected to whether that end-to-end integrated system produces the net result. I don't care if it can do a transaction, I care if it improves my market share. I don't care if it does a fancy optimization of my supply chain; I care if my inventory turnovers go from five to six. And I'd like to measure you that way.'" IBM's "On-Demand Computing" developed out of the company's efforts not only to respond to this shift in the needs and desires of its customers but also to enable its customers to adapt to shifts in their own customers' requirements without rebuilding their entire systems.

Dr. McQueeney listed and commented on some key technical elements of this offering:

- it provides customers, who IBM believes will insist on open systems, choices in components;
- it delivers higher levels of integration, relieving customers of handling that integration themselves;
- it is virtualized, so that a customer can blur boundaries between companies—for instance, running a business process out to a key supplier and back;
- it is autonomic, because system complexity has driven up total cost of ownership, as so much expertise is needed to handle management and maintenance.

As an example of how it works, he ran through a schematic representation of the management of a company's 401k plan. Web Services would be invoked when, for instance, an employee wants to check on the status of his or her investment elections or to make changes in the mix, precipitating a work-flow process that goes across company boundaries. There is a second tier of connections, to suppliers of investment vehicles like stocks and mutual funds, at the level of an external administrator, where services are aggregated. "Those relationships are specified by a quality-of-service metric stated in business terms, not IT terms," Dr. McQueeney stressed. "It's not 'so many processors and so many gigabytes;' it's 'how many tenths of a second of response time is acceptable on a transaction like this?'"

What has been accomplished is, in fact, the virtualization of interactions at a business-process level: Pieces of business process have been picked and chosen, then swapped in and out so that a cross-enterprise business process has been assembled "on the fly" using Web Services. "And if we want to then change that business process," Dr. McQueeney stated, "we can respond very rapidly, in real time on an on-demand basis, rather than having to rebuild the system." Additional advantages are the concealment of complexity from the user—who is mainly interested in the financial performance and security of the plan, not in how the transactions are executed by a supplier—and strong security that not only protects the system against hackers but protects individuals' privacy by making sure that participants get access only to the data that they require to do their job.

In summation, Dr. McQueeney declared: "Yes, we have a Moore's Law continuing at the device level. Yes, we have a stronger Moore's Law at the system level. And, yes, that produces a lot of tremendous benefits and will continue to do so." He emphasized, however, that how efficiently computer power is delivered, which delivery model or models are used to deliver it, and how it is exploited to create business value merit the industry's careful consideration. As to the choice of metrics, he advocated looking at "how many return-on-investment cases from individual businesses justify big IT investments?" and "how well did they pay off?" He suggested that academia or government might have an interesting role to play here as a trusted third party that would aggregate and analyze the pertinent business data, which tends to be more sensitive than technology data and which companies would be more reluctant to share with one another, then report back in such a way that no company's competitive position would be compromised.

DISCUSSION

Moderator:

Steven Landefeld

Bureau of Economic Analysis

Dr. Landefeld began the discussion by mentioning that he was struck, upon looking at one of Dr. Triplett's tables of features used in computer hedonic functions, that it dealt largely with individual devices and their characteristics to the exclusion of "big picture" qualities such as integration, reliability, and downtime. He raised the question of how researchers might move beyond the characteristics in existing models to measure systemic performance characteristics. The simple addition of component parts may underestimate or overstate worth and is an inadequate method of valuation.

William Raduchel, a member of the STEP Board, noted that hardware and software on their own, according to "the old rule of thumb," seldom make up 5 percent of the cost of building a large system such as that Dr. McQueeney referred to in the example of the credit-card processing company. For this reason, he pointed out, indexes used in current metrics account for a "tiny, tiny portion of the total cost" at the level of the actual business process. Moreover, the remainder of the cost goes into the books as general and administrative [G&A] expense, so that "95 percent of the cost of a system that might totally change the business process for a company is [considered] G&A overhead and doesn't show up as investment." Dr. Raduchel asked Dr. McQueeney whether IBM keeps records of project cost breakdown that would make possible a study of the true economics of a collection of large systems-integration projects.

Dr. McQueeney, while expressing reluctance to enter into a discussion of accounting methods, acknowledged that raw data are available that could be assessed as Dr. Raduchel was suggesting by a "suitably insightful person." But he stated that, whatever technical categories the outlays might be placed into, those making investment decisions "do look at the big picture." In addition, he said, such data reflect "the competitive capabilities—in fact, usually the core competitive capabilities"—of the customer, and their sensitivity would prevent IBM, and most likely discourage the customer as well, from sharing them. It was for this reason that he had proposed that the academic or government community might have a role to play in such research.

Dr. Raduchel, returning to his earlier point that hardware and software go into national accounts as 5 percent of the value of the large systems in which they are incorporated, asked whether it is then accurate to consider information technology to be only 5 percent of the overall economy in line with an earlier statement by Dr. Jorgenson. "If you took the full value of all the projects—that is, multiplied it by 20—then [IT is] no longer a small sliver [of the economy]," he said, adding that measuring their impact in aggregate could be an interesting factor in forecasting productivity. Noting that such investment had "dropped dramatically" over the past two years, once the investment crest powered by Y2K had passed, he expressed the concern that the commonly accepted five-year productivity prognosis might be overly optimistic.

Dr. McQueeney said he has observed that customers' interest in improving their business processes has not waned but that they have been taking a different approach designed to avoid net cash outflow. "They will frequently come to us and say, 'I desperately need this business-transformation project, and what I'd like to do is to package that along with an efficiency-of-the-infrastructure project that will generate some operational savings that I can then reinvest,'" he explained.

Dr. Raduchel, returning to the subject of complexity, noted that America Online, where he spent three years as Chief Technology Officer, employed over 20,000 servers located in "four rooms about 10 miles apart" and linked by over 100,000 physical connections. He commented: "Nobody truly knows what every one of those connections does, I assure you."

Dr. McQueeney responded by relating a recent experience he had had with a large agency in Washington whose desire to use "really cheap unilevel hardware" had led it to add another appliance server for each new function—"to the point where [its] mid-range system infrastructure has 4,700 servers." In his opinion, consolidation through the use of common images would allow the same functions to be handled by between 20 and 50 servers. "Think of the savings in managing the complexity of that," he suggested.

Dr. Raduchel in turn brought up the case of Google, which builds all its own hardware from Taiwanese components. That such a course "would turn out to be the optimal solution for somebody" indicated, he said, that "the world is changing in ways that no one would have predicted."

Mark Bregman of Veritas Software Corporation expressed the concern that the metrics used by the industry may be on the way to irrelevance because of changes in customer behavior. Just as the amount of copper wire sold to make generator coils is probably obsolete as a measure of the electrical power industry, he argued, "looking at microprocessors and disk drives as a way of measuring IT value really does cause an increasing amount of the cost to show up as G&A or overhead. When someone goes to an IBM and pays one fee for the whole service utility, they really are capturing all that in an investment in IT business value; when they buy chips and boards and assemble them into boxes and you only measure the cost of the chips and boards, all that other investment looks like overhead." Dr. Bregman singled out the appropriate placing of an aggregation point that moves very rapidly as one of the main challenges in looking at the information technology industry over a period of 30–50 years. "It's not just a matter of looking at the whole stack," he stated.

Dr. Landefeld noted that statistical agencies like BEA struggle with this very problem. "We are trying to measure the value of investments in in-house software," he explained, "but we can't value it in terms of the value of the output—the cost of the inputs is the best we can do."

Dr. Bregman pointed out that changing such definitions makes it hard to compare over time—which, he acknowledged, is the statistical agencies' "whole game."

Victor McCrary of the National Institute of Standards and Technology asked Dr. McQueeney what metrics the industry uses to evaluate its scientists. Noting that researchers who work at the pre-development stage have traditionally been judged by their publications, presentations, and patents, he said his IT lab is seeking other ways to evaluate both short- and long-term research. He commented that notions of "economic value added" and "business value" have "worked back into the R&D community."

Dr. McQueeney corroborated the importance of this issue, stating that IBM has worked "incredibly hard" on it for three decades, during which the marketplace has increasingly provided "inputs to the core research portfolio and core research deliverables." But while IBM's effort to align research with development and product is not new, over the previous five years "the influence of the customer has been reaching further and further back into our product plans, into our deep development, and, ultimately, into our research," he said. Within the previous six months the company had made known that key researchers would take part in consulting engagements in critical areas of research, particularly those concerning optimization and mathematics, "partly to deliver value out to the customers from our research lab, partly to bring a worldview of the bleeding edge of the customer environment back in to influence the portfolio." IBM scientists, he said, now "understand that they're expected to have world-class credentials and world-class standing within their technical communities but that the end goal is impact on the business and, therefore, on our customers."

Panel II

Computer Hardware and Components

INTRODUCTION

William J. Spencer
International SEMATECH, retired

Dr. Spencer, invoking the Silicon Valley saw that "the cost of every integrated circuit will ultimately be $5 except for those that cost less," observed that rapid decline in semiconductor cost and consequent growth in the speed and density of processors and memory have allowed software and systems designers "to get very sloppy." Therefore, as a technology wall looms for hardware, productivity can be expected to increase in software and systems, where a great deal of capability remains untapped.

Endorsing Dr. Jorgenson's proposed road map for computers, software, and communications, Dr. Spencer recalled a 1991 Dallas meeting at which 250–300 engineers assembled for two or three days under SEMATECH's leadership to get the Semiconductor Road Map off the ground. Today the project is run by the consortium's successor, International SEMATECH, with its original editor, Linda Wilson, remaining in charge, but participation by engineers and scientists has increased by an order of magnitude, to around 3,000. And whereas the Roadmap started out in a paper version that was revised every two years, it is now in electronic form and receives continuous updates. Placing SEMATECH's expenditure for data collection in the early 1990s at $1 million annually, Dr. Spencer remarked that the Roadmap is not an inexpensive endeavor.

He then promised the audience that the following two presentations, on microprocessors and magnetic storage, would go to the heart of computer performance. These important technologies, including LCDs [liquid-crystal displays], have provided productivity advances that have "driven the rest of the electronics revolution." Recalling a past prediction that a clock speed of 2 GHz would make possible relatively good voice-recognition capability, he noted that the industry was getting close to this and called upon William Siegle of Advanced Micro Devices to chart the future of processors. Thereafter Robert Whitmore of Seagate would talk about magnetic storage, which, Dr. Spencer said, has been progressing even more rapidly than semiconductor capability as measured in cost per bit.

PROCESSOR EVOLUTION

William T. Siegle
Advanced Micro Devices

Dr. Siegle began by crediting the Semiconductor Roadmap for the speed of the information technology industry's recent advance. He offered two reasons for what he regarded as a direct causal connection between the road-mapping process and the acceleration that had taken place in the decline of logic cost:

- Making meaningful improvements in capability requires the coordination of many different pieces of technology, and the Roadmap has made very visible both what those pieces are and what advances are required in different sectors of the industry to achieve that coordination.
- As companies believe that staying ahead of the Roadmap is a component of success and strive to do so, the existence of a published Roadmap heightens competition.

If the industry feels it is moving ahead too rapidly, he jested, "we should just stop publishing the Roadmap for a little while and descend back into chaos."

The aim of his talk, Dr. Siegle announced, was to survey the evolution of the microprocessor and to pin down the factors responsible for it in order to determine what must be nurtured and sustained so that similar progress might be achieved in the future. Focusing on the previous 10 years—which he considered a representative period, and which coincided with his direct involvement in the area at Advanced Micro Devices (AMD)—he posited that advance had resulted from improvements on four fronts: the architecture of microprocessors; the tools that are used to translate architecture into a physical design that can be implemented; fabrication technology; and the equipment and materials sector, the vast infrastructure supporting chip makers. Before launching into a discussion of microprocessors, however, he warned that computing progress depends on achieving equally significant gains in other areas of hardware and in software as well: "While microprocessors are important, you can't make meaningful systems and applications if there are advances in just the microprocessor." He then pointed to a chicken-or-egg question concerning the relationship between hardware and software: Have advances in hardware been driven by applications; have the latter expanded to make use of the hardware capability available; or have the two evolved together?

To dramatize the extent to which microprocessors have developed in little more than a decade, Dr. Siegle compared two AMD-made chips of the sort found in the IBM PC: the Am386, a 386-class product introduced in 1991, and the Opteron, which was scheduled to go onto the market in April 2003. He cited operating frequency, although in his opinion it is an imperfect measure of performance, as having increased more than 50 times, from around 33 MHz to 2 GHz, while the transistor count was jumping 500-fold, from 200,000 to 100 million. The Am386's transistor gate length of 800 nm made it AMD's first logic device to feature a submicron transistor; the chip, rather small by 2003 standards at 46 mm^2, was the company's first with 32-bit data handling as well. The Opteron, in contrast, specifies 60-nm transistor gates and a 180-mm^2 die size while supporting 64-bit processing of both instructions and data. Dr. Siegle noted that the Opteron's gate length is far smaller than 130 nm, the nominal dimension for its technology generation.

Turning to architecture and design, he said significant progress on these fronts has been a necessary accompaniment to improvements occurring in fabrication technology. Analogizing to the high-level concept that a building architect works out based on the needs and wishes of a prospective homeowner, he defined computer architecture as a high-level model of data flow and data processing that enables a prescribed series of instructions to be executed; although not a blueprint from which a device can be built, it is still a necessary starting point. He described the design process as "a matter of translating that architecture into something that a 'fab grunt' can build." Designers, he said, see a concept in the form of a set of mask patterns that will build multiple layers of a semiconductor process; their task is to translate that architecture into electrical models and an actual physical layout of the shapes that will be implemented.

During the interval separating the Am386 and the Opteron, improvements occurred in a number of areas detailed in Figure 8: the efficiency of instruction processing; memory hierarchy; branch prediction, which Dr. Siegle described as a way of dealing with deviations in the orderly flow of executing instructions; and functional integration, which he explained as "ways of spending the increased integration available by putting more and more function on the processor chip." These advances in architecture—most of which have taken place in the United

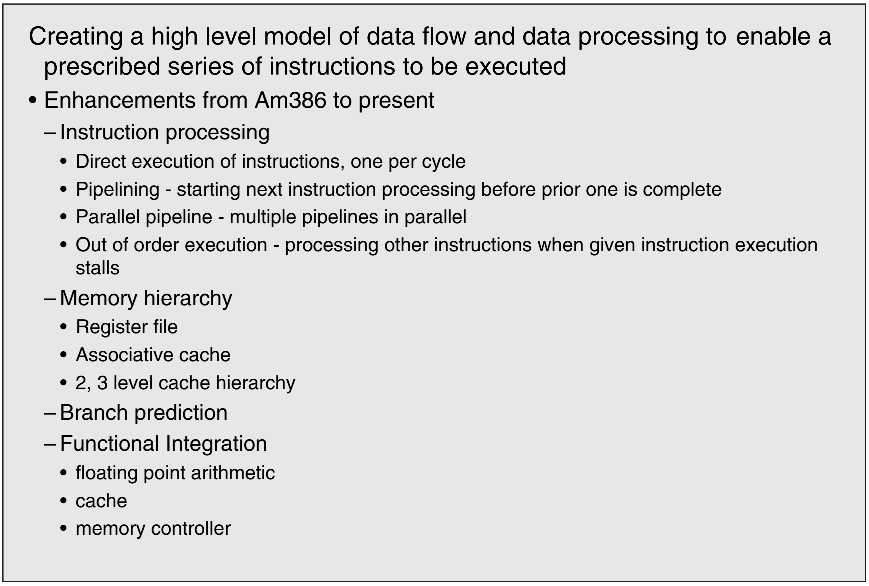

Creating a high level model of data flow and data processing to enable a prescribed series of instructions to be executed

- Enhancements from Am386 to present
  - Instruction processing
    - Direct execution of instructions, one per cycle
    - Pipelining - starting next instruction processing before prior one is complete
    - Parallel pipeline - multiple pipelines in parallel
    - Out of order execution - processing other instructions when given instruction execution stalls
  - Memory hierarchy
    - Register file
    - Associative cache
    - 2, 3 level cache hierarchy
  - Branch prediction
  - Functional Integration
    - floating point arithmetic
    - cache
    - memory controller

FIGURE 8 Architecture.

States—he credited not only to computer firms, “whose business it is to make advancements of this nature,” but also to early-stage work by creative thinkers in university computer science departments. In addition, there have been a number of concepts applied in the microprocessor business, among which he singled out memory hierarchy, that in some cases had been implemented 20 years earlier in mainframes and were borrowed from them.

One change in the design process most notable from AMD’s point of view came in the move from what had basically been transistor-level designs to high-level models that describe a design and are employed in tandem with a variety of tools that excel in providing functional verification that those models will in fact do what is intended. Dr. Siegle characterized several other advances as startling in their effectiveness as well: improvements in tools used in electrical timing and “for the placement of functions so they can be wired up”; increased ability to put the burden on the fabrication area to provide more layers of wiring to hook up all the transistors in place; and the use of higher-performance computers to do the very heavy-duty simulation and computation work that is required to support such design efforts. Looking back on the first version of Opteron, he recalled being “almost dumbfounded that first silicon was good enough to give samples to Microsoft to begin playing with.”

Joining with chip makers and academic institutions in moving design capability forward have been design automation (EDA) companies that supply device makers with tools. Dr. Siegle stressed, however, that in-house capability—"the ability to optimize the technology, to exploit it"—is necessary to get a truly competitive microprocessor. So rather than relying entirely on EDA companies, AMD and, he speculated, most other device manufacturers develop hand-honed design tools of their own that work their way to the general community after a period during which hardening up makes them accessible to "a perhaps less sophisticated set of users." Similarly, in the fabrication end of the business, working with captive capability offers advantages unavailable through an arm's-length relationship with a foundry. And just as foundries, which flourish in Taiwan, are situated largely offshore, there is an increasing tendency for software and design centers to be set up in various parts of the world to take advantage of talent and costs that are more advantageous than they might be in the U.S.—where, nonetheless, most new technology continues to be generated.

Dr. Siegle then turned to fabrication technology, which he described as the realization of the designer's intent through a sequence of wafer-processing steps that is split into two stages. In the first, the development process, a recipe is developed, formulated, and perfected. After a number of trials comes the second, production phase, at which replication of a process is required at high volume. Testing and assembly of the finished wafer, a significant step in its own right, ensues; technical sophistication is growing in this area of manufacturing as well.

Returning to the comparison of the Am386 and the Opteron, Dr. Siegle pointed out that the older chip, a two-level metal structure, used a technique called "wire bonding" in which discrete wires ran from the periphery of the chip out to the package terminals and the subsequent pins. The Opteron, in contrast, has nine levels of metal interconnect that allow hooking up all of its 100 million transistors and can accommodate one megabyte of cache on the device itself. A result of the recent evolution in design and architecture, this change was made possible by a number of breakthroughs, one of the most significant being fine-line patterning. This was permitted in turn by dramatic changes in the equipment used in lithography, a process for putting patterns onto wafers, as well as in resist technology, a term covering both the photosensitive film placed on the wafer to be patterned and etching methods that transfer the pattern. One of the primary lithography tools, the projection stepper, moved in a decade from 436 nm and 363 nm illumination sources down to excimer-laser-driven sources at wavelengths of 248 nm and 193 nm. This improvement was accompanied by a rise in the stepper price from less than $2 million to more than $20 million, he noted, adding that today's price for the lens exceeds the 1991 price of a whole system. As an indication of the magnitude of the progress that has resulted from these and other changes, Dr. Siegle offered a remark he made at the beginning of the current decade: "If you had shown me a picture in 1990 of what we would be building at the turn of the century, I would have thought it was a page from science fiction." Underestimation of the rate of progress is not uncommon among technologists, who, he suggested, are not very good at seeing beyond the next two to three years and tend to err on the side of caution. "Those red bricks in the red brick wall that shows up in the road maps," he noted, "have a tendency of going away as we get closer to them and understand a little bit more."

Dr. Siegle then took up interconnect planarization, a technical breakthrough he rated as equally significant. The thought that this technology, which calls for putting the wafer into a machine and rubbing on it to make it planar, could be used for the fine patterning of wafers initially left people in the field "aghast," he recalled. In fact, chemical-mechanical planarization, in commercial use at AMD from the mid-1990s, enabled an almost unlimited number of layers of wiring to be constructed, whereas the process employed just a few years earlier for doing a mere two and three levels had been relatively painful. While aluminum continued to be used for wiring, tungsten contacts were employed to make connections between the layers, and the planarization method ensured flatness between levels. This approach allowed designers to add layers to the chip, something that Dr. Siegle speculated may well continue almost *ad infinitum*.

With the move to planarization, the stage was set for switching from aluminum to copper, which further unblocked the ability to miniaturize. Pointing to photographs of such copper-interconnected AMD devices as the Opteron and its predecessor, the Athlon, Dr. Siegle called the audience's attention to the planar repetitive structure, featuring connections between layers that make the layers largely indistinguishable from one another. In the miniaturization of devices, he stated, the smaller the design grid at the transistor level, the more levels of wire will be needed to make the connections because one is growing as a square and the other is growing only linearly, creating a fundamental imbalance that requires more and more levels of wiring. So, without this breakthrough, it would have been very difficult to exploit the fine-patterning capability that has resulted in improving transistor density.

Changing the basic materials used in building devices has been critical as well, because the transistor does not just go faster automatically as it is made smaller. In order to keep resistance as low as possible at the transistor-gate level, changes have been made repeatedly: The tungsten silicide material used in the Am386 had evolved into titanium silicide by the mid-1990s and into cobalt silicide by the late 1990s, and it was expected to evolve again, into nickel silicide, over the next few years. The material that insulates one wire from another in the interconnect space has for a long time been "rather plain" silicon dioxide, but in the last few years a significant amount of energy has been poured into reducing the rather high dielectric constant of that material. First came a transition to a fluorine-doped oxide and then, more recently, to a more complex carbon-based organic material; both were aimed at driving down the dielectric constant and driving up performance by reducing the RC characteristics of the network.

Additional changes are on the way. With no more than a few atomic layers separating the gate from the silicon substrate, leakage of current from the gate to the substrate has become a serious issue—one that, according to Dr. Siegle, "has the potential to really represent a spanner in the works"—and needs to be accounted for in the circuit design. Since the materials in current use had been pushed about as far as they could go, researchers were working with more exotic alternatives, such as hafnium-based oxides, looking for a way to continue improving electric field at the silicon surface without the material becoming so thin that it would totally break down.

And progress has been required in other areas: in junction engineering, a term that applies to tailoring the actual profiles of dopants in the silicon so as to ensure acceptable electrical properties as conducting channels grow shorter; in scaling down operating voltage, which has gone from 5 volts to around 1 volt; and in obtaining performance at lower levels of power, which has been accomplished by the introduction of silicon-on-insulator techniques.

Dr. Siegle then turned to infrastructure, which he divided into two areas, equipment and electronic materials. He contrasted the current industry landscape, which includes an extensive network of equipment suppliers that device makers have come to depend on, with that existing when IBM built in house most of the equipment made to manufacture its System 360; the transition, he said, dated to the 1970s. Besides being very mature and sophisticated, this network is global to the point that, with the possible exception of Japan, no region in the world has a set of equipment suppliers sufficiently complete that it can supply a fab. The network—which, he said, "every chip maker in the world is dependent on in one form or another"—spans the United States, Europe, and Japan.

Nowhere near as big in revenue, but equally critical to the device makers' progress, is the network of materials suppliers. They furnish starting materials such as silicon and silicon-on-insulator substrates, high-purity gases and chemicals for fabs, as well as the photoresists needed with reductions in wavelength and feature size and, one of the more fragile elements of the infrastructure, precision photomasks. "If you look at the basic mechanics of the photomask business, it's a rotten business, even worse than the semiconductor business," he said, while adding that, without photomasks, it would be impossible to translate designs into fabrication.

A significant rise in the cost of manufacturing facilities has been a predictable consequence of the escalation of the cost of equipment, and of control and automation technology required to reduce variability, that has accompanied the overall increase of sophistication in the industry. The price of a chip plant has evolved from a few hundred million dollars at the time AMD was building the Am386 to well in excess of $2 billion for a modern 300-mm fab. And other factors have driven up the price of manufacturing as well. The process sequence that in the early 1990s comprised 15 masking levels and was turned in manufacturing cycles of about 30 days had, by 2003, reached 30-plus masking levels and cycle times in the neighborhood of 70 days. As a consequence, "any kind of a blip" requiring a design change can set a manufacturer back by a quarter or two: the time to make a mask, crank it through, verify it, and ship the production. Also

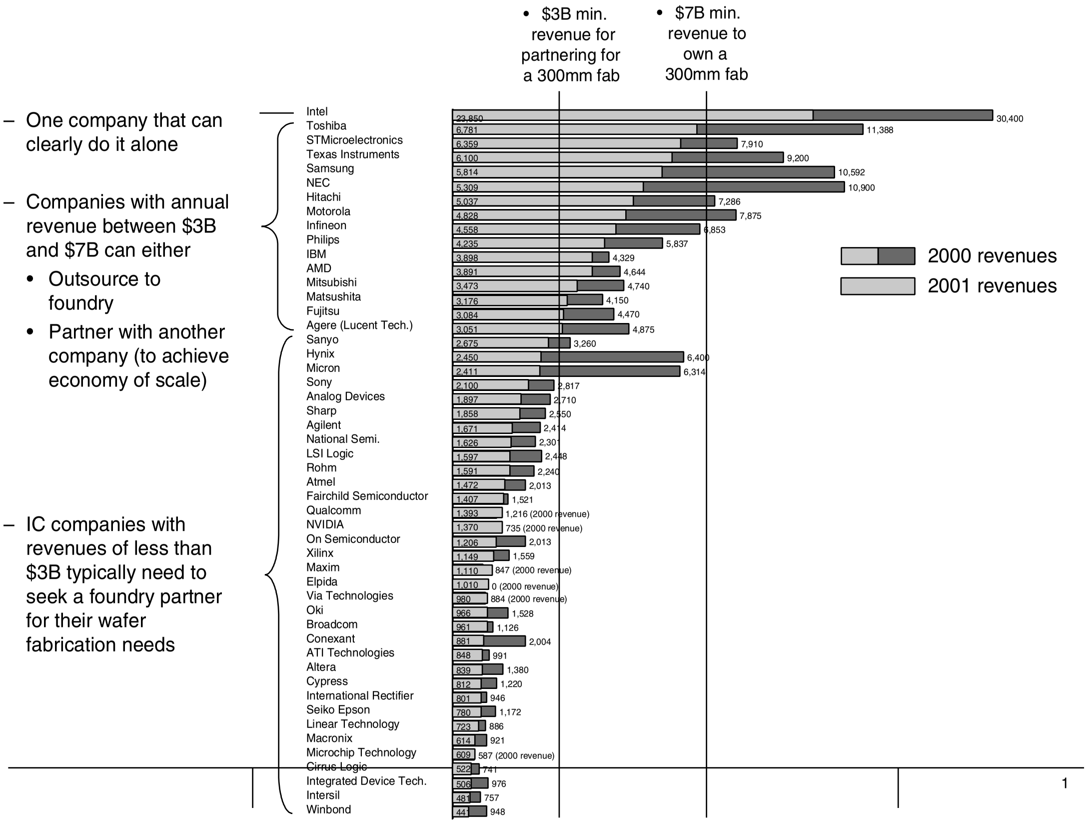

FIGURE 9 Factory investment issues.
SOURCE: Gartner, McKinsey analysis.

driving up cost has been the necessity of driving down the density of manufacturing defects in order to achieve economic yields. Defect density, defined as the number of die-killer defects per square centimeter—with a die killer being anything from a piece of dirt to an imperfection in a photo pattern that causes a circuit to be nonyielding—had to be reduced by around 200 times during a decade in which transistor density increased about 100 times and die sizes grew significantly as well.

The level of investment is increasingly becoming a barrier to entry into semiconductor manufacturing, providing a contrast to conditions in the industry's architectural and design sectors, where nothing carries a price tag equivalent to that of a fab. "If you want to own your own fab, you'd better have a revenue on the order of $7 billion [per year] to make the economics all work out," Dr. Siegle declared. "If you're willing to partner, you can get away with about $3 billion of revenue." Displaying a chart illustrating how much annual revenue a company now needs to be able to afford a megafab, he noted that only Intel among the world's major device makers would be placed in the former category based on 2001 revenues; if revenues for 2000 were used as the basis, the category would include another seven of the largest firms (see Figure 9). While numerous others would be able to afford such an investment in partnership, many more will have to keep older fabs in operation or rely on foundries if they are to remain in manufacturing. (For those entries in Figure 9 that are captives of larger systems companies, another option, of course, is for the system business to serve as "banker" to the captive semiconductor division.) AMD ranks among those who would need to partner on a new fab, "so the burden and the challenge of 'how do we get to this next level of fabrication plant?' is a nontrivial exercise for us," Dr. Siegle said, adding: "It's a pretty sobering picture."

But where have these "key enabling enhancements" come from? Attributing many of the conceptual advances to what are now largely regarded as old-line industrial research labs, he stressed the importance of acknowledging that major advances often have their roots in research done even a decade before they are applied. One example of this long latency is copper interconnect, on which work was well under way at IBM when Dr. Siegle left that company in 1990 but which did not reach the stage of commercial production until around 1997; similarly, work on cobalt silicide, a technology that didn't make it into production until the mid- to late 1990s, was taking place in the 1980s. "University research and research funded by our government colleagues has also been important," he said, "not only in supporting early work in new technologies, but also in creating—at least [in the case of] the university-funded research—human personnel feedstock for our business."

In addition, he praised the level of preparation for work on industry-relevant problems of young researchers who have gained experience at the Semiconductor Research Corporation (SRC) and, more recently, the Microelectronics Advanced Research Corporation (MARCO). "New graduate hires who have finished Ph.D.

work in some of these areas are real plug-in-and-run kind of talent that has been indispensable to enabling not only the growth of the industry but the increasingly challenging work that's being done," he stated. The national laboratories, with intense areas of specialization in such fields as modeling and plasma physics, as well as consortia—not the least of which being SEMATECH, widely credited with helping the turnaround of the U.S. semiconductor industry in the early 1990s—have also played a key role. The Defense Advanced Research Projects Agency (DARPA) has been a particular supporter of lithography, providing funds over a long period for many lithographic-based programs that have helped make progress in lithography and patterning possible. And the large, highly experienced supplier base has brought equipment advances as it has evolved into a widely dispersed, global resource maintaining a close relationship with leading chip producers.

Dr. Siegle then turned to technology integration, which he specified as referring to the combining of the individual elements he had been describing into a modern silicon device. A substantial task requiring a leading-edge fab, technology integration has generally been the province of the major chip makers; it usually has a development cycle of two to three years "after the long lead research is far enough along for integration." AMD and IBM had in recent months announced a joint development arrangement, a way for the companies to deal with the cost, which would be excessive for either on its own, of having an R&D facility to do the work needed to get ready for production.

Summing up, Dr. Siegle cited improvement in microprocessor capability over the previous decade of at least two orders of magnitude—but at comparable prices—in performance, integration level, and density. Those improvements resulted from parallel advances in architecture, design methods and tools, fabrication technology, and industry infrastructure. "Other than because of IP constraints, global dispersal of architecture, design, and infrastructure has already occurred or is likely to occur," he stated. And fabrication technology—whose cost of entry, already high, was growing—would potentially be open only to the largest companies, to multicompany arrangements, and to subsidized ventures. Chip makers' taking advantage of opportunities to participate in subsidized ventures and to contract with foundries in which they do not need to invest is behind the fact that such a preponderance of fabs is being built outside the United States. "Foundry arrangements are very compelling for many suppliers," he concluded.

STORAGE

Robert Whitmore
Seagate

Mr. Whitmore began by noting that, being responsible for all Seagate products in the final three years of their development up through production launch, he had a vested interest in the direction in which technology is taking the industry. He promised the audience a high-level view of a variety of topics, but he said he would begin by introducing the basics of the hard drive and offering an appraisal of the state of the business. Following that would be a review of storage trends; a discussion of the metrics used in the magnetic storage sector and of how they might be translated into a model that would be of benefit in assessing the path of computer productivity; and, finally, a look at what the future might have in store.

First, however, Mr. Whitmore introduced his company and provided a brief overview of the business of hard disk drives, which he described as "kind of the low item on the food chain." Founded in 1979, some two decades after IBM invented the disk drive, Seagate employed 49,000 in 25 countries in 2003 and posted revenue of $6.1 billion in fiscal year 2002. In that period, the company devoted about 10.8 percent of revenue to R&D and another $535 million to capital expenditure in order to sustain the strategy of total vertical integration under which it does everything from making its own wafers to fabricating the product and selling it to the final user. Manufacturing around 200,000 drives per day, a production level that requires world-class manufacturing and design capabilities, Seagate is the world's leading shipper of disk drives: It totaled around 18 million units in the final quarter of FY2002. The company handles development mainly in the United States, with one development center in Singapore, and manufactures mainly in the Far East, although it also has factories in the United States, Mexico, and Northern Ireland.

Turning to the basics of the hard disk, Mr. Whitmore said that, seen from above with its cover removed, a drive would appear as a number of rigid disks stacked one on top of the other with a small arm suspended over them. The fundamental technology that the industry has been scaling for the past half-century is the physics of sending a current to this arm that generates an electromagnetic field and changes the magnetics on the disk, a result he described as "pretty simple—basically, the 1 and the 0—but pretty tricky" to achieve. For, although the technology may appear macroscopic in photos, it is in fact very microscopic. The arm culminates in a head small enough that it would take six to cover a dime; it is able to read tracks equivalent in width to one-twelfth of the edge of a sheet of paper while literally flying above the disk at a distance of 100 atoms of air, or less than one micro-inch. To illustrate such spatial relationships and performance at a more easily pictured scale, he said that an equivalent head projected to the size of a Boeing 747 would be flying at Mach 800 less than an inch off the ground and have the ability to count the blades of grass as it went by. "Kind of a mind-boggling technology," he observed.

Highlighting the competitiveness of the business, Mr. Whitmore noted that the number of hard disk suppliers in the world had dropped from 62 in 1985 to seven by 2003; Seagate and two others are based in the United States, three are based in Japan, and one is based in South Korea. Similarly dramatic consolidation occurred among independent component suppliers over the same period. Amid this "craziness," as he termed it, success has required a large and highly qualified staff such as that of Seagate, which has 3,300 employees in R&D, a quarter of whom hold an M.S. or a Ph.D.; significant outlays for both R&D and capital expenditure; state-of-the-art manufacturing capability; acute focus on OEM relationships; and the ability to contend with both short product life-cycles and huge pricing pressure. Such characteristics have prompted Seagate's chairman and CEO, Steve Luczo, to call the disk drive industry "the extreme sport of the business world" and Clayton Christensen, author of *The Innovator's Dilemma*, to describe Seagate and its competitors as "the closest things to fruit flies that the business world will ever see."

This shakeout has nonetheless brought some stability to the sector, in Mr. Whitmore's opinion, and he opined that in the previous year pricing had begun to stabilize somewhat. Market share over the 61.3 million drives shipped in the final quarter of 2002 went 69 percent to U.S.-based companies: 30 percent to Seagate, 22 percent to Maxtor, and 17 percent to WDC. Of the remainder, Hitachi had 17 percent, Toshiba 5 percent, and Fujitsu 4 percent—giving Japan 26 percent in all—while South Korea's Samsung accounted for 5 percent. U.S.-based companies had 85 percent of the desktop computer market, which accounted for around three-quarters of the more than 200 million drives shipped in 2002, but their lead was only 59–41 in the enterprise (or business) market, and they were entirely absent from the mobile market, which he called "one of the biggest growing markets there is." And even as he signaled that his company would later in 2003 announce steps meant to redress the imbalance in the market for mobile-computer disk drives, he added that Seagate, owing to vertical integration, was the only component supplier left in the United States at a time when non-U.S. companies were investing heavily in components. Here he specified the market segments as ranging from high-end server products like high-transaction Web-based storage to such mass-market products as camera microdrives or the iPod digital music player, the consumer segment being "the up-and-comer." The requirements for the two segments differ, with high-end products focused on performance and reliability, consumer products on ruggedness, size, and power.

Moving to metrics, Mr. Whitmore pointed to capacity, price, performance, and reliability as the main factors for measurement, while saying that other metrics are appearing on the horizon. Beginning with capacity, he displayed a chart tracing the total number of bytes shipped annually by the industry from 1994 through 2002, which showed an evenly paced yet spectacular increase from around 1 million terabytes (TB) in 1998 to between 8 and 9 million TB in 2002 (see Figure 10). A subsequent chart, calibrated in petabytes (PB), projected a compound annual growth rate for storage capacity shipments of 62 percent between 2002 and 2006, which would take the total from between 8,000 and 9,000 TB in 2002 to nearly 60,000 TB in 2006; the number of storage units shipped was projected to post a 13 percent compound annual growth rate during the same five-year period

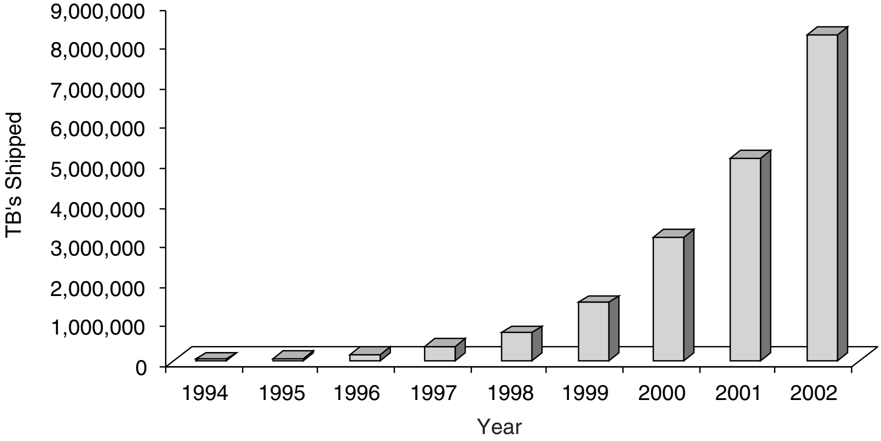

FIGURE 10 Capacity: Past storage growth.

(see Figure 11). "In the 1990s we just couldn't seem to satisfy the need for capacity," he said. "There is continued growth and need for storage, so it's not going away. It's really a question of how do we do it and what are the economics of it."

The next chart showed that the price of rotating magnetic memory on a dol-

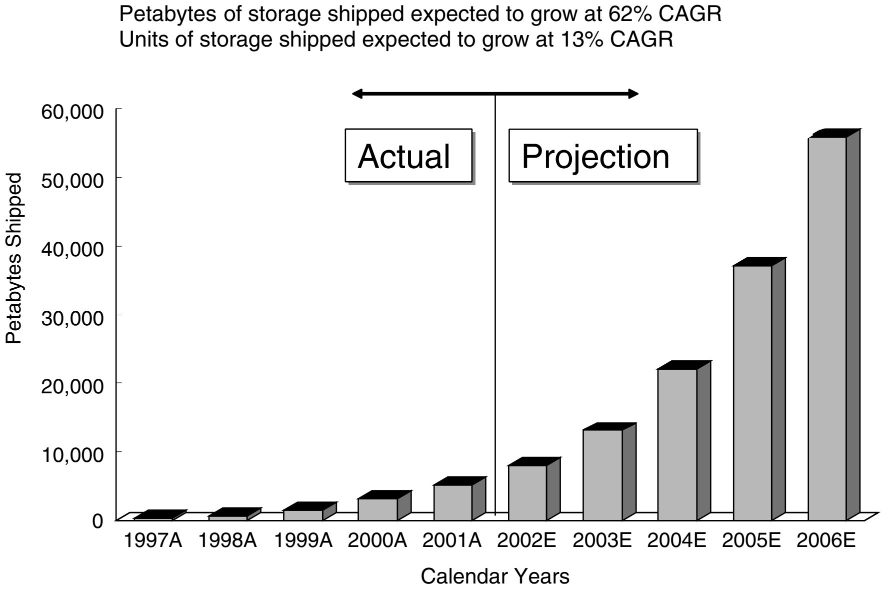

FIGURE 11 Capacity: Future storage growth.

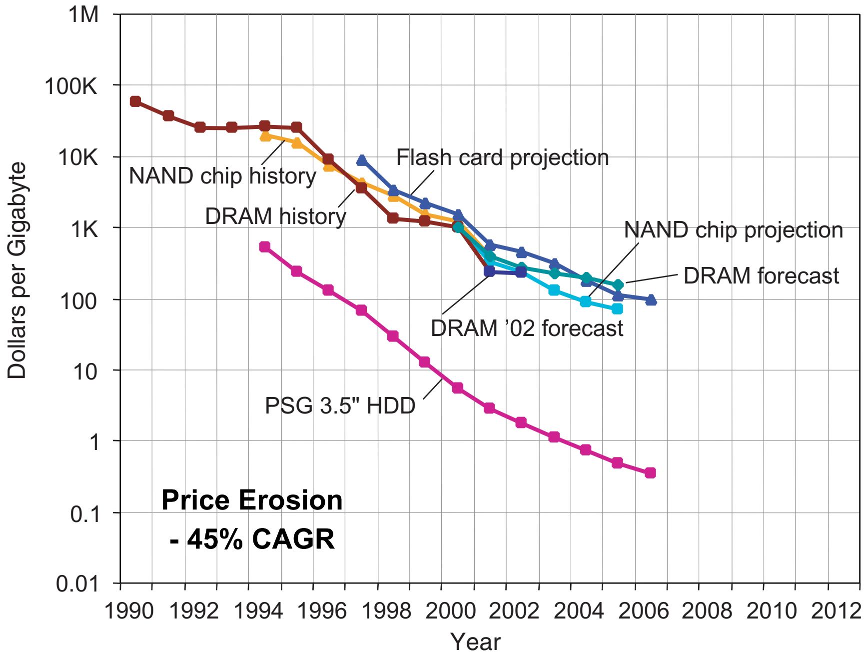

FIGURE 12 Cost: $/GB trends for desktop.

lar-per-gigabyte basis eroded at a compound annual rate of –45 percent between 1995 and 2002, with only a slightly gentler downward slope projected for 2003–2007 (see Figure 12). Silicon storage has been on a nearly parallel downward curve, but its curve began at a price between one and two orders of magnitude higher that that of rotary magnetic memory; the two curves have diverged rather than converged since the mid-1990s—a trend that, according to Seagate's projections, will continue.[10] While Mr. Whitmore acknowledged that using silicon storage will be appropriate at particular price and capacity points, he attributed the consolidation of the disk drive industry to the massive erosion in the price of its product, saying: "It's kind of 'the strong will survive' here, because we're going at a pace that's just burning people out."

Another chart indicated that input/output transactions per second (IOPS) had shown accelerated growth between the late 1980s and the late 1990s, when improvement leveled off (see Figure 13). Referring to the improvements made in this era as significant, Mr. Whitmore said: "Not only are you paying less, but it's

---

[10] The difference in the costs of magnetic storage and semiconductor memory is explained by a vast difference in speed of access to the information.

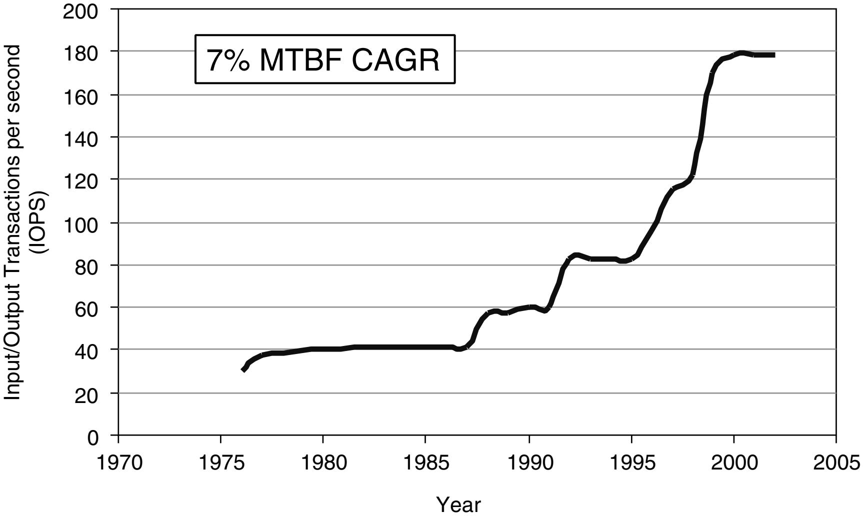

FIGURE 13 Performance: Enterprise growth.

going faster due to spindle speeds and seek times and a lot of different engineering techniques." Finally, coming to the important subject of reliability, he displayed a chart indicating that mean time between failures (MTBF) had grown at a compound annual rate of 25 percent from 1977 through 2001—which he called phenomenal—with progress accelerating in the late 1980s and leveling off around 2000 (see Figure 14). Product currently being shipped is speced at 1.2 million hours MTBF or higher, and the future may well see an MTBF of 10 million hours.

To introduce his discussion of the future of the industry, Mr. Whitmore displayed what he called "the bible for hard-disk storage and growth": a chart illustrating the history of areal density[11] (see Figure 15). Beginning in 1957, when IBM produced the first hard-disk drive, areal density has grown at 42 percent per annum, although growth was limited to 9 percent per year during the period 1975–1990, when it abruptly assumed the pace of 100 percent per year which continued through 2002. He said the plot is made up of a series of S-curves that were described as advances in physics and materials processing took the industry through a number of technology transitions, from ferrite heads to thin-film heads and then

---

[11] Areal density is the amount of data that can be packed onto a storage medium. Areal densities are usually measured in gigabits per square inch. The term is useful for comparing different types of media, such as magnetic discs and optical discs. Current magnetic disks and optical disks have areal densities of several gigabits per square inch.

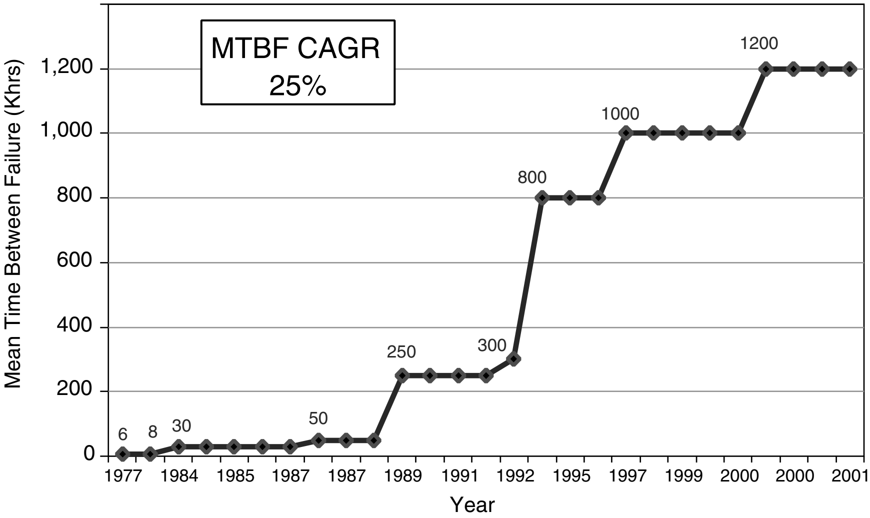

FIGURE 14 Reliability: Enterprise growth.

on to MR heads and to TGMR. The flattening out of the 1970s and 1980s was attributed to lack of demand rather than lack of innovation by Mr. Whitmore, who said that the advent of the PC drove invention and the "whole competitive madness" of the past 15 years.

The question now before the sector is whether growth is going to continue at 100 percent per year. At the moment the industry is in transition from the longitudinal orientation of the bit on the disk to the perpendicular. While the growth of longitudinal technology has flattened out, Mr. Whitmore said perpendicular technology had been demonstrated to work and production of drives incorporating it would begin shortly. He predicted that, with the growth phase of another S-curve coming up, the next few years looked "pretty good"—and that once this new technology started to run out of steam, another would be invented (see Figure 16). Already in sight as a successor to perpendicular is heat-assisted magnetic recording (HAMR), in which a laser beam will shine through the head, heating up the media to several hundred degrees Celsius. This new technology, which has been demonstrated in the lab with the help of funding from the National Institute of Standards and Technology (NIST) and is considered five years from commercial production, will theoretically extend the density of recording beyond 10 Tbit per square inch while leaving the process of reading the data from the disk drive unchanged. "The beauty of this technology," Mr. Whitmore stated, "is that it allows you to take all the stuff you've been doing before and go another round." After an expected five years of HAMR, the industry hopes to deploy a technology

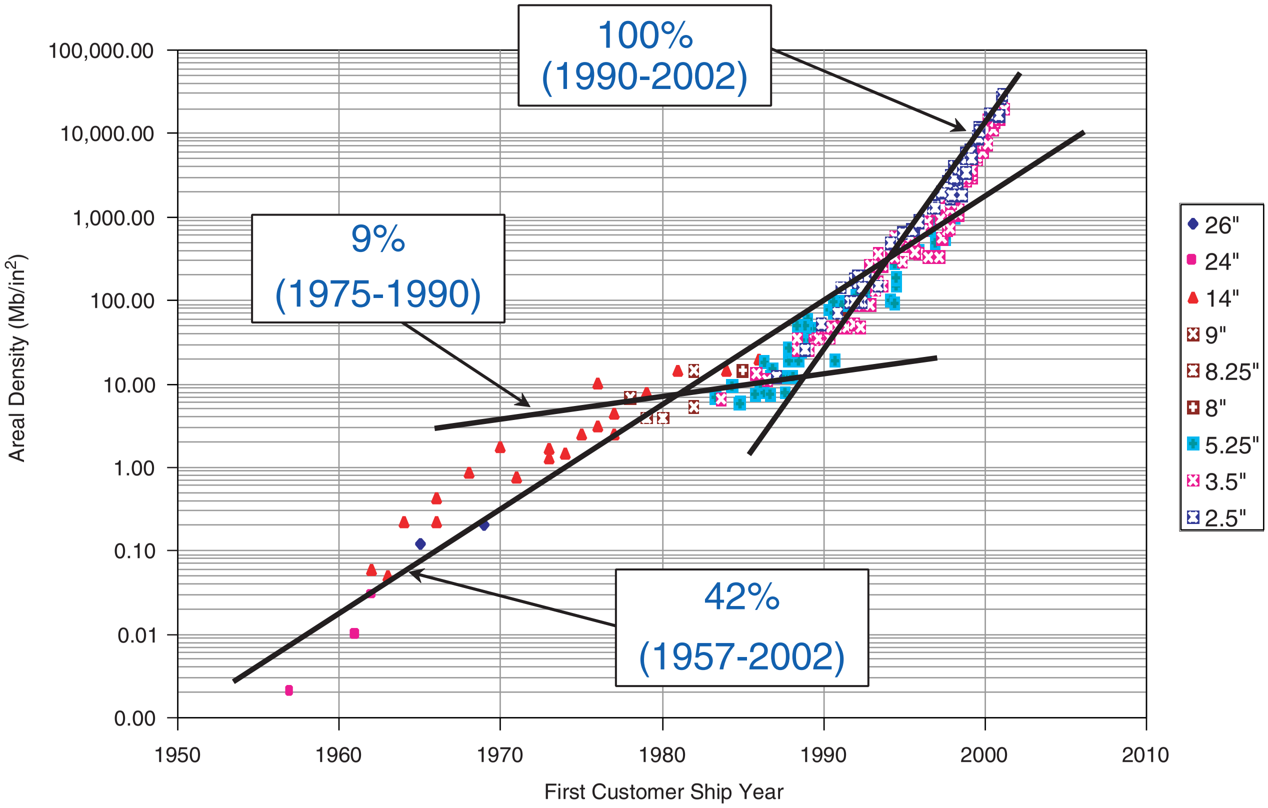

FIGURE 15 45 years of areal density growth.

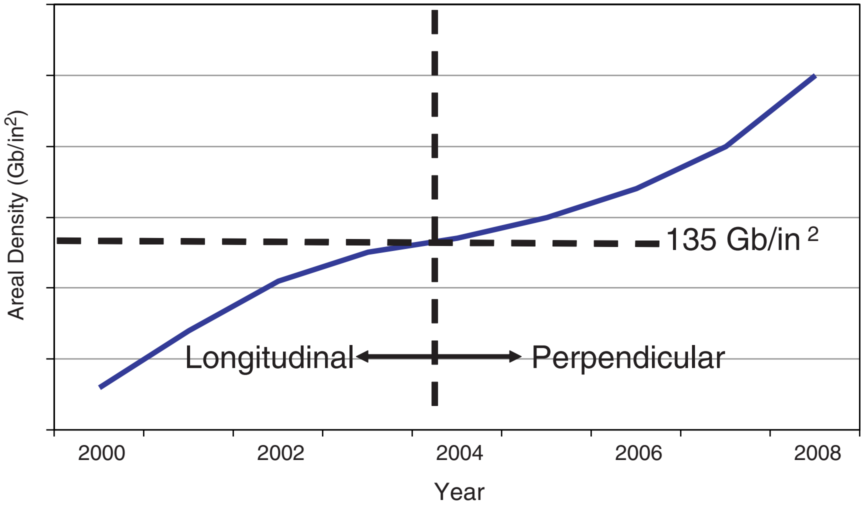

FIGURE 16 Recording technology transition.

known as SOMA for "self-ordered magnetic arrays," which is seen as extending recording density to around 50 Tbit per square inch.

Turning to applications, he acknowledged that the PC market is saturated, although there is room for growth in some developing-country markets, and that while enterprise storage is an excellent business with very good margins, it has its limits. He looked for significant growth, however, to broad applications in smaller-capacity, more consumer-related products: in mobile PCs, a market that is "just starting"; in various hand held appliances, from PDAs and personal audio devices to cameras and multimedia cell phones; and in external storage devices. In addition, computerizing the infrastructure of the home—having a server driving all the PCs and other electronics in the house, as well as putting disk drives in televisions and other devices—is in its infancy but is "becoming real and will continue to grow."

In conclusion, Mr. Whitmore expressed the hope that consolidation has reduced the number of competitors in the "brutal" disk drive business to a reasonable and sustainable level. "We are worried about non-U.S. based companies and their involvement," he said, "but feel we're strategically positioned to handle that." Continuing to invest heavily in R&D is key to Seagate's strategy, but the improved technology that results will enable growth only if it is employed in marketable products and applications for its use can be found.

DISCUSSION

John Gardinier, a self-described "retired science junkie," noted that he had seen an ad in a recent *PC Magazine* offering a $1,000 computer with RAID technology in the storage and asked whether Mr. Whitmore saw any sensible reason to employ RAID in a personal computer—and, if so, what it might be.

Mr. Whitmore explained that RAID is an architecture for arranging disk drives so that there is redundancy in data protection and pointed to different schemes, from a higher-end enterprise system to a method of packaging exemplified by Google that calls for arranging drives in a low-cost system. While acknowledging that the technology was very hot at that moment, he said he did not know why it would be advantageous for use in a personal computer. What those using it were really after was an inexpensive integration scheme that could compete with products offered by companies like EMC and Sun, which are more expensive.

Kenneth Flamm of the University of Texas at Austin, a member of the Steering Committee for Measuring and Sustaining the New Economy, prefaced his question by asking how many of those in attendance had backed up their home computer in the prior 60 days. He cited the lack of hands as a possible explanation for a market for RAID.

Taking the point of view of an economist attempting to ascertain the contribution of the PC's components to its functionality, Dr. Flamm noted that the technical details, while interesting in that they reveal something about where the industry is going in the future and how industry insiders think it is going to get there, are utterly irrelevant from the point of view of the user, who does not care how many layers of interconnect are on the chip. He therefore asked Dr. Siegle to speak more about the functionality that has been added to the microprocessor above and beyond the pure improvements deriving from fabrication technology. He then put the same question to Mr. Whitmore regarding the disk drive, observing that when prospective buyers look at disk drives, besides the size of the disk they are interested in rotational speed and the amount of cache included on the disk—proxies for IOPS, which vendors calculate but do not supply to users.

Mr. Whitmore responded by outlining contrasting business strategies, both of which Seagate is planning to pursue: making the disk drive more powerful by putting such features as multiple performance interfaces and cache onto it, and making it cheaper by taking functions off. By way of example, he said that Linux or the logic of the MP3 player might be placed on the hard drive. "These are things that are very available," he noted, adding that "the difficulty in getting them is more the relationship with the OEM suppliers we deal with than the technology to do it."

Dr. Siegle observed that quantifying true performance is no simple matter, even at the processor level, and that it is probably best done by using a variety of benchmarks. Noting that frequency increased 50-fold over the period he had addressed, he said that about a third of the improvement was attributable to the architecture and the other two-thirds to transistor speed. But, stressing that that is merely a description of frequency, he said that application benchmarks are needed to take into account how much work gets done with each click of the clock. "AMD has started to use a performance metric for marketing processors in order to get away from merely the speed-related issue because of the difference between processors and how much gets done in a given clock cycle," he said. In his judgment, the subjects of measuring performance and of where the various contributors to performance improvement come from need additional attention.

Michael Borrus of the Petkevich Group pointed to earlier remarks that only 5 percent or so of a PC's capability is employed in most applications and that Google's processing capability comes not from high-priced, high-performance components but from less sophisticated components cobbled together to achieve high functionality. He then asked the speakers: "Is there, in your industries, a different technological and market trajectory associated with the lagging edge which those at the leading edge are not paying attention to, but which the [STEP Board's Measuring and Sustaining the New Economy] project ought to be paying attention to because of the impact it could have on the economy?"[12]

---

[12] Technical and market trajectory means a line of technical advance that delivers products with specifiable cost/performance parameters that are sold to a specifiable market of customers. Most technology industries are characterized by conventional, accepted lines of technical advance and most firms in the industry produce products premised on that line of development. Leading edge typically then refers to the latest generation of products delivering the latest advance in cost and performance (usually high performance at an initial high cost). In that context, "lagging edge" refers to cost/performance characteristics that are far away from the conventional leading edge—e.g., potentially much cheaper or with quite different performance characteristics—and thus typically used for completely different purposes.

The simplest example is the line of technical advance that Moore's law characterizes, in which processor speed doubles every 18 months or so, resulting in a well-established technical and market trajectory for microprocessors with specifiable performance and features whose relatively high initial costs decline with the scale of production and that are sold at predictable, declining price points over time to PC makers and other customers. Intel's newest, most advanced microprocessors would then characterize the leading edge, typically produced with the latest, most expensive process technology, and capable of outstanding performance.

By contrast, one example of a lagging-edge trajectory can be found at Berkeley and other places in work on simple semiconductors that can be printed on plastic using cheap laser or real-to-real printing techniques rather than the very high cost, capital-intensive process used to produce microprocessors. These are potentially very low cost and low performance but usable for simple sensor networks embedded in structures, toys, product tagging, and other applications that could never afford a leading-edge microprocessor. A lagging-edge trajectory can still be quite innovative as plastic semiconductor concepts surely are—it is lagging only in the sense that it is aiming for very different cost-performance points than the leading edge of the accepted line of technical advance.

The lagging edge can lead to whole new industries with profound economic impacts, or can disrupt established industries. In this sense, the scrap-iron processing minimills were a lagging-edge technical trajectory 25 years ago, capable of producing only a very limited range of steel products with

In response, Mr. Whitmore returned to his comment that emphasis is being placed on taking features off the hard drive. Sticking to the minimum sophistication needed could boost manufacturing efficiency, because backing off on the technology can get the yields up and the cost down at a faster rate. But while need for continued lower cost drives removing features on the one hand, on the other hand applications still exist that require higher processor speed. This "bifurcation," he said, indicated that there was a business model for both paths.

Dr. Siegle added that, if enough bandwidth becomes available to link the systems that are being used only 5 percent of the time, the potential computing resource will be enormous. He saw the gating issues as getting adequate bandwidth to those systems and people being comfortable with others using their unused cycles.

Dr. Spencer observed that manufacturing of hard drives had moved almost entirely out of the United States and that semiconductor manufacturing was rapidly following along the same path, with foundries, most of which are abroad, taking more of the business. He asked whether, as that occurs, American universities will attract people to work in those areas who will be able to provide the kind of capability that Dr. Siegle described and that Mr. Whitmore indicated are already available in the magnetic storage area.

Mr. Whitmore answered in the affirmative, saying that although Seagate moved its manufacturing offshore long ago, he had not seen any lack of need in the United States for technologists in design, research, or manufacturing, and he did not anticipate that changing. While the need for higher-level skill sets in the magnetic storage industry had been flat, it had by no means tapered off.

In contrast, Dr. Siegle called attracting enough U.S. students into university programs that are relevant to the semiconductor industry a 20-year-old problem. "Somehow we've managed to deal with that adequately," he said, but he added that "the hazard here has been that we have become dependent on foreign nationals who are coming to our universities, being trained, joining our work force, and it's becoming increasingly attractive for them to go back home." A certain level of capability needs to be retained in the U.S. if its firms are to remain on the leading edge of the business.

---

inferior quality compared to the huge, scale-intensive basic oxygen furnace steel making of leading Japanese producers. But the minimill trajectory evolved, becoming more and more capable and competitive with traditional steel-making techniques, eventually disrupting a large chunk of the steel market.

Panel III

Peripherals: Current Technology Trends

INTRODUCTION

Kenneth Flamm
University of Texas at Austin

Dr. Flamm, praising the quality of the panel's speakers, called on Mark Bregman to begin the session.

THE PROMISE OF STORAGE SYSTEMS

Mark F. Bregman
Veritas Software Corporation

Dr. Bregman pointed out that he would be speaking about software, asserting that the area had been overlooked in the discussion of productivity in the computer industry. He argued that this omission accidentally amplified Dr. McQueeney's point that, to see where productivity gains are coming from, it is necessary to look above the hardware level.

Introducing his company, Veritas Software, he said that with about $1.5 billion in annual revenue it is the number-one player in the storage software business. This sector covers all software used to help store, access, or manage data, providing among other things better data availability, data protection, and disaster recovery. Because its software products were independent of any specific underlying hardware platform, he noted, the business stayed clear of the debates over RISC vs. SISC, Sun vs. Intel, or the merits of various operating systems. He characterized software storage as a fairly well-defined yet large niche and said that it had posted a "dramatic" 41 percent compound growth rate over the previous five years.

In his own attempt to deconstruct the computer, Dr. Bregman recounted, he had begun with Moore's Law, then turned to the communication between the elements that do the computing as reflected in Gilder's Law: Bandwidth grows three times faster than computing power.[13] This latter formulation, he commented, "really just says that bandwidth is growing dramatically—getting cheaper, more available—and that that's driving something." He then cited a yet-unnamed law stating that storage achieves 100 percent growth in density annually, saying that translated into better cost at a dramatic rate. "The improvement at the storage element level is happening faster than it has been in the microprocessor level over the last several years," he asserted.

To illustrate the implications of this advance, he offered a quote from the November 25, 2002, issue of *BusinessWeek*: "E*Trade finished yanking out 60 Sun servers that cost $250,000 apiece and replaced them with 80 Intel-powered Dell servers running Linux that cost just $4,000 each." Such instances as this factor-of-50 improvement in underlying capital investment were proving a source of concern for suppliers of storage hardware, whose business had declined by a factor or three over the previous two years, reflecting a major change in the landscape of the marketplace for storage systems. "Just to be fair, not all of this decline is coming about purely because of disk-drive technology," said Dr. Bregman, who pointed in addition to changes at the storage-subsystem level arising from the combined influence of improvements in communications, in computing capacity, and in software that utilizes and links the underlying disk-drive technology.

Addressing the issue of actual business benefit derived by the customer, Dr. Bregman posited that simply quantifying end users' investments in information technology might not be the right way to gauge their levels of productivity improvement. He noted that the industry was pervaded by nervousness resulting from the recent slowing of the purchase rates of both software and hardware, which had reflected the vigorous growth of the economy through 2000. He sug-

---

[13] According to George Gilder, "bandwidth grows at least three times faster than computer power." This means that if computer power doubles every 18 months (per Moore's Law), then communications power doubles every six months. Bandwidth refers generally to the amount of information that can be transmitted over a connection over a fixed period of time. See *www.netlingo.com*.

gested, however, that it is not so much the amount of hardware and software acquired as the manner in which they are used that "really drives labor." Turning specifically to storage, he pointed to a 2001 study by the market-intelligence firm IDC indicating that 29 percent of total enterprise spending on storage went to hardware, with 13 percent going to software and 58 percent going to labor. Since it is classed as overhead, a good deal of the labor content is not measured very well, and yet "it's the driving factor in the investment to get productivity" from IT acquisitions, he argued, adding: "If you're sitting here as a CIO being told, 'Do more with less,' the big nut to crack is labor."

Given the amount of data currently stored in the United States and an estimate by Forester Research that a single storage administrator could typically handle about 700 gigabytes of capacity, Dr. Bregman put the number of storage administrators in the country at 8 million in 2003. This brought to his mind the old story that the telephone system was going to collapse because everyone was going to become a telephone operator, a threat whose only solution was to adopt automated switching, thus changing the paradigm. Automation was the solution in the case of storage as well, but it would come, in his opinion, from software rather than hardware.

In order to keep the discussion firmly focused on the business drivers rather than the technology drivers, he turned his attention to the issues that preoccupy corporate chief information officers on a day-to-day basis. The CIO, whose responsibility was more and more frequently extending beyond technology to the business itself, faced two conflicting, or apparently conflicting, demands. The first demand was to increase service levels, driven by both employees' growing need for online access to applications from inside the enterprise and—even more important—customers' growing need for access from outside. The second demand, particularly over the previous couple of years, was to shrink the IT budget. "'Do it all,'" is what Dr. Bregman said the CIO had been hearing from the CEO. "'Increase the performance, increase the service levels, give me a better result—but spend less money.'"

Expanding on the first demand, that service levels be increased in line with burgeoning need for online access, he posited that if a firm were running its payroll or accounting system using its own IT infrastructure, which wasn't visible to the employees directly, and the payroll ran and didn't complete, it could be run again—with the worst consequence being that employees would receive their paychecks an hour late. But in the case of an online transaction, going offline could mean losing the transaction, and perhaps even the customer. Officials with an options exchange in Chicago had told him that if they were offline for a matter of minutes, they would lose millions of dollars; if they were offline for hours, they would permanently lose a part of their business, which would go to a different exchange; and if they were offline for a week, they would be out of business altogether because the work would be taken up in Philadelphia, London, or elsewhere and would not come back. CIOs were thus being asked to raise availability and quality of service in an increasingly complex environment. "Throwing more hardware at the problem," Dr. Bregman declared, "doesn't solve that problem."

What this first demand, coupled with the second—the need to save money—translated to for CIOs of many IT organizations was the quest to get more efficient use of the hardware and storage they had while making staff both more effective and more productive. Software's value here, Dr. Bregman said, was that it could provide increased serviceability and service levels and at that same time produce better effectiveness and efficiency. While noting that the traditional approach to making a system highly available was to build redundant hardware and highly available hardware platforms, he said that now high availability was increasingly being provided with software—which made possible the use of lower-cost and potentially less reliable or available hardware and relied on software to manage overall system availability. This could be observed in those storage subsystems in which disk drives were put together with RAID and other software solutions.

But it could be observed in high-end servers as well, he said, citing as an example the changes made at eBay since a major outage had brought its auctions to a standstill several years before. "You don't hear that anymore," Dr. Bregman noted, asking: "Is that because somehow magically the servers have gotten to a much lower mean time between failure?" In fact, eBay's servers were continuing to fail just about as often as they did when the outage occurred, but the company in the meantime had set them up in a clustered configuration so that when one failed, operations went on without interruption.

Software could be used not only to shore up availability but also to improve performance. Dr. Bregman conceded that the notion might seem "odd"; experience with the PC, he acknowledged, suggested that each succeeding version of a software package ran more slowly and that only the advent of the next-generation microprocessor could restore or improve system speed. He stated, however, that high-quality software could in reality optimize efficiency, particularly in storage systems: Intelligent software can manage storage and improve performance efficiency, and "you can do things like drive the storage subsystem to look ahead and fetch information from a disk—which is intrinsically slower than DRAM memory—and put it into a cache." As an example of the dramatic difference software could make, he pointed to a Veritas file system designed to run on top of UNIX operating systems that, when substituted for the UNIX file system that Sun Microsystems shipped with its Solaris, provided a 15-fold performance improvement in "real-world, transactional applications" without additional hardware. He further pointed out that, although software is certainly not free, its price is far more negotiable than that of hardware because there is no intrinsic manufacturing cost, which he called "a real leverage point" and "something that has been overlooked in the debate about IT productivity metrics."

Finally, software could provide value by managing complexity, achieving lower cost both through labor efficiency—better tools, so that a storage adminis-

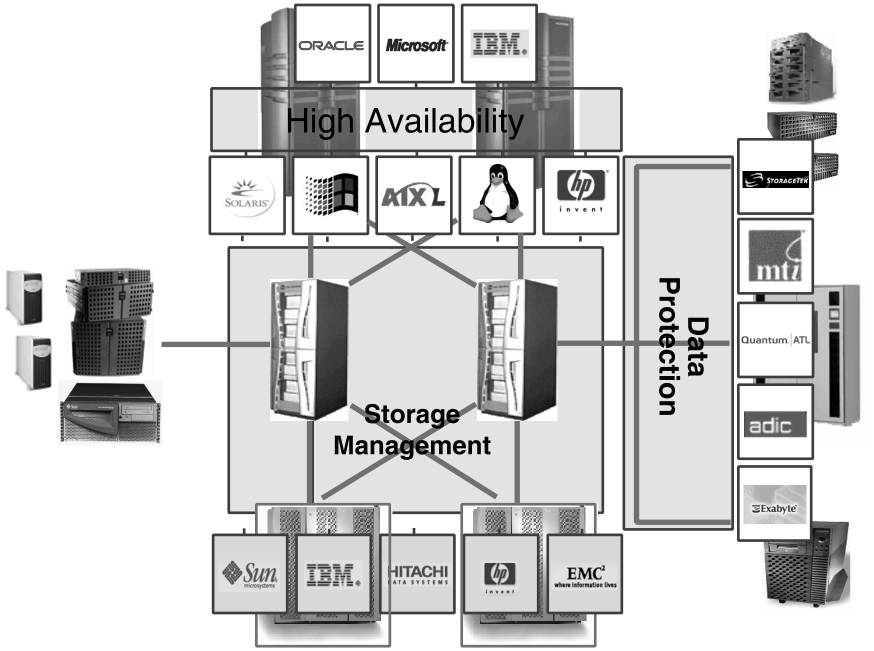

FIGURE 17 Better availability and performance.

trator could manage a distributed storage facility much larger than a mere 700 gigabytes—and through better hardware utilization. Typical storage utilization in enterprises was only about 40 percent of its capacity; firms could acquire very large amounts of storage as a result of buying another server for each new application required, but the storage remained in islands. Through technology whereby Veritas had put high-density storage together with high-performance communication, the entire storage capacity could be connected into a single storage network accessible from all servers, so that actual storage needs could be met across the network. Different categories of storage might be established, so that a high-performance application would be placed on a high-end array and low-performance needs could be met at lower cost. "The software can manage that in a way that takes the labor out of it at the same time," Dr. Bregman noted.

Displaying a chart he described as "my version of the complexity picture," he called the audience's attention to a server and its backup, interconnected for failover, that were partially hidden by the label "High Availability" at the top (see Figure 17).[14] Below them, in the middle of the field, were storage directors, de-

---

[14] The term "failover" refers to a backup operation that automatically switches to a standby database, server, or network if the primary system fails or is temporarily shut down for servicing.

vices that allowed connection of the storage shown at the bottom to any of the servers; these storage directors were also redundant in order that high availability and high throughput could be achieved. On the right-hand side were tape and other devices used for data protection, backing up the data and in some instances storing it offsite. All of these elements needed to be connected or managed through three categories of storage software: data protection software, the traditional backup or data-replication software; high-availability software, which manages the failover between systems and the dynamic workload provisioning between systems; and storage management software, which handles the virtualization of that storage pool, the management of those file systems, and the performance of the storage system. He called these three areas "critical to being unable to unlock the value that people have invested in by buying all this high-performance hardware."

Focusing on the storage capacity represented at the bottom of the graphic, Dr. Bregman noted that it could be a pool of heterogeneous storage of different categories acquired from different vendors—some being high-density storage array, some the low-cost, low-performance capacity known as JBOD for "just a bunch of disks"—connected to a network. Storage-virtualization software allowed applications to treat that storage as a uniform pool and to select what was needed. Newly emerging, he said, was server-virtualization software, which performs a similar function for the server pool, allowing applications to be distributed across that pool of servers and not only to be provisioned as needed but even to be provisioned dynamically as needed, "so if an application sees a spike in demand, it can grab more resources out of that pool, perhaps at the expense of an application of lower priority."

To conclude, Dr. Bregman presented a matrix showing eight markets in which Veritas and its subsidiaries were active: availability software for storage, servers, and applications; performance software in the same three fields; and automation software for storage and servers. His company led the market in all three varieties of availability products, he said, through clustering, backup, and data replication. The storage performance category was characterized by very specialized, very high-tech kernel-level software that he called "the software world's equivalent to microprocessor design"; a company Veritas was acquiring, Precise, offered similar products at the server and applications levels. Automation or utilization-improvement software provided automated server provisioning and automated storage provisioning.

Software of this kind, Dr. Bregman contended, had escaped the attention of those studying which metrics might reflect the relationship of IT investment to productivity. "My single message," he stated, "is that as we think about trying to understand what is driving the new economy—or simply the economy—in terms of IT investment and IT return, we should not lose sight of software, and particularly of infrastructure software." While it might be easier to think about application-level software, such as ERP [enterprise resource planning] application systems, as a business process, he warned, a great deal of value and innovation was being created in the system infrastructure software of the type that he had been discussing.

GRAPHICS

Chris A. Malachowsky
NVIDIA Corporation

Mr. Malachowsky, NVIDIA's co-founder and Vice President of Engineering, introduced his company as the largest fabless semiconductor company in the world, having been the first to reach $1 billion and then $2 billion in annual revenue. Employing around 1,600 and headquartered in Santa Clara, CA, NVIDIA focused on graphics and media processors—which he called a "volume niche in the computing space"—selling into the professional workstation, mobile device, consumer electronics, personal computer, and chipset markets. The company concentrated on the intellectual property that went into the design of its chips: It defined the formula according to which they were manufactured, but it outsourced their production, then forwarded the devices on to its customers. NVIDIA, he said, took responsibility for manufacturing quality without managing the manufacturing process directly. The risk of outsourcing was "mitigated somewhat when you keep your expertise in the area in which you've outsourced," he commented, adding that "when you take ownership of that, you're not losing a skill, you're just not investing in the infrastructure of that process."

Although NVIDIA itself was marketing a card it made for the workstation market, in the PC market it was selling integrated circuits that went onto add-on cards rather than selling the actual add-ons. The company designed a device along with the PC board it would go onto and sent its own components to a contract manufacturer, which procured the other needed components, put the board together according to NVIDIA's formulation, and tested it according to NVIDIA's standards. Having numerous distribution channels—some production was going to OEMs like Dell, Compaq, or HP, some to value-added resellers—had ensured the company a diversity it had found valuable in an economy in which customer segments could grow at varying rates. Recalling a juncture "when the branded PCs were challenged by no-name guys who were using similar components and selling cheaper," Mr. Malachowsky remarked that "when the economy got tough, people seemed to go back" to the Dells and Compaqs, perceiving them to be "a safer investment."

Turning to graphics as a technology, he remarked that, in contrast to the rest of the computer, the graphic element is one that "you can see," that allows "you [to] touch progress or sense progress." Although full accuracy in visual representation—which is complete correspondence between a visual representation and that which is being represented—had not been achieved, graphics capability had been improving—and, in any case, the goal of graphics is not literal accuracy but rather the ability to "fool the eye" into seeing what the artist wants to convey. Still, an enormous amount of computation was required to make "something that was a still photo last year into something that can spin and twirl and that you can interact with this year," Mr. Malachowsky noted, adding that Moore's Law had allowed his industry to do more with the resources available as time had progressed.

The film industry provided an illustration of how graphics was in the process of catching up with some of the other digital media: Movie-making chores that previously required farms of thousands of machines had become possible using consumer PCs at dramatically lower costs. Describing the advance in the sophistication of images as semiconductor complexity grew to the level of 120 million–130 million transistors per chip, Mr. Malachowsky noted that without the storage software Dr. Bregman had discussed, "we're nothing." He displayed a chart tracking the rapid increase in power and complexity through four generations, each separated from the next by about 12–18 months, of one of NVIDIA's product lines: the GPU, or graphics processor unit, in which independent logic elements predominate over memories and caches.

A second chart documented product performance improvements between the second half of 1997 and the first half of 2003, which ran at an annualized rate of 215– 229 percent. He touted the graphics industry's exploitation of not only semiconductor technology and architecture, but also algorithmic advancement, to further its development.

NVIDIA had moved its compute resources onto what he characterized as "small, cheap, no-name, small-form-factor, space-efficient processing" in 1U and 2U packages with dual or quad processors that had "a couple gigabytes" of memory on them.[15] The company was in the process of buying, every other month, about 250–300 of these units, which cost little enough "that you don't have to worry about fixing them"; and even though it was taking out older units, the computing capability of its farm was doubling every six months. "We have, I think, two guys that run 5,000 computers," Mr. Malachowsky said. "You plug a new one in, it configures itself, it downloads its operating system, it puts itself in all the queues, identifies itself to everybody—pretty impressive." The company was allocating its computing resources between a Sun farm—"traditional workstation-class compute resources"—and thousands of smaller Linux-based machines, and it was the capability of the Linux kernel[16] that was keeping the Sun

---

[15] This refers to the size of a relatively small server. A server is a host computer on a network. It houses information and responds to requests for information (for example, it houses Web sites and executes their links to other Web sites). The term *server* also refers to the software that makes the act of "serving information" possible. See *www.netlingo.com.*

[16] A *kernel* is the central module of an operating system. It loads first and remains in memory to control memory management, disk management, and process and task management. It's called a kernel because it stays as small as possible while still providing all the essential nutrients required by other applications. See *www.netlingo.com.*

farm active.[17] "As 64-bit computing comes about, the processing elements are available," he noted. "As the kernel gets stable, and the processing elements are pervasive, and manufacturers test basic PC components with very large amounts of memory, we'll switch over. All the pieces are there, but it needs to stabilize long enough for it to run." Jobs needing to run for four or five days were therefore directed to the Sun farm. "Twenty-gigabyte databases doing place and route of chips—not large by some standards, but you need it to stay running and you need it to have a large memory access, a large dataset," he explained. NVIDIA had, over the preceding three years, purchased $200 million worth of CAD (computer-aided design) tools, which are indispensable to the design business. It also owned $45 million worth of emulation, a technology used to map a design onto programmable hardware elements that then act like a chip. "It's slower," he conceded, "but it's operating on hardware and it gives you the ability to plug it into a real system."

According to Mr. Malachowsky, graphics was one among only three classes of semiconductors that could continually take advantage of Moore's Law without being regarded as a commodity, which he named and remarked on as follows:

- Processors: "You can always include more cache or provide multiple parallel processors on a die—give me more transistors, I can build you a better mousetrap"
- Field-programmable gate arrays (PGAs): "You give me more transistors, it will have a larger programmable device that can satisfy more applications"
- Graphics: "You give us more transistors. We'll do better algorithms. We'll do them more in parallel. We'll get more real. We'll offload more of the software burden."

Adding "too many more transistors" was, in the case of most other semiconductors, commoditizing: "You're competing just on price, not on the value." In the graphics sector, in contrast, transistors were consumed as quickly as they became available and performance increased commensurately.

Beginning a quick overview of the evolution of computer graphics, Mr. Malachowsky noted that expensive machines were required to produce images of the pre-1987 first generation, which nonetheless consisted of little more than line segments in a "wireframe" design. In the second generation, dating from 1987 to 1992, more colors and solid backgrounds were added in the interest of verisimilitude. Even though emphasis on literal accuracy was limited and the viewer would

---

[17] A *farm* is a group of networked servers that are housed in one location in order to streamline internal processes, distributing the workload between the individual components. In other words, a server farm expedites computing processes by harnessing the power of multiple servers. See *www.netlingo.com*. In this instance, the speaker is referring to Sun Microsystems' network.

not have confused these images with reality, he remarked, "You could now get lost in it. You could start focusing on the training aspects of the flight simulator, or exploring geographies, or whatever you were using it for." Introduced in 1992 and lasting through the end of the decade was the succeeding, third generation, in which "texture mapping," described by Mr. Malachowsky as a "decaling process," allowed textures and not just colors to be imposed on surfaces so that images approximated nature. With hardware doing all the mathematics, a two-dimensional picture could be transformed, mapped, scaled, rotated, and projected onto a three-dimensional surface, with the result that the eye believed it was seeing something real.

Through these first generations, semiconductors captured specific algorithms in order to make the images' appearance conform to the preferences of mechanical engineers at large industrial corporations who used CAD. Application writers in NVIDIA's industry, which produced video games, wanted their own look, however. "They didn't need it to be real," Mr. Malachowsky commented, but "they wanted shadows to be emphasized, things to be glowing, things to look one way or another." To achieve this, they built customized algorithms in software on top of the standardized hardware. A more recent trend was to put programmability into the hardware: Once the hardware, its fixed function removed, was made more flexible, developers could get the distinctive look they wanted and still take advantage of the acceleration offered by the wealth of transistors.

The phase marked by programmability started off with a very low-level language; on the order of an assembly language, it was both primitive and tedious. In the previous year, however, a number of companies had come out with compilers which operate on higher-level language descriptions of the effects desired, which raised the level of abstraction, making programming at this level easier and more widely available. This allowed developers to focus on the look and feel they desired in their applications while leaving implementation to the device manufacturer. Mr. Malachowsky offered three-dimensional graphics as an example of what was possible with programmability. Three-dimensional graphics typically involved decomposing the three-dimensional object into triangles, which are planar and therefore suited to having algorithms mapped onto them. Pointing to the turning, rotation, and scaling of the vertices, he said, "You can do all sorts of math on the vertices, and then eventually you're going to figure out what all the individual elements or fragments within the triangle are." Once a program was written covering all of the fragments, different looks could be achieved simply by making changes within the program. The graphics processor now became a programmable unit that could mimic the capabilities of a proprietary software, and the added capability arising from the software had opened the door to increased creativity, he noted. "We still haven't re-created reality," Mr. Malachowsky admitted. "We're not even close, although we could fool you if somebody spent a whole bunch of time thinking about it before they got around to rendering things." He said that, to rise to another level of realism, many of the simplifications that had been effected to make things algorithmically tractable would need to be relieved.

Mr. Malachowsky was reluctant to make detailed predictions for the graphics industry, other than to say that its role as one of the interfacing elements of the computer ensures its further existence. Alluding to the rise in resolution of displays, the subject of the program's next talk, he said the higher the resolution, the more elements one might want to process, something that offered graphics a future beyond the PC "as more and more displays and computing elements find themselves embedded in all sorts of applications." The graphics industry would pursue the higher end of the market—with its "'hell-and-brimfire, make-it-the-best-you-can-at-any-cost' type of solutions"—because that was the segment driving innovation, which then bled into the mainstream product line. As the largest consumer of wafers from the world's largest fab, NVIDIA drove technology, said Mr. Malachowsky, adding that the graphics industry as a whole was "an important player in technology push." His expectation was that, since the graphics industry had not yet progressed that far down the learning curve, it was in for significant improvement over the succeeding 10–15 years, at which point semiconductors incorporating graphics technology would turn up in a wide variety of embedded applications.

FLAT PANEL DISPLAYS

Dalen E. Keys
DuPont Displays

Dr. Keys began his presentation by underlining the tendency of humans to be visual and by arguing for the enabling role played by the display for both software and graphics. But he joked that his conviction in the display's significance had been shaken earlier in the day, first when he found that only seven lines of Dr. Triplett's lengthy paper on "Performance Measures for Computers" referred to displays and again when David McQueeney proclaimed the display to be "good enough." He had decided nonetheless to carry on with his presentation, in which he was speaking not only as chief technology officer of DuPont Displays but also chairman of the U.S. Display Consortium (USDC).

Commencing an introduction of DuPont, which as he noted was not generally considered to be a display company, he posed the rhetorical question: "Why is it we're in this realm?" He pointed out that the 200-year-old company had started as an explosives company specifically focused on black powder, then a century later had become a chemical company, moving into polymers and energy as well as providing raw stock for its own chemicals. As it was moving into the next century, he said, DuPont believed displays to offer it an opportunity to build from some of the chemical and polymer capabilities it had developed over the previous century. "We have some background," he said, pointing out that the science of displays matched many of the company's skill sets.

In 2002 DuPont formulated five growth platforms—Agriculture & Nutrition, Electronics & Communication Technologies, Performance Materials, Coatings & Color Technologies, and Safety & Protection—in the second of which displays was housed. The Electronics & Communication Technologies platform comprised four business units: Electronic Technologies, covering electronic materials, where DuPont was the world's second-largest supplier; Fluoroproducts, of which an outstanding example was Teflon®; Display Technologies; and Imaging Technologies, where the company had focused on professional imaging for the printing industry. DuPont considered its experience—with components of color fidelity, printing, and pixelation, as well as with materials from its fluorochemicals business—an asset that it could bring to the display sector, which Dr. Keys described as being "in the midst of some technical revolutions" destined to enable graphics, software, and consumer applications. DuPont was active in several areas of displays: LDCs (liquid-crystal displays), although as an enhancer rather than as a fabricator; OLEDs (organic light-emitting diodes); thick-film materials, which it provided to plasma-display producers, being qualified with 90 percent of OEMs worldwide; and field-emission displays, for which the company retained enthusiasm because of new approaches opened up by the carbon nanotube.

Dr. Keys then turned to USDC, a consortium formed in 1993 and supported by the U.S. Army Research Laboratory, whose mission is to support its 15 member companies and their affiliates in building a world-class U.S. display industry. Both large and small companies, as well as a variety of display technologies, were represented among USDC's members, something that, he said, "creates some very interesting dynamics in the group." The consortium funds projects predominantly outside its membership companies to advance the technology manufacturing infrastructure.

Beginning the technical portion of his presentation by focusing on the interface of the display with the consumer, Dr. Keys praised as "great" the marriage of the cathode-ray tube (CRT) with the computer, which had held up for over a decade and determined users' expectations regarding image performance. Over time, however, portability became an issue and its solution, the union of the LCD with the PC, rated a "good" but no better even though it had enabled the laptop computer. "What about power?" he asked, observing that although the LCD's image quality was showing continual improvement, it was "not to where it needed to be." In addition, he faulted its design, pointing out that 80–90 percent of an LCD's light is wasted.

The display had come increasingly to represent the product of which it was a feature, Dr. Keys asserted. Recalling the cries of "'I love your computer!'" coming from classmates of his daughter as they got their first look at her new system, he stated: "They're not looking at the computer—they have no idea what the computer does—they're looking at the screen. That's also true for your cell phone: If the display goes out, you hate the phone; you want to yell at the service provider. And the same with your PDA." Out of this came a question: What, these days, is the computer? Data from Korea showed that much of the use of the mobile Internet through cell phones there was unrelated to making calls; instead, downloading pictures and other forms of entertainment, accessing travel and financial information, and e-mail were key applications—in all of which displays played a key role. The display's assumption of additional functionality, a trend that was expected to become more pronounced, sowed confusion about what the nature of the computer is and which market display manufacturers were attempting to serve. While the display accounted for about $53 billion of the communication industry's approximately $1 trillion in annual sales, Dr. Keys noted, the lines in the industry were being blurred.

This was particularly acute in such portable devices as the cell phone or PDA, in which the display takes up a lot of space. In addition to providing the user with information, there was a possibility that other functions might be built in. "Can it sense me as the user? Can it take a photograph? Can it generate power?" asked Dr. Keys, noting that these questions and their implications were generating considerable research activity. Although flat-panel display sales had suffered a dip of slightly more than 10 percent in 2001, they were back up by some 30 percent in 2002—a far stronger showing than that of the integrated-circuit sector—and the forecast for growth in subsequent years was positive. A negative for the producers, in contrast, was pricing. The fact that the consumer PC market was the lowest-price market was no source of contentment among the display fabricators, who had taken to looking to other applications and market segments in search of profitability. One advantage display makers enjoyed over chip makers was that the former could change their substrates almost on a routine basis—there had been seven changes over the previous 15 years—with the result that profitability was enhanced. Attempting to come up with an equivalent for Moore's Law that might apply to his sector of the industry, Dr. Keys detailed the increase in display size from one product generation to the next. For Generation 0, dated 1987, the figure he gave was 270 × 200 mm; for Generation 1, 1990, 300 × 350 mm; Generation 3, 1995, 550 × 650 mm; Generation 5, 2002, 1200 × 1300 mm; and Generation 7, 200x, 1850 × 2100 mm.

A chart on which various applications were located as a function of their image quality, or resolution, and their screen size demonstrated that the computer, whether desktop or notebook, ranked high in both parameters (see Figure 18). Because the technology and resolution were becoming similar in many applications, Dr. Keys said, it had become "easier to back off" technologically, and he cited as a trend that manufacturers "target the high resolution and back off." Showing a chart categorizing display technologies, he said LCD is a direct-view technology that is nonemissive, meaning that it emits no light itself but depends on a backlight (see Figure 19). While most familiar to those in the audience was

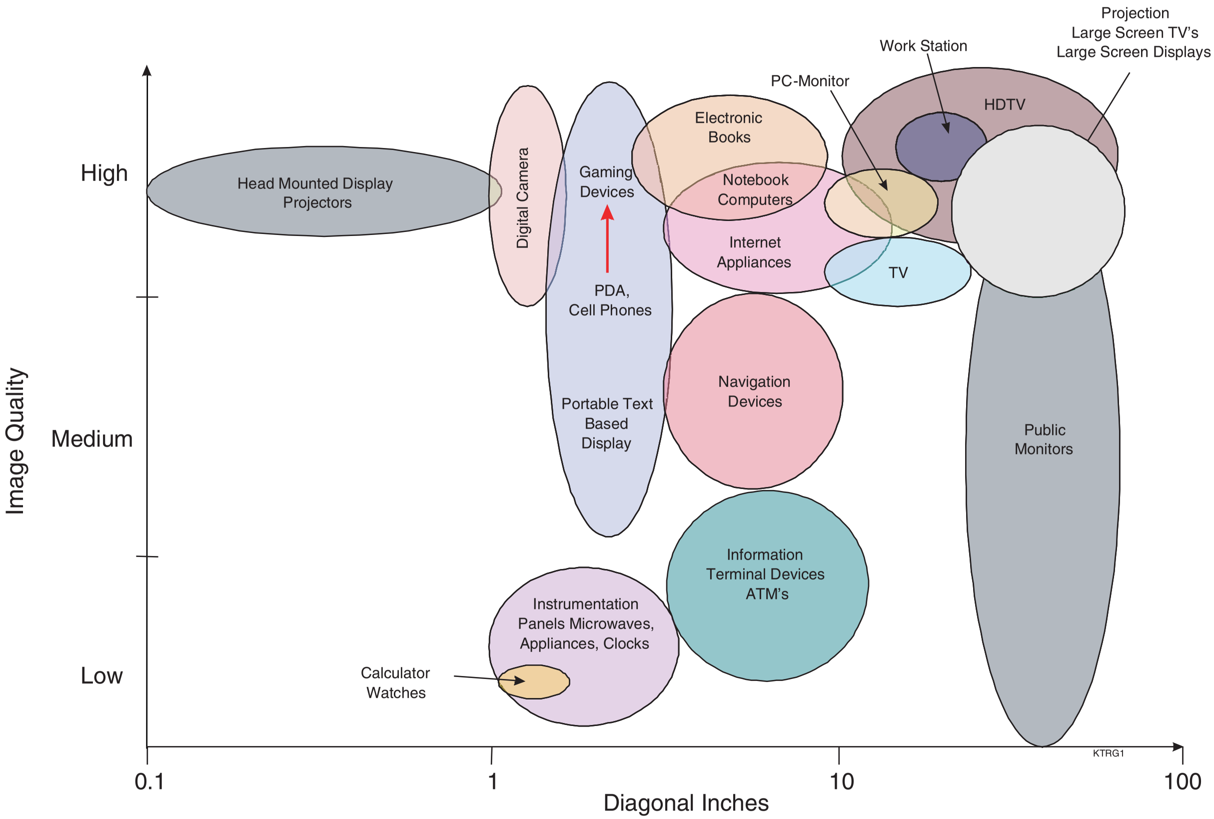

FIGURE 18 Display applications.

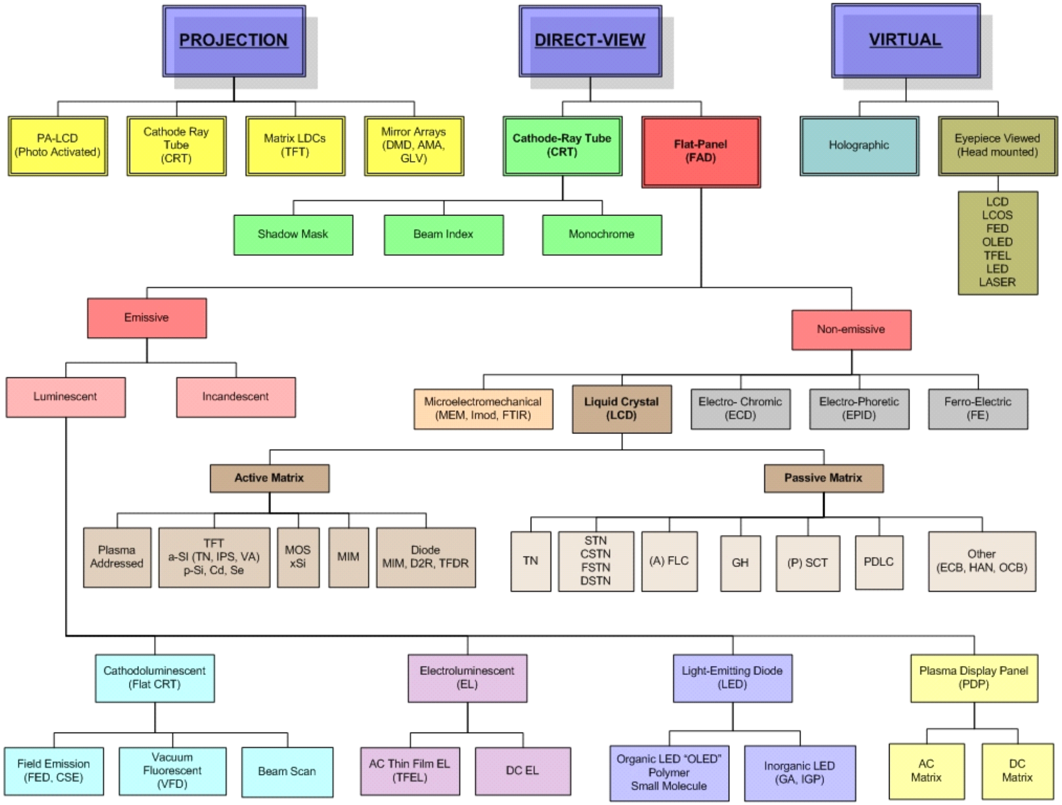

FIGURE 19 Display technologies.

the realm of the active-matrix LCDs, he pointed to a "tremendous explosion in the technologies available depending on the application," adding that their number was "growing yet." He singled out OLED as causing a great deal of excitement in the industry even though it was still in its infancy.

Turning to the current state of affairs, Dr. Keys declared his industry "unhappy with performance," "not where we want to be to meet consumer demand," and "especially not happy with cost or return." In the years 1995–2001, earnings from LCDs were –$2 billion despite capital investment of $34 billion and cumulative sales of $61 billion. "We really have to do something," he acknowledged. "Whether it's with the existing technology or new technology, we're seeking ways to turn this around because the opportunity is there. We as consumers want more." On the cost side, the industry was looking at more automated procedures that might move it away from lift-based processing and toward other ways of putting down materials in a one-pass operation that Dr. Keys called "more printing-like—Gutenberg-like, inkjet-like."

Also under consideration was moving away from glass, a change that might not be a low-cost alternative in itself but would be worthwhile if the manufacturer could become more automated. Mounting a display on a uniform or some other sort of clothing—"wearability" being the buzzword—might improve the display's durability. The possibility of making a flexible display was what brought DuPont into the business: The company saw revolutions in applications resulting from the manufacture of a plastic display using roll-to-roll techniques. While these novel applications had yet even to hit the drawing board—it would be "up to the designers to unleash their imagination," he said—the company believed this "dream" would be realized.

The roll-to-roll manufacturing concept—envisioned, in fact, as more in line with the sheet-fed printing press than as a pure roll-to-roll technique—centers on integrating components from the flex-circuit industry with ink-jet printing from the graphic arts industry to get rolls of materials that could be "sliced and diced" into displays. Expected to be very compatible with both roll-to-roll manufacturing and a flexible product was the OLED. As an emissive device, the OLED does not need the backlight that is always on in the LCD, which not only removes a level of structural complexity but also offers an operational improvement: Since in the OLED light is on only when needed, the LCD's light losses are eliminated.

In conclusion, Dr. Keys said that the display industry in the United States "must move away from LCD technology to have a future hope, not only from performance but also from applications." In part because of the industry's problems achieving profitability, and in part because no killer market application existed as yet for a flexible display, industry would not make the necessary moves alone and needed the government to participate, he said. "The manufacture paradigm related to cost, as well as the form factor, will come from the above functions," he explained. "So we would be betting a lot of money, from an industrial point of view, on a hope." Having seen what was coming, the U.S. Display Consortium had begun to redirect its funding, devoting 31 percent of its outlays to each of two leading areas, flexible substrate processing and "specific materials and components." Dr. Keys said he anticipated spending on flexible displays to continue at significant levels, although the consortium still believed in projection displays and certain aspects of the microdisplay and continued to fund some of them.

DISCUSSION

Jack Triplett, the day's first presenter, asked the three panelists to summarize the performance measures used in their industries to assess technical progress, as well as to indicate where such data might be obtained. Also, in reference to Dr. Bregman's contention that software had escaped the attention of those studying which metrics might reflect the relationship of IT investment to productivity, Dr. Triplett stated: "It isn't actually that we're unaware of the importance of software and the crucial part of the investment in software in looking at the whole computer. It's that we don't have a good metric for measuring its performance." If one were available, he said, far more material on software would be incorporated into studies of computer productivity.

Responding first, Dr. Bregman pointed to the existence of both internal and external metrics as a complicating factor. Taking up the former, he said that whereas progress in the microprocessor industry might be equated to more transistors per square centimeter, in software progress was no longer measured in lines of code. One of Veritas's most profitable products, from which it had earned in excess of $100 million, consisted of 14 lines of code. "Because core intellectual property is not found in a physical embodiment," he noted, "it is very hard to measure." Externally, software was measured by the practical impact of its use; crucial in the current environment, where money for new investment was scarce within IT organizations, was how much it allowed customers to reduce operating costs. "We often see returns on investments, or payback periods, that are within days to weeks," he stated, while nonetheless conceding that systematic metrics did not in fact exist. Moreover, he questioned how eager the industry was to see systematic metrics adopted. "If we all had systematic metrics," he posited, "then it would be clear: You'd publish who's the best according to some abstract measurement, and all the customers would buy that, and the game would be over." Still, he acknowledged that it was problematic to neglect the things one could not measure and said the industry should persist in trying to find a way to measure them.

In the graphics industry, said Mr. Malachowsky, there was a marketing view of performance in addition to internal and external views. At its founding NVIDIA wanted to show the value of its products by differentiating them on performance; since there was no benchmark, however, its first goal became to create one. Its benchmark not only turned out to be too narrow, measuring only speeds and feeds, but it also failed to measure end-user value. Currently, NVIDIA was measuring itself internally on "very engineering-specific, design-specific things" such as bandwidth utilization factors, and externally according to application benchmarks. "Synthetic benchmarks that came out of the CPU world just don't cut it," he said. "They're just not real enough, not end-user-experience enough." The application benchmarks used were to be found at hundreds of Web sites run by "teenagers, scientists, [and] frustrated computer geeks" who were attempting to survive financially by serving clients in need of performance measures for applications they thought valuable.

Dr. Keys began the recitation of an extensive list of metrics by naming technical specifications that allowed benchmarking from display type to display type: diagonal size of the display and pixel count. More complex was measuring power consumption, because how the display was being driven and how long it was in its "on" state had to be factored in. Switching speed, which indicated whether a product could do video-grade displays, was also a common metric. The measure of luminance, while technical, related to the user although it was complicated by the issue of consumer preference. Manufacturing metrics were based on yield and on the size of the substrate used in the manufacturing process. Indicating that there were many more metrics to name, he referred the audience to the Web site of the USDC for "references to competitive information and studies" incorporating such information.

Reflecting on the preceding remarks, Dr. Bregman observed that many external metrics or benchmarks used in industry were aimed at being able to compare products at a particular moment, making them unsuitable for the type of longitudinal study that economists seeking to establish the connection between product performance and productivity were engaged in. Industry had been known to change benchmarks from year to year. The result was a fundamental conflict between what industry was trying to deliver and what economists were trying to ascertain.

David Peyton of the National Association of Manufacturers asked Mr. Malachowsky about the relationship between offshore and onshore manufacturing for NVIDIA, a virtual manufacturer. He asked how much of the company's outsourced work was done in America and how much abroad; whether that ratio had changed in the previous several years; and, if so, what factors had driven the change.

Stressing that he was leaving his personal area of expertise, Mr. Malachowsky answered that all of NVIDIA's semiconductor manufacturing was done by TSMC, the Taiwan Semiconductor Manufacturing Co. While noting that TSMC had an interest in a joint-venture fab in the U.S. Pacific Northwest in addition to owning facilities in Taiwan, he guessed that TSMC did most of NVIDIA's work in Taiwan itself. The majority of NVIDIA's product, perhaps as much as 70 percent, was shipped to manufacturers outside the United States that then incorporated the components in products and reshipped them in turn. In addition to its manufacturing, a very high percentage of NVIDIA's testing and packaging were done overseas, whereas all the design, test development, and manufacturing process development were done in the United States. Mr. Malachowsky expressed doubt that the design function would move offshore, but he acknowledged that "more and more of the manufacturing does seem to be showing up overseas, in multiple countries where the products don't have to come back into the United States and be shipped back out." He pointed to the existence of a penalty, which he recalled as a tax disadvantage, attached to bringing goods in and out of this country that made doing so an economic burden; because latency of manufacturing, from component delivery through end user delivery, was "actually a big deal," shortening travel time was an important issue as well.

Recalling earlier allusions to internal technical metrics on the one hand and to external metrics such as "usability" and "customer likeability" on the other, Victor McCrary of the National Institute for Standards and Technology asked Mr. Malachowsky and Dr. Keys when color and fidelity could be expected to follow word processing to the level of WSYWYG ["what you see is what you get"].

Mr. Malachowsky said that graphics products were being held to a higher and higher standard in that they were being measured according to the electrical fidelity of the signals sent to the display, something he called a relatively new development. Elaborating, he said that some of the newer displays received a digital signal from the processor and "the display processe[d] that digital information internal to the display," with the fidelity quantized such that small perturbations originating in the processor did not show up on the display. In this sense, a subsystem of the display is responsible for the display's quality. On the analog side, he noted, in driving a CRT, the quality of the signal transmitted by the semiconductors through a board, a connector, and a cable into the electronics of the display dictated a large part of what showed up on the screen. Now algorithmic mapping had been added to offset the transfer function that the display would insert so that the user could get WYSWYG; the processor provided the mechanism to improve that, and applications software would make use of it. More often than not, the developers of the Web-based benchmarks analyses mentioned earlier had decided that the quality of that signal was an important metric for the consumer, and as a consequence NVIDIA and its competitors had begun devoting increased attention to it. But the innovation would have to take place in a narrow bandwidth envelope to avoid problems with the FCC, because the wires act as antennas, and the better the fidelity of the information, the bigger radiators they become.

Dr. McCrary indicated that his concern about image quality was focused in the realm of health care, and specifically on the fact that lack of fidelity could lead ultimately to liability.

This question brought to Mr. Malachowsky's mind an incident that took place when he was responsible for the first color accelerated display at Sun Microsystems. The display's introduction at Lawrence Livermore Laboratory was met with the comment: "Thank God, you finally got color to be as fast as grayscale!" Inquiring into the motive for this statement, Mr. Malachowsky was told: "Even though it's now fast enough to be usable, I can't use it because you don't have enough shades"; the Livermore researcher, it turned out, was in medical imaging and was concerned that he might misread a tumor because of a density interpretation. As a technology advances, it does not necessarily cover all the applications of the prior technology, he commented. Noting that NVIDIA's chips had gone from 8-bit color to 128-bit color with 32-bit floating-point components, he speculated that "maybe now that guy will say, 'Thank God, you've done that, because now I can use a color display at all.'" A large number of transistors are "tied up" with 32-bit color components, he noted.

Dr. Keys, remarking that he had spent his first 14 years at DuPont in its x-ray film and graphic arts film business, declared that the display industry was "still at the Gutenberg stage." He denied, however, that another 200 years would have to pass before that quality was reached. LCDs had reached a pinnacle in color fidelity, he noted, but acknowledged that fidelity was highly dependent on the color filter and that having to add additional layers causing color shifts and angle dependencies that are a detriment was a disadvantage. To do what Dr. McCrary was referring to would, Dr. Keys guessed, require 5–10 years to get a material set with the color fidelity to enable progress beyond consumer applications and into the medical realm; and, he added, contrast capabilities reside not just in the materials set but also in how the display is fabricated and driven.

Kenneth Walker, another of the day's presenters, pointed out that the system described was an open-loop system and that the discussion involved the comparison between transmitted and reflected color. "Although you might lay the same densities of color from any set of standard measures," he cautioned, "you're going to get different effects in terms of how the human eye perceives that color."

NIST worked on standards for color and had representatives come from the printing industry asking for the standard for red. Dr. McCrary could, he said, "give them a laser line and say, 'O.K., that's about red,' but that standard for color has not been further articulated since the color diagram over 60 years ago."

Agreeing, Mr. Walker noted that a closed loop would have to be built in order to provide a system for digital cameras and similar products. It would be necessary to have a monitor such that when a signal was sent out of a graphics card that was intended to be Red 23, it could be measured to determine whether or not the phosphors in the monitor had aged to the point that the color had shifted, so that the card's manufacturer could learn how to change the signal in the card to get it back up to Red 23.

Corroborating this statement, Dr. Keys exclaimed: "If you sit in the back of a jet and look down the stream of LCDs, how many of them match? None!" He asked in addition: "When I print a document in L.A., does it match the document I'm printing in Chicago?" Remarking that the problem had not yet been brought to the display industry, he argued that there was a tremendous opportunity to build on what companies in the printing business had learned in controlling and matching color.

Dale Jorgenson, the STEP Board's chairman and host of the conference, asked whether road maps existed in the speakers' segments of the software, semiconductor, and display industries and, if so, where they might be found.

Dr. Bregman answered that there was neither an equivalent to the SIA road map in Veritas's part of the storage software industry, nor was there the same sense of cooperative effort. "If we had a road map, we wouldn't tell you," he said, "unless you were a prospective customer, and then we might show you a little bit of it." Because barriers to entry were much lower in the software than in the hardware sector of the industry, revealing a software road map would be tantamount to handing over intellectual property—so that the road map was not an appropriate solution there. Nor were internal road maps adhered to very well, since a brilliant idea can make its way from an engineer's mind to the market "within a quarter or two," most of which would be taken up with testing rather than with product development. As a result, road maps in the sector generally focused on customer requirements rather than being technology driven; they were intended to guide engineering teams as to what management thought was important from a customer point of view.

Dr. Flamm, noting that NVIDIA had very few competitors, asked what use the firm might make of a road map.

Describing the market for graphics processors as half in the hands of Intel—with NVIDIA and ATI, a Canadian firm, owning the other half—Mr. Malachowsky characterized a typical road map as "soft and fuzzy, and quantized into non-specifics: 'photorealistic,' 'cinematic quality,' 'speeds and feeds at double or triple algorithmically.'" To the extent that his own company's road maps were accurate, it was only because "we're progressing like that and we'll call it that when we get there regardless." On an industry-wide level, there might be projections for a few years out or for a generation or two out for each product segment, but there was nothing analogous to the semiconductor industry's road map.

For displays, Dr. Keys recommended Stanford Resources' work on the display industry, as well as trends reports by Display Search, based in Austin. Kanzai Research Reports of Japan also rated a mention. The USDC had recently released a report on flexible displays.

Mick Silver of Cardiff Business School harked back to Dr. Triplett's initial point: that to measure productivity, price change needs to be measured but also to be corrected for change in performance. This could be particularly difficult in software, he argued, where technology undergoes "sea changes." Although it was relatively simple once a software product has been on the market for successive periods—assuming the quality to have stayed constant, one could just measure the price—the trick was to measure the price change between the period before introduction and the period of introduction, as well as to measure its value to users compared with the value of the technology used before it was introduced. "This is where the productivity change comes in," he emphasized, "between the old technology and the new." Econometricians attempted to estimate such hypothetical prices—what the price might have been before introduction—but it was not easy to do. Locating the heart of the productivity question in the initial period when leaps are made in software, Dr. Silver suggested that the way forward might come through employing market-research techniques to work out how users value the productivity change and through asking what the price might have been had a product been available but not sold.

Concurring, Dr. Bregman chose as an example the introduction of software to simulate crash testing in the auto industry. "What's the productivity value of that software?" he asked. "If we have to look inside the whole system, how did they satisfy this task in the past—including the cost of the cars they destroyed—vs. how do they do it today?" Similarly, the value of software within the information technology sector itself could not be measured without looking at the overall task. If it was managing data, then how was it done before the software became available? By employing hundreds of people. "Now," he asked, "what's the relative economic value? But we don't tend within the industry to look at it that way except to the extent that we build an ROI model for that transaction to sell it. We don't look at it over time."

In the graphics industry, declared Mr. Malachowsky, "if you were looking at price, you'd be deceiving yourself or leading yourself astray." In an industry that started off with gross margins in the 5–6 percent range, NVIDIA was now posting an average gross margin of around 35 percent. "The cost is something you might track," he admitted, "but we're in a business, we're a public company, we're trying to make more money, more margin. And prices don't necessarily reflect costs. We're trying hard to have them not be a linear extrapolation."

In software, where the incremental cost is zero, Dr. Bregman pointed out, that was an even larger issue.

Mr. Malachowsky added that NVIDIA's virtual manufacturing business model, under which it owned the manufacturing line but not the infrastructure, was designed for extracting more value out of the product chain. The firm takes more margin from the manufacturers than they would have had before, but the total chain still spends so many dollars. "We just try to keep all of it," he said. "We want our users to pay more over time, even for the same thing."

Allowing himself the final question, Dr. Flamm called attention to the presence of two representatives of the visual arts on the panel and raised a point he looked upon with some amazement: that in past discussions of technological advances in graphics and displays, the high-end graphics market had been the scientific and design market. Noting that NVIDIA was active in what was basically a consumer market, he asked whether the high-end, performance-driving piece of the market had shifted from scientific and engineering users to consumers and whether that was also true for displays.

Both Mr. Malachowsky and Dr. Keys answered in the affirmative.

"Has that had any impact on the ways you're driving your product forward: the aspects of resolution or what kinds of shapes you paint on the screen?" Dr. Flamm asked.

Again answering in the affirmative, Dr. Keys said that upon entering the display market only five years before, DuPont had laid out its own road map. But that road map, which included what he called "a substantial period of monochrome display supply," had been abandoned. In just the previous two years, there had been revolutions in the industry, in the move to full-color displays in hand-held devices, and in the advent of HDTV and the infrastructure to support it. "That drives all of us, as consumers, to want that latest thing," Dr. Keys observed, "so patience with monochrome just doesn't exist." Even for day-to-day, utility functions, technology had moved toward the high end. DuPont was now testing technologies on such lower-end applications as cell phones and games before implementing them in higher-end computing products. "But," he noted, "everyone wants the same capability."

The economies of scale of the consumer world "have removed the economic motivation from the very high end" and caused "the real high-end players to focus on narrow niches now," added Mr. Malachowsky.[18] The overall advancement of technology, in which the entire industry moved forward to satisfy needs arising in a variety of applications—some of them very high end—had given way to concentration on small market segments. High-end developers had moved into simulation and modeling. NVIDIA had entered the workstation market two years before and, at the medium and low end, had established dominant market share—producing around 20 million processors per quarter—and had put in a minimum feature set in order to be in a position to supply extremely fast machines for very focused, high-end manufacturing customers.

Dr. Flamm observed that the technology going into the high end and that being developed for the consumer market were the same, to which Mr. Malachowsky assented.

---

[18] The term, economy of scale, describes situations where the average cost of producing an item falls as the number produced increased. In other words, the more of an item a company makes, the less each item costs to make on average.

Panel IV

Peripherals: Current Technology Trends, *continued*

INTRODUCTION

Michael Borrus
The Petkevich Group

Characterizing the meeting as "phenomenally interesting," Mr. Borrus said the day's presentations had made clear that the difficulty of measuring productivity growth in computing arose at least in part from the fact that value and functionality were constantly shifting. NVIDIA, for example, was trying to capture more of the value added by migrating software functionality to its chips; component manufacturers were trying to migrate system functionality into the component by adding processing capability to displays, to magnetic-storage components or to networking components; and Veritas was trying to migrate network- or system-management functionality down to the software. An inherently difficult technical measurement problem was thus exacerbated because "who's capturing the value keeps moving around based on changes in business strategy and the ability to execute."

The current panel would offer two more interesting and diverse examples. The first of these, optical-storage technology, had not displayed a pace of technical advance equal to that of magnetic or solid-state storage, with the result that its applications were defined and, in some sense, limited. A question to keep in mind during this presentation would be where new applications and market growth might come from. In contrast, the technology of laser and ink-jet presentation seemed almost infinite in its potential applications—particularly on the ink-jet side, where the printing of words on paper had led, somewhat surprisingly, to the spraying of genetic material onto gene chips and to the manufacture of flexible plastic circuitry and organic semiconductors.

The first speaker, Ken Walker, was the veteran of start-ups in Silicon Valley and elsewhere and had until shortly before held the post of Vice President for Technology Strategy at Philips Electronics.

CD/DVD: READERS AND WRITERS

Kenneth E. Walker

Despite being "one of the casualties" of Silicon Valley's recent downturn, Mr. Walker said he was "still bullish on the future." He proposed a quick review of the state of the art to begin his talk on developments in optical storage, which he described as an established business that was not so much technology-driven as operationally driven. Pursuing the theme of the migration of technology, he noted that value was moving away from the creation and sale of drives and into the integrated circuits necessary to create drives and read the data, as well as into the optical pickup unit: that combination of a solid-state laser and plastic lensing that reads the optical device. While DVD and CD readers had become standard on PCs, certain limits in those devices' capabilities were starting to be reached. Top-of-the-line CD devices, available from mass-market appliance and CD vendors, were rated at 48X to 52X (or 48 to 52 times the speed of the original audio compact disc), the equivalent of around 200 km/hour—a speed approaching the reigning physical limit for CDs, since operating at higher speeds would cause the disc to shred within the drive.

But if physics had placed a wall before the industry, a new set of capabilities had come along in extensions to the rewriteable CD-RW referred to as "Mt. Ranier" or CD-MRW. With Windows XP and a CD-MRW drive, Mr. Walker explained, it was no longer necessary to erase everything on the disc in order to add something to it; instead, material could be dragged on and off. Predicting that the optical device's future would be as "the next floppy," he remarked that after a long tenure the floppy was dead—PCs were being shipped without them—and the CD was taking its place. At the same time, however, a battle was taking shape between read/write and rewrite. Drives labeled DVD–RW and drives labeled DVD+RW were being made according to very different standards of rewriteability. In DVD–RW, a rerecordable format, a disc could be used a thousand times, but adding anything required erasing it and rerecording. In DVD+RW, which resembled CD-MRW, additions could be made incrementally and sequential erasures were possible, with whole segments able to be erased and reused.

Finally, the next generation of DVD, based on a blue rather than a red laser, had been shown over the previous two years in a consumer electronics perspective.

To understand the conditions prevailing in the electronics industry in a way that would further a discussion of the deconstruction of the computer, Mr. Walker claimed, it was imperative to consider the "fight back and forth" that had taken place over the years between consumer electronics and compute electronics. Many features of computer monitors, from size to resolution to aspect ratio, had been shaped by the early use of television screens as monitors; also playing a role was the fact that computer firms were obliged to deal with the same manufacturers that produced TV tubes, since glass handling was the part of the process that required the greatest capital expenditure.

The impact of consumer electronics on optical storage could be seen in laser discs. In the 1970s, when the volume of computers was basically nil and the volume of televisions quite high, a great deal of work was done toward creating a video laser player for a long-playing disc designed to rival the LP. Digital technology began to come into consideration around the middle of the decade; the original video discs, even though they were laser based, used analog technology, as did the original audio laser player, or ALP. Philips spent two years looking at an analog long-playing disc; thinking quadraphonics was going to be big, the company was seeking a large-size disc with a capacity of two hours. When it saw, however, that quadraphonics wasn't going to work, Philips opted for stereo, which could be accommodated on a compact disc. The first public demonstration of an 11.5-cm disc took place in 1979, two years after the idea was adopted. Considering that the development covered most of the period 1970–1979, Mr. Walker stated, change from an R&D perspective was not that fast.

In 1980 Philips went to Japan and approached all the Japanese electronics companies. Sony, the only one that thought there was any real promise in the CD, joined with Philips and, together, they created a standard called "Red Book." When the first CD products came into the market in 1983, the two firms talked it up and managed to convince all the other firms to jump in and convert their catalogues. And because the standard belonged to Philips and Sony—the size of the disc was copyrighted, the name "CD" was copyrighted—everyone in Japan had to come to them to license the intellectual property, and they made a lot of money. The standard was set for CD audio in 1980 and product reached the market in 1983, the same year that the first data version, the CD-ROM, appeared. The recordable CD for computers, the CD-R, followed in 1993; the CD-RW in 1997; and the Mt. Rainier, which allowed the CD to be treated more like a floppy disk, in 2002.

The other Japanese companies' conclusion from their experience with the CD—that they should never again allow themselves to be held hostage over standards—led to a certain fractiousness at the appearance of the DVD. To avoid being left out in the cold, the firms joined with one another in a consortium, the DVD Forum, to which each contributed intellectual property; license fees from non members were to be divided up by members in proportion to their contributions. So that "everybody got to play," Mr. Walker recalled, a wide variety of standards, "none of which actually worked the same," was created: There was a separate standard for video, for data, for rewritable data, for audio, and for recordable and rewritable data. Each was promoted by different companies and different groups, hugely increasing the problems of compatibility. That the different devices were variations of what was all fundamentally digital technology reflects an American view that "did not fit the Japanese view of how consumers use products," he said. The Japanese firms resisted pressure from U.S. PC makers to treat everything uniformly as data until the end of the 1990s, when they realized that DVDs for PCs had outsold consumer DVDs by a 5 to 1 margin. The DVD Forum, having at last seen the light, in November 2002 released a draft standard for what was being called the "DVD multiformat," which corresponded to a DVD multidrive that was to read and write all DVD standards.

The rapid pace of change in CD and DVD technology "gets in the way of our being able to do a nice, clear curve of price-performance improvements," Mr. Walker stated. The decline in price of an audio CD player between introduction in 1983 and 2003 had been dramatic; the rapid drop in price that followed the advent of the CD-ROM in the late 1980s had occurred simultaneously with an increase in capability. The 12X CD had been in the market less than six months when it was displaced by the 16X, a phenomenon that, considering the price of R&D and of tooling up, made it a challenge for producers to break even. As the technologies advanced, with some superseding others, speed became harder to compare: The CD-R had a dual specification, speed to write vs. speed to read, while the CD-RW's, read/write/read, was tripartite. "CD-Rs can read CD-ROMs, CD-RWs can do CD-Rs and CDs," he remarked. "That's why it's difficult to create a single bar or graph that says, 'This is how one technology has performed versus price and another.'"

Turning to future computing standards for optical storage, Mr. Walker predicted that the CD-RW would replace the floppy at the low end; that rewrite speeds would reach an upper limit of 48X–52X with standard media, although special media might achieve more; and that the DVD+RW—which was supported not only by its creator, Philips, but by Sony, HP-Compaq, Dell, Microsoft, and IBM—would prove the winner over Pioneer's DVD-RW, even though the DVD Forum was pushing the latter. Sony, meanwhile, was following its usual practice of hedging its bets with a combination device compatible with both CD+RW and CD-RW. The rewrite speed limit for the DVD matched that for the CD, around 200 km per second, but it was designated as only around 16X because the much higher density of DVD data yielded a higher data rate at that same speed.

Another source of technological development was the quest for copy protection on the part of the recording and motion-picture industries. One promising option was the embedding of chips with copy-protection identification in a CD so that it could not be rerecorded by home users. "If you buy a copy of Microsoft

Word, the error-checking card will have to be read by the driver," Mr. Walker explained, adding: "It will be a mess, but the paranoia that's coming out of Hollywood is going to drive a lot of things around embedded copy protection." The big technological shift, however, would be the advent of the blue laser for DVD, whose wavelength would allow a more tightly focused beam and thus offer superior density; it would operate with a 0.5-micron gap, compared to a 1.6-micron gap for a CD and a 0.72-micron gap for a DVD. This would make possible a huge increase in storage capacity from the 4.7 gigabytes that then represented the standard side of a single-sided, single-layered DVD. Although up to 22 gigabytes could already be stored on a double-layered, double-sided DVD, the appropriate comparison would be to the 25–27 gigabytes possible using a blue-laser-based platter of the same size. Again the consumer side was driving the computer side: in this case, demand for DVD products compatible with high-definition television. A fight was also brewing over this technology, however, as Microsoft and Warner Brothers were arguing for better compression using existing DVD media and red-laser devices. According to Mr. Walker, Microsoft wanted to own the codec because, if it was Windows-Media based, Microsoft would get a royalty on every DVD player sold. Warner's interest was in reselling its catalogue: Blue-laser technology was being set up to be inherently recordable as well as readable, but if the winning technology were not recordable, the company would be in a position to resell its entire catalogue to anyone acquiring an HD set.

The blue-laser disc could be expected to operate on a single rewrite standard from the start, a consequence of the lessons learned from the DVD experience. But variation might come in the size of the disc: Companies had been looking at "small form-factor optical," a 3-cm disc using a blue laser in a drive about the size of a matchbook that could accommodate a removable piece of optical storage holding a gigabyte. Makers of advanced phones and digital cameras had shown an interest in the technology, possibly as an alternative to IBM's microdrive.

Beyond the blue laser not much technological development was in store for optical storage's next decade, in Mr. Walker's personal view. Recalling that it had taken 20 years to make "simple" red lasers function at their current level, he pointed out that only non-visible light in the form of ultraviolets would offer any significant advantage over the blue laser. The relatively low price that users were willing to pay for optical-storage devices and the difficulty of mass production would combine to limit their commercial potential. In addition, there were problems inherent in the technology itself: At some wavelengths ultraviolet lasers are no longer reflectable in standard mirrors but burn right through them, complicating the construction of a device.

Returning to Mr. Borrus's point concerning the shifting of functionality, Mr. Walker ventured that competing technologies were likely to take on some of the roles thought to belong to optical storage. He called Mr. Whitmore's presentation on magnetic storage "compelling" and noted that he himself used a 1394-based hard drive to back his home computer systems rather than optical drives—even though he owned DVD burners—because it was "just easier and simpler than feeding disks into a system." He called solid-state, static RAM "great" for use with the USB dongels that had become popular for transferring files and rated as "fascinating" the new magnetic RAM technology: "Like standard RAM, it doesn't need any energy to maintain its state, but it's also extremely fast. So you can turn your computer off, walk away, leave it away from power for a week, come back, plug it in, and everything is right where it was before—there's no wait for the state to be written out of the hard drive" as with some laptops. Pervasive networking, which like optical storage could be used to share files and to move them back and forth, was another source of competition. "If the network is fast enough, and I can get to my music from my hard disk at home through my cell phone, do I need to carry an MP-3 player with me?" he asked.

Looking to the future, he stressed that a fight was in the offing over whether red or blue lasers would be used for high definition. DVD and CD reader/writers would be pervasive in mass-produced distribution, as the cost of making mass quantities of discs was in the tenths of a cent. CDs also held value for those who didn't want to be concerned about electromagnetic pulses taking out their content and those who created content on their computers that they wanted to play on a home stereo, although this advantage might be attenuated with the advance of pervasive networking, whereby everything in someone's consumer electronics deck would be connected to his or her computer. Continued Asian control of the optical storage business was to be expected. There would be limited R&D beyond the blue-laser technology, with royalty revenue models decreasing and the price steadily eroding. The feature set would become more consolidated: "We're going to end up with a single drive that can read and write CDs and read and write DVDs," Mr. Walker declared. He saw the questions determining which media consumers use and why as "How big a piece of optical do you want to hand around to somebody?" and "What quality of video are you going to want to play on the TV set?"

Prominent among the lessons learned from the experience of the previous two decades or so was that a shift had taken place in the value equation. At the industry's inception, it was the drive manufacturer who had the power and collected the profit. But building a drive was no longer challenging because the availability of subcomponents was such that it had become quite easy to assemble them into a drive. "If you are a drive manufacturer, your focus must be on execution and assembly," Mr. Walker observed. "And if your focus is on execution and assembly, and not on worrying about the integration, these people start forward-costing." The cost of a new generation of CD or DVD then works out no longer to have the premium for a new speed that was earlier associated with a new technology. As a result, he said, integrated companies pursuing technology innovation were no longer able to recoup their investment in R&D. Value was now coming from the intellectual property: that is, from the chips. Philips, he noted, had helped itself by having its semiconductor group make chips; but the semiconductor group, in line with its obligation to make money, started selling chips to Philips's competitors—"and, well, there goes that business."

In conclusion, Mr. Walker noted the main reasons behind the fact that the optical-storage business had what he called "a very Asian center of gravity." When the CD came out, he recounted, the manufacturer and designer mechanisms were jumped on by Japanese companies, while "almost no one" in the United States and only Thomson and Philips in Europe responded. The commitment made by Japanese manufacturers meant that the component makers all set up shop in Japan, with the result that their R&D was located there as well. By the time the PC started to use optical-storage devices, a large number of PCs were being built and designed in Taiwan, where the government encouraged PC manufacturers in their desire to move into the sector. Taiwan's Industrial Technology Research Institute (ITRI) spent liberally on technology-development programs, enabling the formation of companies like Lite-On and Media Tek; this made for a synergy of location, he noted, as the two firms "happened to be right across the street from each other." Lite-On subsequently became one of the leading providers of optical drives and devices, Media Tek the leading chip provider in Taiwan. Drive manufacture had been moving to China; with the customers, the PC and consumer electronics companies, having moved or moving their manufacture of PCs and DVD players to China, the component providers were moving there as well.

Mr. Borrus then introduced as the next speaker Howard Taub, Vice President and Director of the Printing and Imaging Research Center at HP Labs. Dr. Taub, a member of the core group that had managed the invention of thermal ink-jet technology at Hewlett-Packard, was the holder of a great number of patents in that and related areas.

LASER AND INK-JET PRINTERS

Howard Taub
Hewlett-Packard Labs

In prefatory remarks, Dr. Taub said he would go beyond the title of his presentation to speak about a third printing technology and about digital publishing, fields he held to be of great interest at that moment. And, embracing Dr. McQueeney's earlier statement, he asserted that it was no longer possible to talk exclusively about technology, but that technology, infrastructure, and business issues had to be considered together.

Showing a chart that indicated the actual and projected breakdown of market share among various printer technologies from 1998 through 2006, Dr. Taub noted that only a thin line represented color page printers and opined that, despite appearances, the workplace had not made the transition to color laser printers (see Figure 20). Market share for laser printers, designated as "monochrome page

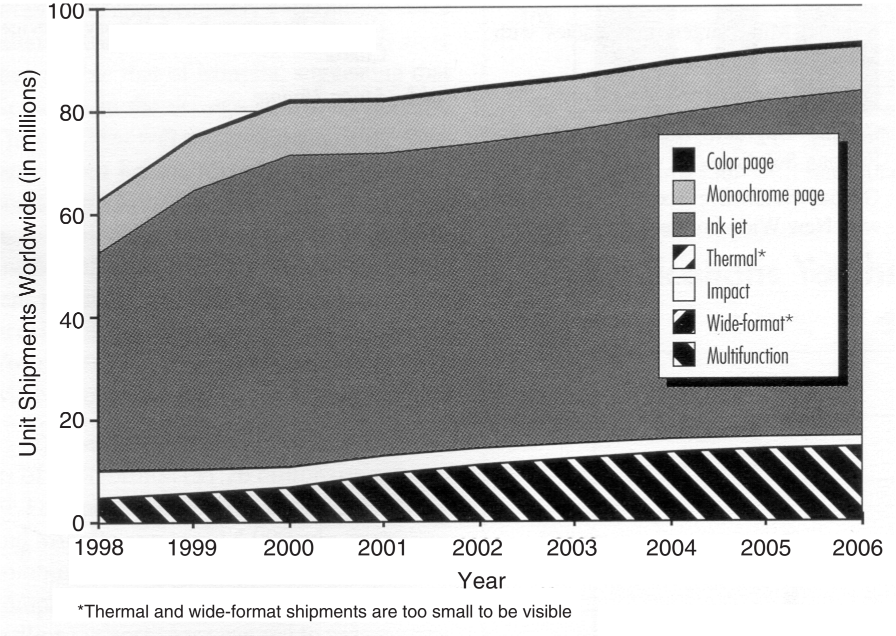

FIGURE 20 Why focus on ink jet and laser? Shipment projections of various printer technologies.
SOURCE: Lyra Research, Inc.

printers," was fairly substantial, but it was obvious that the ink jet had held or was expected to hold a consistent two-thirds of the total printer market over the entire period of the chart. In addition, ink-jet technology was used most often in multifunction printer peripherals, which as a category had passed monochrome laser printers to move into second place sometime in 2001. Sales of thermal printers and wide-format plotters were not large enough even to show up on the chart, while those of impact printers were visible but dwindling into insignificance.

Among a score of ink-jet technologies on the market, there were two clear leaders, piezoelectric and thermal; piezo, around since the middle of the twentieth century, was the less popular of the two technologies. "Laser printer" was a misleading term because some of these machines wrote using LED arrays rather than lasers. While this class of printer was best characterized as "dry-toner electrophotography"—with electrophotography, or EP, designating the kind of printing—Dr. Taub himself thought of the laser printer as being defined by the fact it used a toner particle of around 5 microns or more in size. A significantly different technology although it was also based on electrophotography, "liquid-toner EP" used toner containing particles that were much smaller than the particles usually used in a laser printer—of submicron size instead of 5 microns—which afforded capabilities in terms of quality and speed that were hard to achieve with a dry-toner printer.

He displayed a chart that located ink-jet, liquid EP, and laser printing along axes corresponding to the parameters "faster," "better," and "cheaper," which Dr. McQueeney had stated to be de rigeur in any forward-looking assessment of technology (see Figure 21). Recalling his first work in ink-jet printing, on IBM's 6640 in the 1970s, Dr. Taub said that only $5,000 to $10,000 of the value of that 92-character-per-second, $30,000 machine was accounted for by the printer itself, with the rest residing in its massive paper-handling mechanisms. In the 1980s, HP had come out with its first ink-jet printer, the ThinkJet, which had a price of $500 but was a special paper printer and was not letter quality; Dr. Taub had managed the research project at HP Labs that invented the technology. It was not until around 1987 that "the real bigger winner," the DeskJet, made its appearance. A black-and-white, plain-paper machine printing 300 dots per inch, it cost $1,000; shortly afterward, the company was able to offer a color version at around the same price. In 2003, $49 would buy a printer similar technologically but of much better quality, able to print out photos that were "almost as good as good-quality Kodak photographs," he said.

Although he suggested that a corollary to Moore's Law could be imagined for printers, given the foregoing evidence of quality improvement accompanied by price decrease, Dr. Taub declined to predict whether they would indeed get

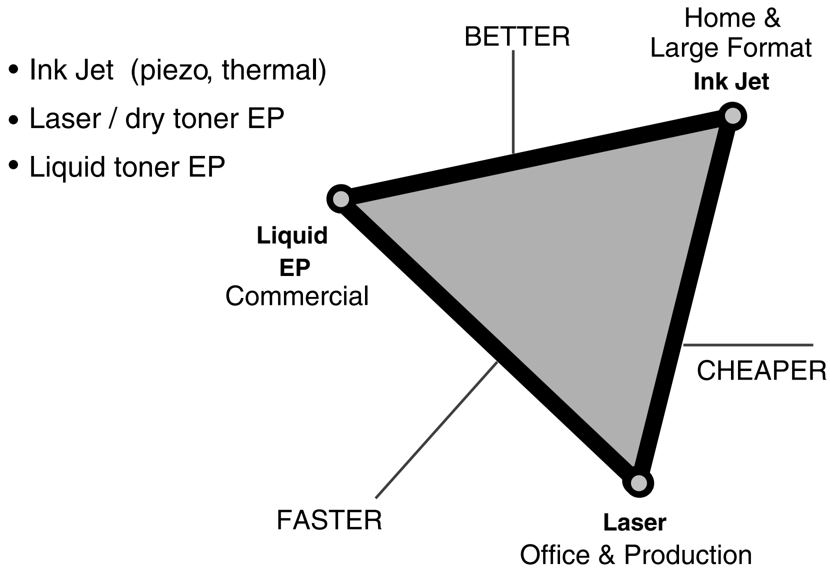

FIGURE 21 Printing technologies.

cheaper. The printer was sold on a "supplies" business model like that of the razor-blade business. "We call [printers] 'sockets,'" he said, and "just want to get as many out there as we possibly can and then sell supplies for them." Some were already given away: There were Lexmark units available for $29, with the cartridge costing more than the printer. As to the "faster" parameter, an ink-jet printer could be made to go faster simply by adding more nozzles. The meaning of "better" tended to center on resolution, although only some of that improvement had been of real significance, the rest being the result of a competitive "game." After moving early on from 92 to 300 dots per inch (dpi), HP put out a 600-dpi printer, then increased that to 1,200 dpi. HP stated at that point that higher resolution would not constitute true improvement, but its market share began to drop as competitors went to 2,400 dpi, which it then matched and later bettered at 4,800 dpi.

But pushing resolution any further would make "no sense at all," he said, because what is "really critical about the quality of the print is the size of the spot that you're putting on the paper—and if you can't make that smaller, then going to higher and higher resolutions doesn't mean anything. It's more marketing than anything else." The spot size had in fact gone down considerably over the years, from 240 picoliters for the ThinkJet through about 100 picoliters for the DeskJet and finally to about 3 or 4 picoliters, a size approaching the point at which the eye can no longer see an individual spot on the page. And photographic printers had been using extra inks to dilute the cyan and magenta such that the "final little spot" had become invisible. "We are pretty much at a point where the quality of the image that you can print is about as good as you're going to get," Dr. Taub stated. "There's probably a little bit more you can squeeze out of it, but not very much more." The notion of "better," therefore, had gone on to acquire other dimensions, encompassing more features than resolution: connectivity, in the shape of the ability to plug a card into a photographic printer and print pictures out; ease of use; and industrial design.

Comparing the attributes of the main printer technologies, Dr. Taub called the ink jet "cheaper and better," capable of providing high-quality images and text very inexpensively. Faster and more expensive ink jets did exist—for example, Scitex made a $3.5 million ink jet for print books and high-speed graphics—but the bulk of ink-jet devices were not at that level of performance. He characterized the laser printer as "faster and cheaper," while the liquid-toner EP, used more for commercial printing, got the designation "better and faster." Each thus had its own niche: Ink jet, delivering high-quality and color at low cost but somewhat slowly, fit well into the home market as well as being the technology of choice for large-format plotters; the laser, leader in the office, was in production printing as well; and liquid EP was the right device for offset-quality commercial printing.

Describing the printers' technologies, Dr. Taub said the piezo ink-jet printer had a chamber full of ink with a small piezo crystal on it; a voltage pulse could be applied to the crystal at the rate of 10,000–20,000 times per second, causing a deformation that in turn squeezed the chamber and drove an ink drop out of it. The thermal ink jet worked on a similar principle, but instead of a crystal on the outside of the chamber it had a resistor inside. Applying a pulse of current to the resistor drove the temperature of a thin layer of ink above the resistor up to 400 degrees Celsius, which in a matter of microseconds created a bubble over the resistor that pushed on the ink in the chamber and forced a droplet out. In going up to 400 degrees in a matter of a microsecond, the heat flux from the resistor was comparable to that at the surface of the sun; for that one microsecond, a one-square-meter heating element could actually consume the energy produced by a medium-sized nuclear power plant.

Dr. Taub began his explanation of the laser or dry-toner EP printer by describing a drum that was continuously being cleaned as it rotated and then had a uniform-layer charge put down onto it by a charge roller. A laser scanned back and forth across the drum, turning on and off as it did and discharging areas of the drum that it passed over when it was on. In this way, it added the information to be printed; if the charges were visible, they would reveal a pattern made up of charges that resembled the text and images to be printed. At a toning station, toner was thrown at the charge patterns, sticking where the drum was charged and not sticking where it was not charged. The image pattern formed by the toner on the drum was then transferred to the paper, which passed through a fuser roller that, employing heat, fused the image very robustly onto the paper.

The liquid-toner EP printer works basically like a laser printer, using a laser to create a charge pattern on a photoconductive drum. This pattern is then “developed” using the liquid toners to create a toner image on the photoconductor. However, while the dry toner laser printer transferred its toner directly to the paper, the liquid-toner printer transferred the toner to a soft-rubber roller which in turn transferred the toner to the paper. This process, which uses an intermediate transfer to a rubber roller, is similar to the process in offset printing where the ink is transferred to a rubber “blanket” before the final transfer to the paper. As a consequence, the liquid-toner EP printer shared with the offset press the capability of printing on very irregular surfaces.

Moving on to the technologies’ attributes and issues, Dr. Taub characterized ink jet as a simple and highly scalable process. He posited that a one-nozzle ink-jet printer could produce a very inexpensive, very low-performance printing solution; at the high end, $3 million ink-jet printers were capable of printing newspapers. Over this enormous range, printing nozzles could be added to increase speed and color capability as desired. Ink-jet printers costing $100, $1,000, and $20,000 shared the same fundamental technology. One of the principal challenges facing the technology was improving image durability, with respect to both light-fastness and water-fastness. The former, after a number of years of work, had largely been solved: Durability under normal use, represented by a framed document on an office wall, was being quoted at around 17 years, rising to 70 years or more if the document were placed in an album or otherwise kept out of the light. Waterfastness continued to need work, as in some cases inks would run if they got wet. Another challenge, particularly when a large number of nozzles was involved, was dealing with what happened if a nozzle went out. Redundancy and maintenance systems offered solutions to this problem, although the increased complexity that accompanied the move to larger systems could "make things a little Rube-Goldbergish," Dr. Taub acknowledged. Not only reliability but also drying speed became an issue with more complex systems because the amount of fluid on the paper that needed to dry quickly increased.

Laser printers had achieved excellent text quality, and they were starting to provide good image quality as well. High-speed printing; durable print, both water-fast and light-fast; and relatively low operator intervention were also positive attributes. Among the challenges outstanding were that the gloss of the printed material was shiny in contrast to that of the paper, whose gloss is dull. Because in offset printing the printed areas and the nonprinted areas had the same gloss to them, no one in the trade would have mistaken laser-printed copy for offset copy. In fact, with most dry-powder toner printers the layer of toner tended to be 5–15 microns thick as compared to 1 micron for offset printers; this not only drove up the operating cost, it left a very glossy layer no matter what the paper looked like.

Liquid-toner systems afforded the highest text and image quality, and the gloss of the paper matched the gloss of the print because the layers were at 1 micron, as in offset printing. Because the toners were trapped in the liquid, processing speed could be cranked up beyond the capabilities of dry toner. Dr. Taub showed sample images from HP dry-toner and liquid-toner printers for the sake of comparison; he pointed out that the dry-powder toner produced fuzzier print because of its larger toner particles. Comparing color spots produced by a liquid-toner printer against color spots from an offset press—such spots being a basic tool used for judging print quality—these produced very similar results. As to problems and challenges, the liquid-toner systems were easier to use than a printing press but more complicated than something like a Xerox DocuTech, and they required more highly trained operators than the latter. Also a question was whether this technology was suitable for the office environment, since it involved solvents and, therefore, containment systems for the solvents. And as it incorporated a cooling system, it was not only power-hungry but could not simply be plugged into a standard outlet, requiring instead a special power source.

But besides having the capability of providing quality indistinguishable from that of an offset press, this technology was digital—something, Dr. Taub observed, that "really changes things significantly." Looking back a few years, to the late 1990s, at the distribution of work load in the printing market, he cited figures on the order of several hundred billion pages per year each in office printing and office duplication, virtually all of them monochrome; in contrast, the commercial printing and commercial publishing markets reached 3 trillion and 8.3 trillion pages, respectively, with higher-value color accounting for 50 percent of the former and 90 percent of the latter. HP's business was "3 or 4 percent of the total number of pages being printed," he remarked, noting that even if his company had 50 percent of that market and it represented $20 billion annually, it was "still only 3 or 4 percent." Growth, therefore, meant looking into other areas of printing, particularly as color, absent from the office printing and duplicating market, commanded higher prices.

HP responded by paying around $1 billion for an Israeli company that made a liquid-toner printer, Indigo, to jump into the commercial printing business rather than waiting for ink-jet technology to develop to the point that the company could enter using in-house technology. This move to hedge its bets had ruffled some feathers at HP, the market leader in both ink-jet and laser printing; the ink-jet division in particular thought itself capable of competing with any technology, including the one being acquired. Dr. Taub said the Indigo Digital Press represented the best of offset litho, thanks to fast, high-quality printing and flexibility with media; and the best of laser technology, being capable of short-run, fast-turnaround printing and, in fact, "everything you could want from a digital press."

To illustrate the future of his sector, Dr. Taub displayed a magazine created for a children's party as an example of the custom publication that can be produced using the Indigo press. Each girl's copy had her own photo on the cover, and half of the magazine was made up of poses of that one girl, the other half being one picture each of the other girls at the party, all taken by a professional photographer engaged for the occasion. A digital press made such a product possible: Other than on the Indigo press, the only way it could even have been approximated would have been using prints of the photographs themselves, which would have been extremely expensive. This capability opened the door to such unprecedented applications as home-delivered commercial magazines and direct-mail advertising whose content was tailored to the interests and preferences of the individual recipient, something that would save time for the reader and money for the sender. Physical inventory could also be reduced: The setup of a print run on an offset press was so costly that there was a tendency to print excess copies in order to avoid the expense of setting up a second run. Reports placed at between 40 and 60 percent the amount of offset-printed material—in the form of books, magazines, marketing literature, and most likely newspapers as well—that was discarded, while a significant quantity of the remainder needed to be stored in warehouses.

Although admitting its source was unclear, Dr. Taub cited an estimate that by the end of 2006 over 50 percent of marketing documents would be printed digitally. To illustrate the possibilities this would open up, he displayed a brochure from an auto distributor in California that had been delivered to an HP colleague within several days of the latter's having filled out an information request card and mailed it in. The brochure pictured a car of the model and color that he had indicated he was interested in and provided, in addition to detailed data on the car itself, a comparison of its characteristics with those of a directly competing model.

The addressee's name was embedded throughout the brochure's text, which was customized to the point of including the name and contact information of a local sales representative, so that the brochure appeared to come from a particular dealership even if it was in fact generated as part of a national marketing campaign. In addition, Dr. Taub pointed out, the brochure incorporated a method of measuring the effectiveness of the campaign: the offer to the addressee of a free oil change for bringing the brochure along when visiting the dealership. "Imagine if, having expressed interest in this particular car, you got something in the mail that looked like it was made for you in response to something that you had inquired about specifically," he said, adding: "Very powerful."

To corroborate this observation, he cited a study on direct-mail response rates by Frank Romano and Dave Broudy of the Rochester Institute of Technology showing that using the addressee's name to personalize a black-and-white brochure increased response 44 percent, from a typical level of around 3 percent to between 4.5 and 5 percent. A similar increase in response rate occurred when a full-color brochure was sent in place of a monochrome brochure but neither was personalized. A full-color brochure with the addressee's name added produced a response rate that was 135 percent higher than the nonpersonalized monochrome brochure had achieved, and using additional data to customize the brochure so that it resembled that received by Dr. Taub's colleague yielded a 500 percent increase in response over the baseline.

Displaying a chart summing up his points, Dr. Taub commented that at the lowest level of sophistication—one-to-many marketing in which there is no customization—a major advantage of digital publishing was the ability to print on demand (see Figure 22). This eliminated waste in two ways: by making it economical to print only as many copies as needed at any particular moment, and by making it possible to update or correct content between the resulting shorter print runs. The next level up, that of customizing or "versionizing," was that at which a publication's content was tailored to a city or dealer in order to obtain a result more relevant to a specific customer set while lowering distribution costs and improving timeliness. The level above that, one-to-one marketing, corresponded to the auto brochure described earlier. At the highest level, that of "event-driven" publications, a dealer could have a personalized brochure addressing an individual customer's concerns prepared while that customer was inspecting product on the dealer's premises and present that brochure to the customer at the end of the visit.

Dr. Taub then discussed what it took to keep a press running at all times, a key to profitability. In the case of a conventional Heidelberg press, among the needs were an uninterrupted supply of correct plates, rapid plate changeover, and alignment of the plates color-to-color on the various stations of the press, a time-consuming operation which was what he said "you really pay for when you submit a job for an offset press." With an Indigo press, all that was needed was an uninterrupted stream of correct data after an initial calibration. As many jobs

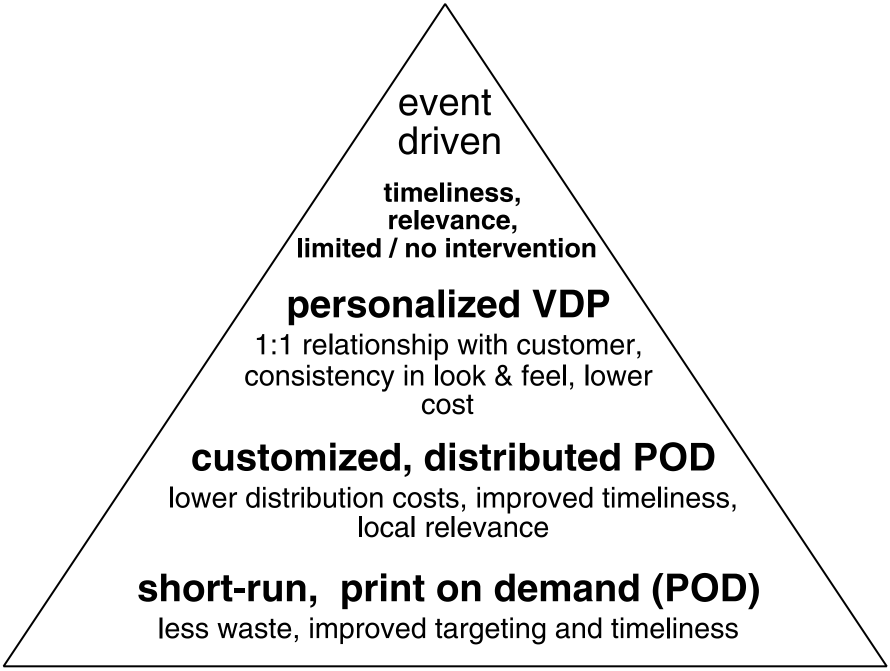

FIGURE 22 Opportunities for digital publishing.

could be sent through the press as could be lined up one after another, with the first sheet out in principle being usable. “When you finish one job, you just print the next job,” he said, “so this really stops being a technology or press-related event and starts to be basically an information-technology problem.”

Turning to the process of feeding data to the press, Dr. Taub presented an explanation of how a contemporary direct-selling campaign was developed. Beyond simply printing text or images on paper, the campaign designer needed to integrate customer relationship management (CRM) data and content data into the design of the campaign or its materials. Delivery to customers could be effected in printed form, on a laptop, or via PDA or cell phone, presentation being a major issue in campaign design; flexibility up front would enable delivery using a variety of outputs without the material’s needing to be redesigned. Based on feedback, CRM and campaign-management systems could be adjusted to increase the effectiveness with which the data could be used on the next campaign.

The greatest challenge for digital publishing, and one of its most telling metrics as well, appeared to be cost per page, an issue encompassing reliability, usability, and labor. Digital publishing was “clearly a lot more expensive” than offset printing, Dr. Taub admitted, adding that for the former to become "palatable" to potential customers cost had to be driven down "considerably." In addition, digital publishing would have to offer the same options offered by an offset press, such as working with different sizes of paper and providing a variety of finishing capabilities. Yet even more important was Web-efficient work flow, which would allow the press to be in constant operation. Databases would need to be integrated in order to create custom publications efficiently, which meant taking one source of content and generating multiple outputs from it. This was no small chore in the case of large companies with multiple sources of customer data—HP, classing among them, had five major databases devoted to this. Owing to the length of the path the data would have to follow, security protecting the marketing campaigns from the eyes of competitors would become a very important consideration as well. "It's a little bit like Napster," he noted. "Once the things become digital, it's much easier for them to escape, and you need to have a level of security that companies will be comfortable with."

Finally, the significant investment required to go digital was complementing the inertia that tends to stand in the way of change. "We need to help the commercial printers and enterprise customers make the transition," said Dr. Taub. "They need to understand the benefits of digital" beyond being shown its capabilities, so clear from the marketing techniques he had been discussing. Absolutely necessary was convincing them that the correct measure of digital publishing was value per page rather than cost per page, because "you get so much more impact from a customized page than from the standard offset page."

Mr. Borrus noted that the personalization techniques described by Dr. Taub would raise interesting issues for those who were trying to measure productivity increases, as they would be obliged to apply a deflator to printing.

Panel V

What Have We Learned and What Does It Mean?

Moderator:
Carol A. Corrado
Federal Reserve Board of Governors

David F. McQueeney, International Business Machines
William J. Raduchel
Marilyn E. Manser, Bureau of Labor Statistics
Kenneth Flamm, University of Texas at Austin
Jack E. Triplett, The Brookings Institution

Dr. Corrado praised the value of the day's discussions for economists like herself whose reason for attending had been to gain an understanding of the prospects for technology. She introduced Dr. Manser, head of the productivity group at the Bureau of Labor Statistics, as the only member of the panel who had not spoken in the course of the day's program, then called upon Dr. McQueeney to begin.

Dr. McQueeney observed that while many of the day's speakers had alluded to Moore's Law, all had offered different views of how it applied to their industries, each of which was unique. He raised a question about the future of the information technology industry that he had often discussed with colleagues at IBM: "Are we going at some point to slow down the rate of innovation, or are we not going to slow it down but cross some 'good-enough' thresholds, so that some parts of the industry will become mature whereas other parts perhaps will not?" If it were the latter vision that turned out to be the more prescient, he stated, "it means we've picked the low-hanging fruit and filled up the valleys with rainwater, and perhaps the way you create value changes." The industry had in the past been able to create value by applying microscopic, core-level technology at the bottom of the food chain and having the value trickle up to the top, the only place where it matters to customers; it was this, he noted, that might be in for a change.

A second point that had come up in the course of the day's discussions was that business value was no longer contained in one place: in one country, at one site, or in one company. Through web services, a way of finding and executing business capabilities electronically on the web, any geographical connection between the place where one business process was executed and the place where another linked business process was executed had been completely broken. As an example, he cited companies' using global resources to enable them to locate a help desk in a different part of the world from its technical operations. This phenomenon raised the question of whether value—to the customer, in the IT industry, or from the point of view of economic competitiveness—was any longer localized. "I think that it is not," Dr. McQueeney declared.

A third issue identified by Dr. McQueeney was that of whether investment in the consumer electronics industry would drive future innovation on the commercial side of IT. He noted that IBM's single highest-volume processor-chip business was in game systems, where the company was partnering with Sony and Toshiba. Since the microprocessor that gets the most design resources is the one that is most efficient and most advanced, he said, even though IBM could charge more for a mainframe microprocessor—and it might be able to make the best system, because the system is very complex—there was "no way that that market could afford the engineering bill to make the best microprocessor." Mainframe microprocessors 10 years hence would probably resemble processors developed for the "consumer-gaming/handheld part of the industry," whose customer base was so much larger than the amount of development investment that could be made in it as a fraction of revenue was much higher. Including displays and graphics chips with microprocessors, he said that the leisure-spending or disposable-income market provided "a tremendous source of investment for things that are at the bottom of the food chain for the commercial side of IT." It furnished new ways to innovate that had not existed when the IT industry was rather self-contained in its investments.

Speaking next, Dr. Raduchel pointed out that the IT industry was not static and said huge changes could be expected in the future. Noting that the industry was "very supply-driven," he rated the attempt to measure the contribution of IT to economic productivity as a very difficult challenge. Because advancement wasn't "being driven necessarily by customer" but by the fact that staying alive in the industry meant heading off competitors by putting out the best product possible, the process was characterized by technologists' efforts to do their best within the laws of physics. Such a process "doesn't lend itself well to hedonic price indexes" and, in general, poses obstacles to arriving at easy output measures.

Identifying a second challenge by the phrase "the network is the computer," Dr. Raduchel observed that the industry was prone to "talking about computers as isolated devices"—which, he added, "they are not." Networks were causing massive changes in how systems that real-world people used were structured and built. Alluding to a STEP Board plan for a conference on telecommunications, he cautioned that developments in switches were just as important as those in the PC for overall delivery to customers. Pointing to an earlier reference to the change brought by Dense Wave Division Multiplex (DWDM), he called "the biggest misjudgment in the capital market in the history of the world" the financial sector's failure to reckon with the potential invention of a technology that would increase the capacity of existing fiber-optic cable by a thousandfold. "Every time you touch this problem, you begin to realize how big it becomes," he observed. Moreover, none of it would be of any consequence unless software were brought into the picture, another major challenge. Software was often accused of being sloppy, but sloppy software might still be of great value in solving a problem. And sometimes the cause of its sloppiness—that many brains went into developing it—was also the key to its effectiveness.

Dr. Manser began by noting that a great deal of attention had been paid to the importance of the high-tech sector in explaining the productivity growth and the productivity speedup that occurred during the latter part of the 1990s. High-tech equipment affects labor-productivity growth in two ways: through the use of high-tech capital services (that is, the flow of services from the stock of high-tech equipment and software) throughout the economy by producers in all sectors, and through productivity improvements in the industries that produce high-tech equipment. BLS data for the non-farm business sector showed that, in combination, those two high-tech effects accounted for roughly two-thirds of the speed-up in labor-productivity growth that occurred in the latter part of the 1990s relative to the 1973–1995 period, a result she called "striking."

BLS also produces labor-productivity and multifactor-productivity measures for industries in the U.S. economy according to their three-digit classifications. Labor-productivity change is measured by relating changes in real output to changes in worker hours. In multifactor-productivity change, which is somewhat more complicated to measure, BLS relates real output to changes in not only worker hours but also capital services, intermediate purchases, and, in some cases, worker skills. For the period 1987–2000, generally regarded as a period of good performance, the rate of growth of labor productivity in the U.S. non-farm business sector was 1.8 percent per year on average, which was generally regarded as strong. BLS calculates productivity using data from a variety of sources on, among other factors, revenues, prices, and labor hours. In the four-digit industrial classifications that comprise the computer and semiconductor area, the strongest measured labor-productivity growth was 42 percent per year for semiconductors and related devices, a rate she called "pretty phenomenal compared to the 1.8 percent for the economy as a whole." The data also showed labor productivity growth of 37 percent per year on average for electronic computers, 17 percent per year for computer peripheral equipment not elsewhere classified, and 15 percent per year for computer storage devices.

Dr. Manser had pulled these figures together before attending the conference because she wanted to know whether what she heard from the day's speakers would raise questions about the measures BLS had been using, controversy being constant in the measurement community over how the Bureau was measuring output in the high-tech sector. Her conclusion, particularly in light of the "phenomenal things that had happened in these industries over time," was that nothing that had been said suggested that the BLS numbers were overstating labor-productivity growth. While the way prices and real output are measured is very important to BLS in general, also important is using price measures that are consistent over a long period, because looking at productivity means looking at trends over time.

Turning to the question of workers, about which she noted little had been said in the course of the day, Dr. Manser said that BLS charted what it calls "labor composition," which concerns the impact of changes in workers' skills on productivity. BLS's official series on labor composition broke out data only for the overall business sector and for the non-farm business sector. It would be important to examine the impact on productivity of changes in the composition of labor for the high-technology industries, she said, adding that she hoped this could be done in the context of work BLS was embarking on jointly with the Census Bureau, where a "very rich" new dataset had been developed, to expand understanding of labor composition.

Dr. Flamm said he would offer two quick comments, the first of which was that there had been a certain tension among some of the presentations. A number, Dr. McQueeney's and Dr. Bregman's among them, had focused on the problems of measuring technical advance and productivity in the large, complex systems used by large organizations. Dr. Flamm wondered whether the big volumes and big dollars in computer hardware and software were at that moment centered in the large-enterprise, complex-system market or in the small-to-medium business and consumer market. If the market were conceived of as being bifurcated in that way, then measurement problems for many of the issues discussed would be far less arduous for the small-to-medium enterprise and consumer than for the large, complex systems, whose owners are writing their own code and constructing their own storage systems. Dr. Flamm proposed the following as a characteristic to distinguish the two markets: In the small-to-medium enterprise and consumer business the storage administrator was the user, whereas in the market made up of large, complex systems there were storage administrators who were not the users. He said that Dr. McQueeney had been unable to tell him which market was larger and that his own guess was that most of the dollars were actually in the simpler, stand-alone systems.

Dr. Triplett countered that the services industries were the big buyers.

But Dr. Flamm, insisting that he would like to see the numbers, reiterated that the problem of measuring big, complex systems was different from that of figuring out what the PC and the software used by his accountant cost and what the change had been over time.

His second comment was that perhaps one way of approaching the issue of how to go about figuring out which characteristics were the right ones to measure—which was, he noted, one of the objectives of the day's meeting—would be to set aside the notion that there was some fixed set of characteristics for everything that would need to be documented consistently over time. Following up an earlier statement by Dr. Silver, he suggested instead coming to grips with the fact that characteristics' relevance might shift as technology advanced. "There are waves or generations of technology," Dr. Flamm noted, "and there's one set of characteristics that are appropriate for one wave or generation, but then the world shifts and everyone quickly forgets about the previous wave." Once one technology has given way to another, the businesses involved are no longer concerned about collecting data used in metrics that applied to the previous technology—unlike economists and business statisticians on the one hand and, on the other, the government, to the extent that it needs the data to enforce export controls. Concluding, he noted that export controls induced businesses as well to collect the older data and suggested that persuading the government to take a list of characteristics drafted by the STEP Board as the basis of its export-control system might open up a very fertile source of data.

Dr. Triplett began his remarks by suggesting that the tension of which Dr. Flamm had spoken might have resulted from the variety of the professional groups represented among the conference presenters and participants. With the focus of both the presentations and the questions differing to a noticeable degree, it was possible that the bridge between technologists and economists had not been completed, but he judged the session to have been productive in any case.

Dr. Triplett said that hearing experts in the field say that software was very hard to measure had been a learning experience for him. "I thought I knew that when I came in here," he observed, "but part of my concern was that I thought maybe it was hard to measure because I didn't know enough about it. So the good news is that it wasn't my fault." The bad news, however, was that there had not been as much progress made measuring software as had been desired. He pronounced himself "not too sanguine" on significant results' being obtained from analyses of the impact of software changes on customers, saying that in the complex world of the business environment it would be hard to hold constant all that would need to be held constant in order to perform the experiment of putting in a software innovation and seeing what its effect was. "I'd like to see somebody try it," he said, "but if I had government research money to hand out, I'm not sure that would be one I'd think would be really profitable." The idea had, in fact, put him in mind of a presentation he had seen on measuring the contribution of consulting firms to their clients' profits—a "great idea," but one that ignored the complexity of the world in general and, in particular, the manifold reasons for which firms hire consultants.

Dr. Triplett then turned to the benchmarks, saying that the "possibility of using real information on the cost of doing a job with different kinds of methodologies" had been discussed in economics literature since at least the 1980s. A first problem with benchmarks was that in time-series comparisons, which are what interest economists, only the change in the cost of what has been done in the past can be benchmarked. "You can figure out what you did yesterday, and you can figure out the change in the cost or the speed of doing that today," he said, but "what you can't value is doing today what you couldn't do yesterday." In the history of the impact of the computer, the most significant gains have not come in cutting the time it takes to type a letter but in making possible things that either simply had not been feasible before or had not been feasible in the same way. It is the task itself that has changed.

A second problem with benchmarks was that many tasks must be combined in some way. Offering an example of an office-productivity benchmark, Dr. Triplett posited as a task the time taken to execute the "replace all" command in Microsoft Word (which replaces all instances of a certain word with another word). Breaking activity down in this way caused data to be gathered on a great number of different tasks from which an aggregate must be derived. But the aggregate was often arrived at simply by averaging the tasks up—in the absence of any valuations for the tasks. And even if various sets of tasks were successfully benchmarked, it might be difficult to weight them across users with different requirements. While praising the potential of benchmarks, he cautioned that still to be thought through were methods of obtaining benchmark information, which was not yet available in sufficient quantities, and of aggregating it.

DISCUSSION

John Gardinier began the question period by commenting that he had recently retired and was trying to start a small business. He said he had seen a significant shift on the horizon that would make more applicable to the real world a basic tenet of economic theory: that in a competitive market correct information is available immediately to all the players. Making price comparisons during the previous week using the Internet, he had found the identical computer peripheral having a price span of over 100 percent; on airline fares he found a spread of a factor of five. "They used to be able to fool me," he remarked. "They can't fool me anymore." Because this information was available, businesses would be under increasing pressure from better-informed consumers.

Dr. Flamm responded that a huge literature was developing on price dispersion on the Internet, which was staying constant or increasing.

Dr. McQueeney, noting that airlines had put more and more optimization into their yield maximization of profit vs. load, said "an arms race" was in progress between consumers armed with *cheaptickets.com* and the airlines armed with their own tools.

Iain Cockburn of Boston University observed that even as competitive pressure and technological development produced huge gains through the rapid expansion of capabilities and decline of prices, there was a cost: "Things become impossible that you used to do." As an illustration, he cited the difficulty of retrieving data stored on half-inch tape or a 5.25-inch floppy disk. Display technology, in what he called a "striking difference," seemed immune to obsolescence and its associated costs.

Dr. McQueeney recalled that, a month before, a group having lunch at his research lab postulated that someone handed them an 8-inch floppy disk with winning lottery numbers that were to be drawn in an hour. The unanimous consensus was that those at the table would not have been able to collect the money because they would not have been able to find a reader in time.

Bill Long of Business Performance Research recounted having received two Web documents as part of a single STEP Board bulletin—one notifying him of the current meeting, the other of a workshop on research and development in data needs—and said that the juxtaposition of the two had, over the course of the day, brought certain questions to mind. Alluding to personal frustration with the quality of data he has collected and worked with, he indicated that he had been intrigued when speakers talked about "the possible use of ROI calculations either by the seller or the buyer" of information technology, as well as about "payback data—'we paid for it in six days or six months,' or whatever." He said, however, that the data referred to had sounded "very proprietary" and characterized one speaker's attitude thus: "We wouldn't tell you even if we had it, and I'm not sure I'm going to tell you whether we have it." Noting it was the job of many in the room to make sense of productivity gains and to determine what caused them, he asserted that success would depend on "some government agency collecting some data it's not collecting now using a classification system that probably doesn't yet exist." He asked whether there had been any progress in that direction and what sort of progress might be expected over the next five to 10 years.

Dr. Corrado responded that comments throughout the day had provided fuel for the view that a broader definition of "business fixed investment" might be appropriate in describing the economy. "I don't know if those of you who were speaking realize that you were supporting the view that R&D expenditures should be capitalized, that is, treated as business investment, in our national accounts," she said. In view of the importance of innovation and technology in today's economy, the measurement of R&D expenditures and considering whether some uses of employees' time represent investment rather than inputs to current production are important topics in productivity research. Although studies have expanded the production boundary to encompass R&D and related outlays, the precise scope of these expenditures, and their measurement in real terms, are not settled issues and remain challenges before us.

Concluding Remarks

Dale W. Jorgenson
Harvard University

Thanking the panelists, Dr. Jorgenson proposed to answer a pair of questions that he had raised at the opening of the conference. He restated his first question thus: “Could we measure progress in the computer industry and had we done it?” Answering both parts in the affirmative, he said it had been demonstrated not only that measuring progress was possible but also that such measurement had become steadily more sophisticated and, in fact, “quite successful.”

The second question was whether it would be possible to use a road map resembling that maintained by the Semiconductor Industry Association to project developments in information technology generally and in the various fields discussed by the day’s presenters specifically. Again Dr. Jorgenson answered in the affirmative, although he cautioned that the institutional framework required could in some instances demand “a good deal of thought.” As evidence for his position, he pointed out that the road map for display devices described by Dr. Keys, which the U.S. Display Consortium had published only a few weeks prior to the meeting, had been modeled on the semiconductor industry’s road map. Printing and storage devices were other areas in which establishing road maps “clearly would be feasible,” he stated, and another was computers, although this last case was one of those in which the institutional framework would be of some concern. “Those are only the questions that I posed at the outset,” Dr. Jorgenson said, “so I’m very happy with the outcome of the symposium. I think we made a lot of progress on these questions and filled in a lot of the missing pieces of information."

Looking to the future, he urged optimism on his listeners as "a good posture in general in research" while also presenting reasons for his own, personal optimism. Dr. Triplett, he reminded the audience, had described the history of research on computer prices—"the economist's way of encompassing all these technical developments that we've been here to describe"—going back to the beginning of the computer's commercial history in the 1950s. A set of measures for computers and peripherals, grounded in work that IBM researchers began publishing in the economics literature in the late 1960s, achieved incorporation into the U.S. national accounts for the first time in the mid-1980s and had continued in use, while also being enhanced and developed, to the present day. These measures had been extremely informative, especially in helping to understand the recent behavior of the economy.

Only the month before, Dr. Jorgenson said, Dr. Corrado had sent him the Federal Reserve Board of Governors' first cut at a set of official statistics for telecommunications equipment, which had been intended to fill in what had up to then been "a 'black hole' in economic understanding." Economists, for example, had not previously understood the role of DWDM, a very rapidly developing technology, but it had now been encompassed. While assuring the audience that the Board of Governors' work would be enhanced and improved as better data became available, Dr. Jorgenson likened the effort to "the beginning draft of the human genome: It's not the thing that you really wanted at the end of the day, but on the other hand it's where you want to start."

Finally, Dr. Jorgenson came to what he characterized as "the great challenge of software." The software industry had been growing at more or less the same rate as the hardware industry since the beginning of the computer's commercialization, he noted, but employment in the production of software had been growing about 10 times as fast. "There's something there that we don't fully understand," he admitted, but he predicted: "The challenge of dealing with the issues having to do with software, although it lies ahead of us, will yield to methodologies similar to the ones that we were discussing today."

Dr. Jorgenson again expressed his gratitude to all in attendance, and especially to the presenters and panelists, for taking part in what he called "this very fulsome discussion of a very important topic." The STEP Board was planning more meetings in its series "Measuring and Sustaining the New Economy," with the next one to focus on telecom, aided by the recent work of Dr. Corrado and her colleagues, and the one following that on software. Encouraging his listeners to "stay tuned," Dr. Jorgenson said he looked forward to their participation in this continuing discussion.

II

RESEARCH PAPER

Performance Measures for Computers

Jack E. Triplett
The Brookings Institution

I. INTRODUCTION

The "Deconstructing the Computer" workshop has the purpose of gaining better understanding of computer performance, especially the contributions of computer components to computer performance. Two groups of professionals are interested in measuring the performance of computers, peripherals, and components. This paper provides a bridge between their interests.

Section II explains, primarily to computer professionals, why economists want to measure computer performance and what economists do with performance measures. Subsequent sections provide background on economists' work on measuring computers and components. As this workshop is part of the STEP Board's "New Economy" project, one of its objectives is obtaining better performance measures for economic uses.

A second audience consists of economists. It is clearly true, as Nordhaus (2002) remarked, that computer performance measures used by economists in recent years have, if anything, gone backward compared with the measures they used 15 or so years ago. We need to ask "why?" We also need to ask: "How much does it matter?"

I review in section III the performance measures used by economists in the earlier computer literature, which covers primarily the mainframe years. Sections

IV and V review performance measures used by economists and by statistical agencies in more recent years, where studies have turned predominantly to personal computers (PCs).

II. WHAT ECONOMISTS DO WITH COMPUTER PERFORMANCE MEASURES

I begin by addressing technologists. Why do economists want to measure computer performance? And what do they do with performance measures? Technologists need to understand how economists use computer performance measures in order to converse with economists on this topic. The questions do not imply that the performance measures wanted by economists are the only performance measures that matter, but fortunately it turns out that what economists want is not that different from what technologists have developed. Indeed, historically, technologists and economists have proceeded in similar directions in measuring the performance of computers. But that gets ahead of the story.

Suppose, to create a simple illustrative example, one computer exists. Call it UNIVAC. Suppose three UNIVAC computers are made in 1952, and they cost $400K each.[1] We are supposing that UNIVAC was the only computer produced in the economy, so total U.S. output of computers in 1952 was $1.2 million.

Now suppose a new computer is developed in 1955 (I call it "new computer"), and that it has higher performance than UNIVAC. Suppose "new computer" sells for $600,000 in 1955, and suppose further that it is the only computer available in 1955, the UNIVAC having disappeared. Ten computers of this new type are produced in 1955, so the economy's output of computers is $6 million in 1955, a fivefold expansion since 1952 in what economists call "current-price" output. This is clear enough, but other aspects of the 1952–1955 comparison are less clear (see Table 1).

First, is there inflation in the computer market? "New computer" costs 50 percent more than UNIVAC, but the new computer also has higher performance. Part of its higher price is just a performance premium. Economists do not want to show an increase in computer performance as inflation. They want to measure computer inflation so that it is adjusted for changes in the performance of computers; in other words, computer inflation should be measured net of the performance premium. The example suggests that computer inflation was less than the 50 percent increase in selling price. How much less? To determine that, economists need a computer performance measure (or more precisely, the performance premium).

What about computer output? "New computer" has higher performance than UNIVAC, so each "new computer" is equivalent to more than one UNIVAC.

---

[1] These numbers correspond to UNIVAC production in 1952. See Flamm (1988), Table 3-1 for the price and page 51 for the quantity.

TABLE 1 UNIVAC and "New Computer," Hypothetical Price and Output Calculations

	UNIVAC (1953)	"New Computer" (1955) (Case One)	"New Computer" (1955) (Case Two)
Number produced	3	10	10
Price, each	\$400 thousand	\$600 thousand	\$600 thousand
Current price output	\$1.2 million	\$6.0 million	\$6.0 million
Performance index (M)	1.0	1.5	1.8
Computer inflation (with estimated performance premium = 1 + .7 (ΔM)	1.00	1.11	.96
Computer "constant price" output index	1.00	4.60	5.17

Computer output must have expanded by a factor greater than the threefold increase in units produced. How much greater? To answer that, economists also need a measure of computer performance.

Economists also want to calculate the productivity of making computers, just as they calculate productivity in other industries. One common form of productivity is labor productivity, defined as output per worker hour. Again, if "new computer" has higher performance than UNIVAC, economists want to calculate output per labor hour in producing computers with a "quality adjustment" that incorporates the improved performance of the new computer. For estimating productivity change, "new computer's" higher performance must be factored into the output measure. Similar statements apply to other economic measures, particularly to computer investment and capital stock.

Thus, for measuring inflation, output growth, productivity growth, the volume of investment and capital stock, and for other economic measurements, economists need a measure of computer performance. It is well known that a bottom-end desktop computer today greatly outperforms anything available at the dawn of the commercial computer age, which was 50 years ago. Counting the number of computers produced will never tell us much about trends in computer output. The great expansion of computer output in the last 50 years is an expansion not only in numbers of computers but also in what might be thought of as "output per computer produced," that is, performance per machine.

We need now to discuss the properties that economists want in their measures of computer performance. To carry this forward, suppose now that we all agree on a measure of computer performance. It should not be too surprising that, as I discuss in the following section, achieving measures of computer performance is not at all a straightforward task. But set that aside, for the moment.

Suppose we have an agreed-on measure of computer performance that covers the UNIVAC and the 1955 "new computer." Suppose that we standardize our performance measure so that the UNIVAC has 1.0 performance units, and the new computer has 1.5 performance units.[2]

Unfortunately, even when computer performance is a scalar measure, we cannot simply divide the value of UNIVAC or "new computer" production by the performance measure in order to compare 1952 and 1955 computer output. Economists need the value of the performance indicator. There are several reasons. An old computer relationship called Grosch's Law indicates that the cost of a computer center does not rise linearly with its computing power. Similar arguments can be made on the demand side: The incremental value of improved performance to the user does not necessarily rise proportionately with an increase in performance. Thus if "new computer" has 1.5 times the performance of UNIVAC, we need some way to value this 1.5 performance improvement ratio. We must know the performance premium, a value measure, not just the increment in performance.

The valuation problem is truly daunting. Likely, UNIVAC and the replacement computer do not appear in the market at the same time. If they do, "new computer" should sell for more, and it is natural to take the ratio of the two machines' prices as measuring the value of their relative performance. All kinds of problems exist with that, which I do not mean to minimize. For example, the high-end user might be willing to pay more than the actual price premium for "new computer" to get a high-end machine, but the low-end user might not be willing to pay the price difference; if so, the price differential only reflects the value of the performance difference to the user who is on the margin between buying the one or the other. But these are essentially aggregation problems (over users), which I set aside because they arise throughout economic statistics of this kind.

A more promising situation exists empirically if there are a large number of computer models, and we have data on their prices and their performance. One can then run a regression, such as:

$$P_i = a_0 + a_1 (M)_i + e_i \tag{1}$$

In equation (1), P is a vector of prices of computers, where models are indexed by the letter i, M is the associated performance measure for each computer, and e_i is the regression error term. Using equation (1), we estimate a_1 and use a_1 to put a value on the performance difference among machines: If UNIVAC has $M = 1.0$ and "new computer" has $M = 1.5$, then the quantity $[a_1(0.5)]$ gives a "quality adjustment" that can be used to value the difference between the two machines.

---

[2] And we suppose, contrary to what is true, that computer performance can be represented as a simple scalar.

Suppose that we estimate a_1 to be 0.7. Computer inflation between 1952 and 1955 ("quality adjusted" for "new computer's" performance premium) is then: \$600K / {\$400K ((1 + 0.7(0.5))} = 1.11, or 11 percent inflation. This number is clearly less than the 1.50 (equals 50 percent inflation) that the unadjusted data would show. If "new computer" has $M = 1.8$, then computer prices are falling: \$600K / {\$400K ((1 + 0.7 (0.8))} = 0.96, or 4 percent price decline.

Turning to computer output, the usual method for measuring output changes (in the national accounts, for example) is to "deflate" expenditures on a product by its price index (information on the U.S. national accounts is in Bureau of Economic Analysis, 2001). To form a deflated measure of computer output, we start from the change in "current price" output, which in our example was (\$6.0 mil – \$1.2 mil) / \$1.2 mil, equal to a 400 percent increase. Deflating that by the price index of 1.11 gives for the "constant price" output change a 360 percent increase between 1952 and 1955. Deflated output grows less than current price output because in this example ($M = 1.5$) computer prices, performance adjusted, were rising.

When "new computer" has a larger performance differential over UNIVAC (in the second example, $M = 1.8$), the price index declines, to 0.96, or a 4 percent decline. Using this declining price index as a deflator results in a "constant price" output change that is larger than the "current price" change (400 × (1/.96) = 417). In national accounts, this "constant price" output measure is sometimes (rather inappropriately) known as "real output."[3]

This simple example illustrates several principles that govern estimation of computer output and investment in the U.S. national accounts. The most important one is that it shows how strongly measures of computer performance influence economic measurement of computer price change, and through the deflation procedure, how strongly computer performance affects measures of computer output, investment and productivity.

Equation (1) is a relation that is known in economics as a "hedonic function," although equation (1) is a very simple hedonic function. The "quality adjustment" outlined in the preceding paragraph is, in essence, the method applied by the Bureau of Labor Statistics (BLS) in estimating price indexes for computers, where quality adjustments for enhanced computer performance are derived from a hedonic function. This example is far too simple, however.

In general, computer performance is not a scalar; it is multidimensional (the implications of this are explored in the subsequent section). Thus, computer hedonic functions look, generally, like equation (2):

$$\ln P_i = a_0 + a_1 \ln (M_1)_i + a_2 \ln (M_2) + \ldots a_k \ln (M_k) + e_i \tag{2}$$

---

[3] BEA also uses the somewhat cryptic term "chained dollars" to represent the same thing, and reports percentage changes in the form of index numbers, under the title "chained-type quantity index."

Each of the k variables in equation (2) is a "characteristic" of computer performance. The current BLS hedonic function for personal computers has more than a dozen characteristics. Each of the coefficients, a_k in equation (2), is interpreted as the value of the corresponding computer performance characteristic. Because I have written equation (2) in a logarithmic form (the hedonic function often turns out to be logarithmic, but not always), these coefficients are not prices denominated in the usual dollars or euros, but dollar and euro prices can be extracted from the coefficients, if desired.

Hedonic price indexes have become the standard economic tool for measuring price change in computers. In principle, they measure the price of computing power.

Getting from the price indexes to the output investment numbers is relatively straightforward and follows the example already presented. Table 2 shows current dollar changes for computer investment in the national accounts, the computer deflator, and the resulting deflated investment numbers from the national accounts, for the years 1995–2002. In 1995 computer and peripheral "current price" investment in the United States equaled $64.6 billion. In 2000, the value of computer equipment investment equaled $93.3 billion. Thus, in current prices computer equipment investment increased by 44 percent.

The computer equipment price index declined by 71 percent over the same 1995–2000 interval (from 131 to 38, using 1996 as the base). The change in current price shipments divided by the change in the price index gives the "deflated" (also called "constant price" or "real") value of the change in computer equipment investment over that interval: As Table 2 shows, this increased fourfold (the quantity index goes from 69 in 1995 to 348 in 2000). The source of the great increase in computer investment in the national accounts numbers is not only the increase in spending on computer equipment, but also the decline in performance-corrected prices for this equipment.

This same point is dramatically illustrated by the post-2000 experience. Actual spending on computer and peripheral equipment fell by 20 percent. But the

TABLE 2 Private Fixed Investment in Computers and Peripheral Equipment

	1995	1996	1997	1998	1999	2000	2001	2002
Billions of current dollars	64.6	70.9	79.6	84.2	90.4	93.3	74.2	74.4
Billions of chained 1996 dollars[a]	49.2	70.9	102.9	147.7	207.4	246.4	239.9	284.1
Chained price index	131.29	100	77.38	56.99	43.6	37.87	30.91	26.27
Quantity index	69.4	100	145.22	208.39	292.64	347.77	338.61	400.92

[a] "Chained dollars" is the Bureau of Economic Analysis name for a quantity index of computer equipment output (see text).

SOURCE: Bureau of Economic Analysis, NIPA Tables 5.4, 5.5, and 7.6.

national accounts quantity index increased by 15 percent over the same 2000–2002 interval, because the price index fell by 30 percent (see Table 2).

Tables 3 and 4 show that these trends have been going on for a long time. The price of computer equipment (computers plus peripherals) has declined 17.5 percent per year over the whole period for which national accounts investment data are available. Moreover, over the whole of the historical period, prices of computers themselves have declined faster than prices of peripherals (this is evident from Tables 3-5). The price of computing power today approaches 1/1,000 of 1 percent of what it was at the introduction of the commercial computer 50 years ago (Table 4).[4] Additionally, the prices of ancillary devices have also fallen, though their performance improvements are often overshadowed by the spectacular progress in computer hardware: *PC World* (March, 2003, page 91) reports that the cost of storage media (disks) has fallen from $16.25 per MB of data stored in 1981 to $0.0008 (8 percent of a penny) in 2003, an annual rate of decline of 36 percent, comparable to the rate for PC computers over the same interval.

Computer price indexes fall because the performance of computers is increasing very rapidly, where their actual selling prices are stable or falling. Accordingly, it is no surprise that computer output in the economy rises not so much because increasing numbers of computers are produced (though this is true) but because the capability of the computers that are produced has increased so much. The great increase in computer investment over the last 50 years as measured in the national accounts is in large part an estimate of the value of increased performance of computers over this interval.

Nordhaus (2002) takes the price of computing back another 50 years, using a different approach. Though the rates of decline in the last half century are greater than in the half century before that, Nordhaus' results indicate that high demand for improvements in computational power has existed over a long time, as well as indicating the extraordinary fruits of innovative ability set to satisfy that demand.

Price indexes for computers transfer directly into economists' measures of the output of computers, of "real" (an economist's somewhat misleading jargon) computer investment and capital stock, and from these the rate of productivity improvement. As examples of the latter, Jorgenson, Ho, and Stiroh (2002) estimate that the contribution of ICT (information and communication technology) investment was responsible for a large proportion of the acceleration in U.S. labor productivity in years following 1995. Triplett and Bosworth (2002) reached comparable findings for the importance of ICT investment to the substantial gains

---

[4] The mainframe index in Table 4 gives a beginning/end value of 3.91^{-05}. But over the period for which mainframe and PC price indexes are available (1982 forward), PC prices have fallen at 21 percent per year in the government indexes, where mainframes have trailed, at 18 percent per year (Table 4). Moreover, studies suggest that the government PC price index records too little decline, certainly over the first part of this period—for example, Berndt and Rappaport (2001, Table 1) indicate that PC prices fell over 30 percent per year between 1983 and 1999. Hence, taking all this together, the round number 1/100,000 in the text.

TABLE 3 Private Fixed Investment in Computers and Peripheral Equipment

	Price Index (1996 = 100)	Quantity Index (1996 = 100)
1959	101372.4	0.000
1960	79593.2	0.00044
1961	58800.8	0.00077
1962	41710.3	0.00143
1963	27395.1	0.00396
1964	22916.4	0.00616
1965	18936.0	0.00979
1966	13272.8	0.02354
1967	10784.1	0.03091
1968	9202.6	0.03696
1969	8332.3	0.05038
1970	7484.4	0.0605
1971	5698.7	0.07458
1972	4592.4	0.11
1973	4354.0	0.12
1974	3554.9	0.15
1975	3288.5	0.15
1976	2746.5	0.23
1977	2390.1	0.34
1978	1616.8	0.66
1979	1339.7	1.07
1980	1045.6	1.69
1981	918.9	2.63
1982	822.3	3.24
1983	685.6	4.92
1984	554.6	8.04
1985	471.5	10.10
1986	406.2	11.61
1987	346.0	14.59
1988	321.4	16.67
1989	300.1	20.27
1990	272.3	20.03
1991	244.6	21.75
1992	209.2	29.40
1993	178.4	37.31
1994	157.3	46.00
1995	131.3	69.40
1996	100.0	100.00
1997	77.4	145.22
1998	57.0	208.39
1999	43.6	292.64
2000	37.9	347.77
2001	30.9	338.61
2002	26.3	400.92

SOURCE: Bureau of Economic Analysis, NIPA Table 7.6, and unpublished BEA data in possession of the author (with more precise quantity index for earlier years). In 1959, the quantity index (1972=1) equals 0 to three decimal places in the unpublished data.

TABLE 4 Price Indexes for Domestic Mainframes and PCs (1996 = 100)

	Mainframes	Personal Computers
1953	791125.1	
1954	682645.1	
1955	605330.6	
1956	516628.7	
1957	456095.6	
1958	412943.3	
1959	354208.3	
1960	260717.8	
1961	197440.1	
1962	143811.5	
1963	109881.7	
1964	83549.6	
1965	60060.6	
1966	22761.6	
1967	16110.6	
1968	14560.0	
1969	14513.8	
1970	13967.4	
1971	10847.4	
1972	8871.7	
1973	9453.9	
1974	8041.6	
1975	7771.7	
1976	7106.5	
1977	5582.2	
1978	2812.2	
1979	2306.9	
1980	1591.9	
1981	1311.4	
1982	1106.3	1549.9
1983	1006.8	1086.5
1984	727.1	937.0
1985	537.1	877.3
1986	486.5	646.4
1987	419.1	582.9
1988	397.1	533.3
1989	346.5	496.1
1990	307.5	415.5
1991	297.6	350.4
1992	277.2	267.9
1993	234.2	207.3
1994	182.1	182.7
1995	144.0	145.3
1996	100.0	100.0
1997	68.6	67.1
1998	49.2	40.3
1999	38.6	25.7
2000	30.9	20.7

SOURCE: Triplett (1989) and unpublished Bureau of Economic Analysis data.

TABLE 5 Average Annual Growth Rate, Price Indexes for Domestic Mainframes, PCs, and Computers and Peripheral Equipment

	Mainframes	Personal Computers
1982–1987	–17.6	–17.8
1987–1995	–12.5	–15.9
1995–2000	–26.5	–32.3
1982–2000	–18.0	–21.3
1953–2000	–19.4	
	Computers and Peripheral Equipment	
1959–1969	–22.1	
1969–1987	–16.2	
1987–1995	–11.4	
1995–2002	–20.5	
1959–2002	–17.5	

in labor productivity experienced in services industries in recent years. Services are the industries that purchase a predominant portion of U.S. investment in ICT equipment. The research results in these two (and other similar) papers could not have been generated without economic measurements that incorporate performance measures for computers. At the moment, economic statistics that incorporate computer performance measures are lacking in most other OECD countries (Wyckoff, 1995; Colecchia and Schreyer, 2002), which greatly inhibits the ability to analyze recent productivity trends and the contribution of ICT investment in countries outside North America.

In sections III–V, I consider the variables that have been used as characteristics of computer performance in computer hedonic functions and the measures of computer performance that one would like to have for economic measurements.

III. COMPUTER PERFORMANCE MEASURES IN EARLY STUDIES—MOSTLY MAINFRAMES

It is intriguing that what an economist calls a hedonic function has appeared, in an essentially equivalent form, in the computer science literature. The earliest research on computer performance measurement grew out of, or was influenced by, research issues in the computer systems literature. Some performance measures were devised as a practical aid to equipment selection. Alternatively, computer technologists wanted to estimate the rate of technical change in computers. It makes no sense to do that without considering both performance and the price.

So where the economist naturally thinks of the price of computers, *adjusted* for performance, the computer technologist thinks of performance per dollar spent on computers. For example, Knight (1966, 1970, 1985), who is actually an economist but was writing for computer publications, was interested in estimating the rate of technical progress for computers and not a computer price index. He estimated an equation that was similar to equation (1).

Economists were interested in a somewhat different but closely related problem: measuring performance-corrected price indexes for computers, using for the most part hedonic methods. Among economists, the early hedonic researchers on computers more or less followed the lead of technologists in choosing their performance measures. The following section is partly based on Triplett (1989).

Performance Characteristics of Computer Processors

From the earliest studies, the performance specification of computer processors consisted primarily of the speed with which the computer carries out instructions and its memory size (main memory storage capacity).[5]

It has always been difficult to obtain a publicly available measure of speed that is both sufficient and at the same time comparable across processors. This remains the problem today.

A computer executes a variety of instructions. The execution rate of each instruction is properly a computer characteristic. Computer "speed" is accordingly a vector, not a scalar.

Applications require instructions in different proportions or amounts (e.g., graphics and office productivity programs). Moreover, different users, even if they employ the same applications, employ them in different frequencies—I use both these two applications, but my usage differs greatly from the usage of a graphics designer. Accordingly, numerous measures of "speed" exist, in principle, because speed is a vector and there are many ways of valuing the speed vector.

Nevertheless, some scalar summary of the speed vector is needed. Three major approaches have been employed by economists in hedonic studies. In considering these, it is well to bear in mind the twin aggregations of the speed vector—one aggregates over instructions (for application speeds); another aggregates application speeds over users who use them in different proportions.

Single-Instruction Speed Measures

In this approach, the speed of one instruction is chosen (in early studies, it was invariably addition time or multiplication time), which then serves as a proxy

---

[5] Phister (1979), Sharpe (1969), and Flamm (1987) contain good statements of the rationale for the specification, and Fisher, McGowan, and Greenwood (1983, pp. 140–141) emphasize its limitations.

for the rest. Single instruction speed measures were prominent in early computer studies (see Annex Table A). Even in the early days, they were not adequate. In analyses of instruction mix frequencies cited in Sharpe (1969, pp. 301–302) and Serlin (1986), additions accounted for only between 13 and 25 percent of total instructions, and multiplications around 5–6 percent. “Logic” or “other” or “miscellaneous” instructions, not easily measured at the time, were the largest category. A single-instruction speed measure will not adequately characterize a cross-section of computers or represent the change in computer performance over time.

Intermediate-Stage Proxy Measures

In this approach, the investigator looks for a machine specification that is correlated with the vector of performance characteristics. In early studies memory cycle speed was a popular proxy speed measure for computers. Memory cycle speed is memory cycles per second, or its inverse defined as the time (in microseconds) to read a word from the main memory and replace it. Memory cycle time is correlated with the speed of other processor operations and therefore acts as a proxy for those other determinants of speed. Closely related measures also appear in the regressions of Chow (1967—memory access time), Michaels (1979), and Fisher, McGowan, and Greenwood (1983—transfer rate).

Another intermediate-stage proxy measure is machine cycle time, also known as “clock speed.” The execution time of the logical portion of any instruction equals machine cycle time multiplied by the number of machine cycles required for that instruction.

Even in the mainframe days, it was well established that the relation between clock speed and instruction execution speed will shift with the instruction mix. Across machines, moreover, the relation varies with machine design. Thus, machine cycle time or clock speed contains the potential for substantial proxy error, both from machine to machine in the cross-section and over time. Economists who used clock speed either (a) did not understand its shortcomings, or (b) understood them, but used clock speed because it was widely available for a large sample of computers.

Benchmarks

Single-variable proxies for a multivariate vector of instruction speeds will always present the problem that the particular proxy chosen may represent very poorly the speed at which a computer performs actual jobs. It is thus natural to measure computer speed by presenting the same job or mix of jobs to various computers and measuring the time actually taken to perform them. Such an exercise is called a “benchmark” or a “benchmark test.” Computer users have often performed benchmarks for machine selection, and benchmark results for standardized or stylized data-processing problems have been published for many years: Phister (1979, p. 100) presents examples for mainframe computers that include filing and sorting problems, matrix inversion problems, and so forth.

An advantage of a benchmark measure is that it measures directly the speed, or cost, of jobs or applications, rather than of the instructions that are required for the job. However, representativeness requires selecting a group (possibly a large group) of alternative computer tasks and running each task on each machine in the sample. If problems are realistic, performing the tests may be expensive. A second problem arises when results of multiple benchmark tests are highly, but not perfectly, correlated (as they generally will be): The researcher must either select one benchmark as a proxy for all the rest or find some way to aggregate them—a problem exactly parallel to the use of single-instruction speed measures.

Weighted Instruction Mix Measures

A weighted instruction mix is formed by examining records of computer centers, or analysis either of "test packages" or of a sample of widely used programs. An internal "instruction trace" provides counts of the frequency of each machine instruction encountered in the programs. Execution speeds for each instruction can be timed (or obtained from published machine specifications). Weighting the speeds of the various instructions by the relative frequencies recorded in the instruction trace yields the weighted instruction mix.

The best-known weighted-instruction mix is MIPS (millions of instructions per second), which was used in a hedonic function for computers by Dulberger (1989). Lias (1980), Serlin (1986, p. 114), and Bell (1986)—and other computer manufacturers—emphasized that MIPS was designed to measure speed for IBM architectures, and might not provide a comparable measure across different machine architectures. Lias (1980, p. 105) put the measurement error that arises from applying MIPS to machines of different architectures at 10–30 percent. For this reason, Dulberger restricted her dataset to IBM and "plug-compatible" computers. Nevertheless, MIPS was fairly widely used across the industry in the 1980s and into the 1990s. Indeed, one often saw MIPS speed measures quoted for personal computers in the early 1990s (e.g., Rosch, 1994, Table 3.1).

Published documentation of the instruction mix used in MIPS is sketchy. Lias (1980) and others indicate that it was based on "IBM Job Mix 5," but the instructions in Job Mix 5 were not fully documented. One presumes the mix was updated from an earlier set of instructions known as the "Gibson mix," which Serlin (1986) dates around 1960. Serlin presents an example of the use of the Gibson mix to estimate processor speed. The only published documentation for the Gibson mix is an article in Japanese (Ishida, 1972).

Perhaps because published documentation of Job Mix 5 did not exist, some confusion has arisen about the nature of MIPS. To obtain a MIPS estimate with a smaller amount of work, approximating formulas were developed. For example, Bloch and Galage (1978) present an approximating formula that involves machine cycles and memory accesses per instruction combined with clock speed and memory access time. It is no doubt true that shortcuts were taken, and published empirical work may be affected by inaccuracies in the computer speed measure they employ.

Kenneth Knight (1966, 1970, 1985) published a weighted instruction mix speed measure designed for scientific purposes. His set of instructions included fixed-point addition, floating-point addition, multiplication, division, and finally, logic operations. Instruction frequencies for the scientific speed measure were derived from traces at a scientific computer center, with some arbitrary adjustments for aspects of the architecture of certain machines (for details, see Knight, 1985, Table 3, p. 117). For commercial uses, Knight collected the mix of instructions that were executed in a sample of commercial programs.

Knight's computing power formula combined the characteristics of speed and memory size into an index of processor "computing power." Some of the parameter values that were assumed in combining memory size with speed are arbitrary. Knight's updated indexes, which extend through 1979, retain the original 1963 weights.

Of the studies in Annex Table A that used an instructional mix measure of speed, only Cartwright, Donahoe, and Parker (1985) report a variable in the hedonic function other than speed and memory size. This evidence suggests that where other hardware attributes have been employed as variables in hedonic functions for mainframe computers, they were correcting in some sense for an inadequate measure of processor speed.

Synthetic Benchmarks

MIPS is sometimes termed a "synthetic benchmark." Others existed, even in mainframe days. The "Whetstone" reflected primarily scientific and engineering problems; the "Dhrystone" was based on systems-programming work, rather than numerical calculations. The "Linpack" measured solution speeds for systems of linear equations. Other special benchmarks existed for, e.g., banking transactions (Serlin, 1986, gives some representative results).

Since finding a satisfactory processor speed measure is the biggest challenge to measuring price and technological change in computer processors, one would have thought that economic researchers would have explored the usefulness of synthetic benchmarks. That did not happen until very recently (see section V).

Various studies investigated performance measures for peripheral equipment. Most of them were also measures of speed and capacity. The performance variables in the IBM price indexes are displayed in Cole et al. (1986). Flamm (1987) contains alternative indexes for peripheral equipment. The available price indexes for mainframe-era peripheral equipment are reviewed in Triplett (1989, Tables 4.11 and 4.12).

Summary

At the close of the era in which economic research on computer performance focused on mainframe computers, two generalizations characterized the state of that research.

First in their choice of speed measures, economists had turned away from simple clock speed and memory access times toward more representative speed measures, primarily synthetic benchmarks. Though there was some sense that future research might incorporate true benchmark measures, in fact that never happened. Indeed, as explained in the subsequent section, when research on computers picked up again in the 1990s, the advance represented by synthetic benchmarking measures was almost entirely abandoned, and economists turned back to simple clock speed as their primary measure of computer performance.

Second, in the mainframe research era, little or no attention was paid to system performance. Economists primarily modeled the performance of separate "boxes" of computer equipment, without paying very much attention to the integration of the equipment. The computer literature of the day was full of discussions of queuing theory and the implications of this for optimization of system performance (see, for example, Bard and Sauer, 1981). The economists' decisions were justified, in part, by their objectives: The boxes were separate pieces of output, typically produced and sold separately, and economists wanted to measure the output of boxes, adjusted for their performance. The computer center manager worried about the optimization problem; economists interested in output measurement and analysis did not have to be concerned about the computer center managers' problem.

In my 1989 survey (Triplett, 1989) I speculated that measures of system performance would show more rapid improvements over time than did the measures of separate boxes. It is true, of course, that technology has reduced the speed and cost of what was done yesterday. But computer users have gained most from technological changes that enable them to do things that were not possible with the previous technology, not merely from doing what they did before cheaper or faster. In my 1989 survey, I used as an example computer modeling solutions to an even then 150-year-old problem in aerodynamics, the solution to the set of Navier-Stokes equations that model the flow of a liquid or gas over a solid surface. In the intervening years, computer aerodynamics simulations have largely replaced wind tunnel tests, flight tests, and so forth. The computer permits new calculations. Its value is not just in doing the old ones faster.[6]

Extensions of computations (one might better say "manipulations of data") into new elements of the computational space are a major part of the contribution

---

[6] This is a common error in critiques of computer measurements. "I just type letters, the faster computer does not increase my typing speed in proportion, so computer performance measures overstate benefits to users." Perhaps for some users, this is true, but not for *uses* that take full advantage of computer capabilities.

of the computer that is not captured at all in existing measures of computer processor (and peripheral) speed and performance. Using even the best benchmarks to measure time series comparisons of computer equipment performance must inevitably measure the cost or speed of doing the jobs that were done yesterday in today's technology.

During the mainframe research era, the personal computer was more or less ignored, even though it was well established even in the late 1980s. As the next section shows, research on PCs did not extend the research on computer speed, and it largely ignored as well most of the problems of modeling system performance.

IV. PERFORMANCE MEASURES IN STUDIES OF PCs

In a sense, the engineering architecture of the PC is more closely aligned with the mainframe computer than is either its performance measurement or its economics. As explained in the following, economists who have modeled the PC's performance have followed, consciously or not, the performance measures used for separate computer equipment "boxes" in mainframe-era research.

A joint project between the IBM Corporation and the U.S. Bureau of Economic Analysis (BEA) developed hedonic computer equipment price indexes for the U.S. national accounts (Cole et al., 1986; Cartwright, 1986). These were the first hedonic computer price indexes introduced into any country's statistics.

The IBM-BEA price indexes covered four products: mainframe computers, disk drives, printers, and displays (terminals). The performance variables in the IBM studies have provided the basis for most subsequent investigations on computer equipment, including price indexes for personal computers (which were actually not included in the IBM-BEA work). The IBM hedonic functions for computer equipment continue, therefore, to provide guidance for empirical investigations of computer equipment today.

The IBM-BEA hedonic indexes were price indexes for computer equipment "boxes"; they controlled for quality change that arose as manufacturers increasingly put more performance into each of the separate boxes. No direct attention was paid to how the boxes—or properly, the characteristics of the separate boxes—were combined into an operating computer center, because an operating computer center was not purchased as a transaction. The buyer assembled a computer center; it was not produced and sold as a unit by the manufacturer.

The PC is, in effect, a pre-assembled computer center. The PC contains separate components that link nearly one-to-one to the individual "boxes" that were the subjects of Cole et al. (1986). For example, the PC's central processing unit (CPU), its hard drive and a display (keyboard/monitor) correspond to separate mainframe-era components. Most of this equipment can be purchased separately; indeed, these items may initially be manufactured by different manufacturers. It is not technologically linked together into a PC in the sense that components cannot be investigated separately. However, from the final purchaser's perspective, the PC transaction typically combines several of them (for example, a monitor and a keyboard are almost always included in the price). Only the printer remains as a separate piece of equipment that is typically purchased in a separate transaction. The transaction, more than the engineering, determines the unit that economists must analyze.

Thus, in modeling the PC one must ask a question that was never confronted in the IBM-BEA studies: Are we interested in the performance of the PC (that is, in the computer system)? Or that of its components? Of course, we are interested, ultimately, in both, for several reasons. But it will be important to keep the distinction between system performance and component performance in mind.

Annex Table B lists the performance variables used in PC hedonic functions. Three points can be made about existing PC hedonic functions.

First, compared with the IBM studies of separate computer equipment boxes, the PC studies omit some performance variables that were included in mainframe-era research (e.g., hard drive speed). Second, their processor speed measure (almost exclusively clock speed, measured in megahertz, MHz) is a step backward from the weighted instruction mix measure that was based, in principle, on the speed of performing jobs. Third, the PC hedonic studies measure the performance of components in the system—or simply the presence or absence of components, such as the video card—rather than the performance of the system, even though what the PC transaction researchers sought to model is the sale of a computer system, not the separate sale of computer components. These three points are developed in the following discussion.

Component Performance Measures

The variables used in three relevant IBM computer equipment hedonic functions are displayed in the first column of Annex Table B. The other columns of Table B summarize the variables used in a number of recent studies on personal computers, in comparison with the variables employed in the original IBM studies.

The studies in Annex Table B may not make up a complete review of research. They have been conducted in a number of different countries and show the degree of international comparability in hedonic research on personal computers. A number of them (mostly those at the right-hand side of the table) show either operational or experimental hedonic functions that have been estimated by statistical agencies in various countries, generally for the purpose of publishing performance-corrected computer indexes similar to those of the U.S.

Most of these studies have succeeded, to a perhaps surprising extent, in combining into one hedonic function many of the variables in three of the original IBM studies (Cole et al., 1986). None of the PC studies tries to combine printers (subject of a separate IBM study) in the PC hedonic function, because the acquisition of a printer still typically remains a separate transaction, even though the printer, too, is sometimes bundled with the rest of the PC. Barzyk (1999) does not include monitors, presumably because they were not bundled into the Canadian dataset used for his research. Dalen (1989) also excludes the monitor from his hedonic regression. With these exceptions, all the PC studies can be viewed as combing into a single hedonic function three of the separate pieces of equipment studied in the IBM work.

All studies measure processor performance with speed and memory size, as did Dulberger (1989), though nearly all of them measure speed with megahertz (a topic to which I return below). Dulberger introduced the idea of specifying the semiconductor type used in the processor (technology dummies). Most PC studies follow this innovation (BLS, Chwelos, Bourot, and in modified form, Moch and Finland CPI).

With respect to the hard disk, all studies use a standard measure of capacity. The difference between megabytes (MB) and gigabytes (GB) is merely a scaling, adopted for convenience because hard disk capacity has grown so large. Of the PC studies, only Dalen uses a hard drive speed variable, although the type dummy variables used by Barzyk (1999) and Bourot (1997) control to an extent for HD speed.

For monitors (displays) all the studies investigate measures of the quantity of information that can be shown on the screen and the resolution of the picture. Because some software producers have taken increasing amounts of the screen for control "bars" and so forth that are not readily hidden by the user, screen size may be an imperfect measure, but it clearly influences the price of the monitor. Other monitor characteristics include flat screen and the thickness of the monitor, which reflect users' desires that the machine occupy a smaller amount of desk space; those characteristics are omitted from existing studies.

In addition to the basic hardware items—processor, hard drive, and monitor/keyboard—a modern PC comes bundled with a number of other hardware features. Many of these are a consequence of the fact that the computer's function is increasingly not "computations" but the manipulation of digitized data, including sound and pictures. Sound cards, video cards, network cards, and so forth may be regarded as other pieces of hardware that are attached to the basic PC components, as are input/output devices such as CD-RW.

Most of the PC studies have included dummy variables for the presence or absence of at least some of these auxiliary functions or, alternatively, for more advanced versions of the functions in cases where some version of the feature has become nearly universal. Annex Table B-2 displays these other hardware features.

Little consensus has emerged among the users about which of these auxiliary hardware features should be added to the PC hedonic function. Additional research will be required to determine the reasons for the differences between the variables included in, for example, the BLS study and the others tabulated in

Table B: Do international hedonic functions for PCs differ because markets differ in the U.S. and other countries, or because of data availability differences, or because of different decisions made by the researchers? And even more importantly: How much difference in the computer price indexes computed by BLS and others results from differences in the variables in the hedonic function?

Performance Variables for PCs: The Dell Data

A recent Dell catalog illustrates the complexity of the bundle of PC computer characteristics and the inadequacy of the representation of the computer in most existing PC hedonic functions. This catalog illustrates how Dell markets computers to buyers. Most hedonic functions for personal computers do not contain nearly so many variables.

Dell advertises megahertz (now rescaled gigahertz, or GHz)**.** But it also advertises the speed of the bus and the size of the cache.[7] Memory size is there, but the specification page also talks about the speed of the memory, 266 MHz or 333 MHz SDRAM or RDRAM, which is a faster form of memory. The size of the hard drive is there, but so also is its speed. Specifications for the monitor and the DVD drive, not just their presence or absence, are included. Different cards are distinguished, the graphics card and the sound card, for example. How is that performance to be measured? Economists have left that out. Similarly, speaker performance must now be modeled in the PC bundle of characteristics, and audio specifications matter, a topic on which there is a minimal amount of hedonic research but nothing that has been applied to the computer bundle.

Then there is the software included in the Dell choices. This is a huge problem. A tiny amount of economic research exists on the performance of software, even though software in the national accounts is a larger component in the United States than is purchases of hardware (in the aggregate economy, though not necessarily for PCs).[8] Software actually included in a Dell machine is more extensive than what is mentioned in the catalog. Little or none of this software is modeled in PC hedonic functions.

There is also, of course, the warranty and the Internet access. Two years ago, Dell offered one-year “free” Internet access from *MSN.com*. Now it offers only six months “free,” but it gives a choice between AOL, MSN, and EarthLink. Is that an improvement or not?

Judged by the Dell catalog, existing research by economists on PCs is not very adequate in the way they model computer performance. The relevant ques-

---

[7] Some PC hedonic functions include information on the cache—see Table B-2.

[8] For software price/performance research by economists, see Harhoff and Moch (1997), Gandal (1994), Prud’homme and Yu (2002) and Levine (2002). Discussions of software measurement in the U.S. national accounts are contained in Parker and Grimm (2000) and Moylan (2001).

tion is: How much difference does it make? Do the omitted characteristics (such as hard drive speed, performance of cards, and quantity of software included in the sale) bias hedonic price indexes?[9] Does the "clock speed" measure of processor speed adequately measure improvements in processor performance?

Even aside from the adequacy of the variables in hedonic functions for PCs, there is another point: Do these variables measure the performance of the PC? Or do they measure the performance of inputs to PC performance?

V. BENCHMARK MEASURES OF PC PERFORMANCE

Ohta and Griliches (1976) introduced the distinction between what they called "physical characteristics," or engineering characteristics and "performance characteristics." In their language, processor megahertz and hard drive speed and size are "physical characteristics." The same distinction has also been discussed in the hedonic literature under the name "proxy variables": Physical characteris-

---

[9] Some economists will no doubt observe that the omitted characteristics may be correlated with the included ones. This requires what is essentially a digression.

Computer performance characteristics may, or may not, be highly correlated among themselves. This is a somewhat more complicated matter than has sometimes been supposed in some of the hedonic literature. As a factual point, correlations among the explanatory variables in the BLS PC hedonic function are in fact rather low (simple R's are almost entirely under 0.3 and some under 0.1, in a hedonic function with some 15 variables). Unless omitted characteristics have higher correlations than included characteristics have among themselves, we can presume that in the BLS model the influence of omitted characteristics on the estimated price premium for improved machines will be simply lost. Some clue to the importance of omitted variables is provided by examination of R^2: The BLS equation gets values on the order of 0.97, where the lowest value in the studies displayed in Annex Table B amounts is only about 0.5.

Inter-correlations among the characteristics of computer performance are higher in other datasets, for reasons that we do not need to explore here. One might presume, therefore, that omitted characteristics in these datasets will also have higher correlations with included characteristics. Most importantly, most of the other hedonic functions displayed in Annex Table B have far fewer characteristics than are used by BLS. This itself suggests that omitted variable problems (whether correlated or not) are far more serious in some computer hedonic studies than others. Differences in R^2 among studies may reflect properties of the data themselves, and not just the number of characteristics in the equation.

Nevertheless, even if omitted variables are correlated with an included variable, failing to consider them will bias measures of computer progress, and will bias the price index, *unless the omitted variable improves at the same rate as the included variable with which it is correlated*. This is not the place to explore what are essentially econometric problems in estimating hedonic functions. My own conclusion, from considerations that are developed elsewhere in the hedonic literature, is that omitted variable bias in hedonic price indexes and in measures of computer performance can be serious, and that omitted variables predominantly result in missing some of the improvement in computer performance, or what is the same thing, missing some of the decline in computer prices. For a similar conclusion on different grounds, see Nordhaus (2002). On the other hand, see the discussion of the work of Chwelos (2003) in section V.

tics have some relation to (are proxies for) the performance that buyers want from a PC, but they do not measure the performance that buyers really want. Variables in hedonic functions should represent what buyers buy (and sellers sell), not technical measures that have some relation or other to the true characteristics that are important for buyers and sellers behaviors.

Benchmark measures have the advantage that they measure machine performance, rather than measuring some proxy for machine performance, or some input that may influence machine performance. Researchers who have tried to incorporate benchmark data into PC hedonic functions are Chwelos (2003) and Barzyk (1999). Chwelos (2003) contains a good discussion of the relation between technical variables such as megahertz and benchmark performance measures.

From browsing e-sites, one would think that quite a number of benchmarks exist for PCs. The impression is illusory: It is a bit like the days when Sears sold Whirlpool appliances under its own name—a good many internet sites repackage benchmark tests from two companies: Veritest's Winstone and Bapco's SYSmark 2002 (revised from the 2001 version). Both of these perform separate benchmarks for performance on office productivity applications and graphics applications. For example, one task included in SYSmark's office productivity application is the time to execute a "replace all" command in Word, and there are a large number of tasks for which times are recorded. The overall score is the average time taken for the tasks included in the benchmark. For SYSmark, the final score is a geometric mean of scores on the two types of applications.

These benchmarks appear far better suited to economists' needs for performance measures than is megahertz, or clock speed. The benchmark is still subject to the problems listed above: Two or three applications benchmarks may not be representative of the range of applications that are important to users. For example, SYSmark 2002 contains only two applications, one for "office productivity" and another for "internet content." How one aggregates across users (in this context, the weights to be applied to the individual tasks within an application) is an issue. Another issue is the weights assigned to applications across users: SYSmark effectively assumes the two applications have equal weight. Winstone remains agnostic on the matter, leaving aggregation over applications included in the benchmark to the user of the benchmark.

More importantly, data on benchmarks may not always be available to economists and are not necessarily consistent over time.[10] The Butler Group (2001) commented that "attempts to market processors using something other than their clock speed have found limited success. . . . Consumers are used to dealing with the seemingly easy to compare clock speed, even though this may not be the greatest performance indicator it has been the only one available." Economists,

---

[10] Chwelos (2003) used overlaps to estimate comparable points to create time series of benchmarks.

too, have had to measure performance with the only data available, which has been megahertz.

One also needs to distinguish a benchmark for the speed of the microprocessor chip, which is the focus of a considerable amount of recent interest in benchmark measures,[11] from a benchmark for the system as a whole.

In the end, the key question is: How much does an inferior or proxy measure of speed matter? Chwelos (2003) compared usual proxy measures of speed (megahertz and so forth) with benchmark tests from PC magazine. He found that the relation between megahertz and performance differs across microprocessor generations.[12] However, the price indexes he estimated differed trivially: An index using benchmarks declined 39.6 percent per year, and one using technical specifications declined 39.3 percent annually, where the indexes used otherwise comparable computational forms (see his Table XII). One reason for this result is that Chwelos' technical specification was unusually rich: It included measures for cache memory, dummy variables for chip generation, and so forth.[13] The same result might not apply to the simpler hedonic models employed in most of the other studies tabulated in Annex Table B.

Yet the results are provocative. Simple measures of processor speed may not be that inadequate, empirically, though they seem inadequate, a priori.

VI. INPUT COMPONENTS, THEIR PERFORMANCE MEASURES AND THEIR CONTRIBUTIONS TO SYSTEM PERFORMANCE

As discussed in sections IV and V, it is not entirely clear whether the performance variables in PC hedonic functions measure the performance of the PC, the performance of inputs to the PC system, or some of both (or neither). However one addresses these questions, a modern PC incorporates many hardware and software components. The performance of many of these components is not measured at all in PC hedonic functions.

Putting these matters aside as unresolved for the present, this final section addresses another important topic: To what extent can the price/performance of the computer be explained by technical changes in the computer's components? This shifts attention from measuring the computer's performance to explaining it. The problem itself and its economic importance are well stated in Jorgenson (2001): Understanding recent economic growth in the U.S. and forecasting future growth requires a better understanding of the sources of technical progress in ICT equipment.

---

[11] For example, AMD released a report by PricewaterhouseCooper listing a variety of benchmark tests on AMD microprocessor products (*www.amd.com/us-en/.htm*).

[12] ". . . A MHz of clock speed from a 286 processor produces less performance that a MHz from a 386, a 386 less than a 486, and so on." Chwelos (2003, p. 14).

[13] Chwelos contains an excellent discussion of the architecture of the PC and the implications of the architecture for measuring price/performance.

Two strands of economic research have approached this problem. One can start from the semiconductor, clearly a major input into the computer itself and into many of its components. This approach is represented by Flamm (1997). Flamm calculates the impact of the semiconductor on the economy, which perforce involves both the impact of the semiconductor on computers and the consequent impact of computers on the economy (as well as the impact of semiconductors on other, non-computer, sectors).

Another approach is to model the price/performance of the computer as an outcome of the price/performance of its inputs. Triplett (1996) decomposes multifactor productivity in computers into three parts: productivity in computer manufacturing, productivity in semiconductor manufacturing (which provides a major input to computers), and semiconductor manufacturing equipment, treated as providing technological inputs (with computers) to the manufacture of semiconductors. This approach builds on the technical change model in Triplett (1985): The production of a technological output that uses technological inputs is treated as the production of a set of output characteristics with a set of intermediate input characteristics. This implies a transformation (production) function of the form:

$$0 = T\,(M_i,\ \ldots,\ M_h;\ V_{11},\ldots V_{1m},\ \ldots.\ V_{r1},\ldots V_{rn},\ Z),$$

where M_i is a characteristic of computer performance (and there are h characteristics of computer performance), and V_{jk} is the kth characteristic of input j (which might be the semiconductor, or the network card, and so forth), and Z is a group of other inputs, which may or may not be homogenous inputs, but whose characteristics are ignored for simplicity.[14] This model is a very general way of modeling what most economists would describe in a much more structured way: Quality-adjusted output is a function of quality-adjusted inputs.[15] Either structured or unstructured ways of approaching the problem require information on the price/performance of the output of the computer production process and the price/performance of the inputs that are incorporated into the computer.

In Triplett (1996) the specification of component inputs to the production of computers was very crude: Semiconductor price indexes that embodied price/

---

[14] "Labor quality," for example, is composed of elements of human capital, with characteristics such as education, training, experience and so forth. This is formally identical to the treatment of technological inputs through hedonic functions.

[15] The advantage of the unstructured way of approaching the problem is that it avoids having to specify whether a characteristic of the video card (say) enters into the performance of the computer in some manner that is independent of the characteristics of some other component, say processor speed or hard disk capacity. The unstructured way is more appropriate in the sense that it can more readily incorporate the engineering knowledge on the relations among characteristics of different components. In practice, however, the information requirements are daunting, and most economists therefore pursue the structured approach, even though it represents implicit and unrealistic specifications of engineering relations.

performance of semiconductors had only recently become available, following the pioneering research of Flamm (1993) and Dulberger (1993). Because only an aggregate price index for semiconductors was available, semiconductors were treated as an aggregate good, even though different types of semiconductors go into computers and they do not always have identical rates of performance improvement.

More recently, Aizcorbe, Flamm, and Khurshid (2002) have produced detailed price indexes for 12 classes of semiconductors and also produced consumption weights among these 12 classes for end uses, including computer production. This is a great step forward in data that can be used for modeling the contributions of component inputs to the performance of computers.

Missing, however, are comparable performance indicators for other component inputs to computer performance. The performance of hard drives does not appear to be driven primarily by advances in electronics, though it may be true that the miniaturization technology that is evident in hard drive technology is similar to, or is driven by, the miniaturization technology that underlies advances in semiconductor performance. Little information on the performance of graphics and networking cards, or CD/DVD read/write, or monitors, or other components has so far been brought to bear on economic modeling computer price/performance. Even storage media themselves appear to have undergone price/performance changes that rival those of hardware.

VII. CONCLUSIONS

Much is still unknown about the contributions of component technologies to the increase in computer performance. Earlier, I noted that the cost of computing power is now around 1/1,000 of 1 percent of what it cost 50 years ago. That estimate, breathtaking as it is, actually does not incorporate all of the aspects of computer performance that one might consider. Even so, an exciting research agenda is to account for the determinants of the great decline in computer price/performance over the last 50 years: Part of it must be technological innovations in computer components, and to quantify those contributions, we need better performance measures of components**.** The Deconstructing the Computer workshop is a step along the route of modeling those determinants.

REFERENCES

Aizcorbe, Ana, Kenneth Flamm, and Anjum Khurshid. 2002. “The Role of Semiconductor Inputs in IT Hardware Price Decline: Computers vs. Communications.” Federal Reserve Board Finance and Economics Series Discussion Paper 2002-37. August.

Archibald, Robert B., and William S. Reece. 1979. “Partial Subindexes of Input Prices: The Case of Computer Services.” *Southern Economic* 46(October):528–540.

Bapco. 2002. "SYSmark ® 2002: An Overview of SYSmark 2002 Business Applications Performance Corporation." Available at *http://www.bapco.com/SYSmark2002Methodology.pdf*. Accessed February 19, 2003.

Bard, Yonathan, and Charles H. Sauer. 1981. "IBM Contributions to Computer Performance Modeling." *IBM Journal of Research and Development* 25:562–570.

Barzyk, Fred. 1999. "Updating the Hedonic Equations for the Price of Computers." Working Paper of Statistics Canada, Prices Division. November 2.

Bell, C. Gordon. 1986. "RISC: Back to the Future?" *Datamation* 32(June): 96–108.

Berndt, Ernst R., and Zvi Griliches. 1993. "Price Indexes for Microcomputers: An Exploratory Study." In Murray F. Foss, Marilyn Manser, and Allan H. Young, eds. *Price Measurements and Their Uses.* Studies in Income and Wealth 57:63–93. Chicago: University of Chicago Press for the National Bureau of Economic Research.

Berndt, Ernst R., Zvi Griliches, and Neal Rappaport. 1995. "Econometric Estimates of Prices in Indexes for Personal Computers in the 1990s." *Journal of Econometrics* 68(1995):243–268.

Berndt, Ernst R., and Neal J. Rappaport. 2001. "Price and Quality of Desktop and Mobile Personal Computers: A Quarter-Century Historical Overview." *American Economic Review* 91(2):268–273.

Berndt, Ernst R., and Neal J. Rappaport. 2002. "Hedonics for Personal Computers: A Reexamination of Selected Econometric Issues." Unpublished paper.

Bloch, Erich, and Dom Galage. 1978. "Component Progress: Its Effect on High-Speed Computer Architecture and Machine Organization." *Computer* 11(April):64–75.

Bourot, Laurent. 1997. "Indice de Prix des Micro-ordinateurs et des Imprimantes: Bilan d'une rénovation." Working Paper of the Institut National De La Statistique Et Des Etudes Economiques (INSEE). Paris, France, March 12.

Bureau of Economic Analysis. 2001. "A Guide to the NIPAs." In *National Income and Product Accounts of the United States, 1929–97.* Washington, D.C.: Government Printing Office. Also available at *http://www.bea.doc.gov/bea/an/nipaguid.pdf*.

Butler Group. 2001. "Is Clock Speed the Best Gauge for Processor Performance?" *Server World Magazine* September. Available at *http://www.serverworldmagazine.com/opinionw/ 2001/09/06_clockspeed.shtml*. Accessed February 7, 2003.

Cale, E.G., L.L. Gremillion, and J.L. McKenney. 1979. "Price/Performance Patterns of U.S. Computer Systems." *Communications of the Association for Computing Machinery (ACM)* 22 (April):225–233.

Cartwright, David W. 1986. "Improved Deflation of Purchases of Computers." *Survey of Current Business* 66(3):7–9.

Cartwright, David W., Gerald F. Donahoe, and Robert P. Parker. 1985. "Improved Deflation of Computer in the Gross National Product of the United States." Bureau of Economic Analysis Working Paper 4. Washington, D.C.: U.S. Department of Commerce.

Chow, Gregory C. 1967. "Technological Change and the Demand for Computers." *American Economic Review* 57(December):1117–1130.

Chwelos, Paul. 2003. "Approaches to Performance Measurement in Hedonic Analysis: Price Indexes for Laptop Computers in the 1990s." *Economics of Innovation and New Technology* 12(3):199–224.

Cole, Rosanne, Y.C. Chen, Joan A. Barquin-Stolleman, Ellen Dulberger, Nurhan Helvacian, and James H. Hodge. 1986. "Quality-Adjusted Price Indexes for Computer Processors and Selected Peripheral Equipment." *Survey of Current Business* 66(1):41–50.

Colecchia, Alessandra, and Paul Schreyer. 2002. "ICT Investment and Economic Growth in the 1990s: Is the United States a Unique Case? A Comparative Study of Nine OECD Countries." *Review of Economic Dynamics* 5(2):408–442.

Dalén, Jorgen. 1989. "Using Hedonic Regression for Computer Equipment in the Producer Price Index." R&D Report, Statistics Sweden, Research-Methods-Development, Vol. 25.

Dulberger, Ellen R. 1989. "The Application of a Hedonic Model to a Quality Adjusted Price Index for Computer Processors." In Dale W. Jorgenson and Ralph Landau, eds. *Technology and Capital Formation*, pp. 37–75. Cambridge: MIT Press.

Dulberger, Ellen. 1993. "Sources of Price Decline in Computer Processors: Selected Electronic Components." In Murray Foss, Marilyn Manser, and Allan Young, eds. *Price Measurements and Their Uses.* Chicago: University of Chicago Press for the National Bureau of Economic Research.

Ein-Dor, Phillip. 1985. "Grosh's Law Re-visited: CPU Power and the Cost of Computation." *Communications of the Association for Computing Machinery (ACM)* 28(February):142–151.

Evans, Richard. 2002. "INSEE's Adoption of Market Intelligence Data for Its Hedonic Computer Manufacturing Price Index." Presented at the Symposium on Hedonics at Statistics Netherlands, October 25.

Fisher, Franklin M., John J. McGowan, and Joen E. Greenwood. 1983. *Folded, Spindled, and Multiplied: Economic Analysis and U.S. v. IBM.* Cambridge, MA: MIT Press.

Flamm, Kenneth. 1987. *Targeting the Computer.* Washington, D.C.: The Brookings Institution.

Flamm, Kenneth. 1988. *Creating the Computer: Government, Industry, and High Technology.* Washington, D.C.: The Brookings Institution.

Flamm, Kenneth. 1993. "Measurement of DRAM Prices: Technology and Market Structure." In Murray Foss, Marilyn Manser, and Allan Young, eds. *Price Measurements and Their Uses.* Chicago: University of Chicago Press for the National Bureau of Economic Research.

Flamm, Kenneth. 1997. *More for Less: The Economic Impact of Semiconductors.* San Jose, CA: Semiconductor Industry Association.

Gandal, Neil. 1994. "Hedonic Price Indexes for Spreadsheets and an Empirical Test for Network Externalities." *RAND Journal of Economics* 25.

Gordon, Robert J. 1989. "The Postwar Evolution of Computer Prices." In Dale W. Jorgenson and Ralph Landau, eds. *Technology and Capital Formation*, pp. 37–75. Cambridge: MIT Press.

Harhoff, Dietmar, and Dietmar Moch. 1997. "Price Indexes for PC Database Software and the Value of Code Compatibility." *Research Policy* 24(4-5):509–520.

Holdway, Michael. 2001. "Quality-Adjusting Computer Prices in the Producer Price Index: An Overview." Bureau of Labor Statistics, October 16.

Ishida, Haruhisa. 1972. "On the Origin of the Gibson Mix." *Journal of the Information Processing Society of Japan* 13(May):333–334 (in Japanese).

Jorgenson, Dale W. 2001. "Information Technology and the U.S. Economy." *American Economic Review*, 91(1):1–32.

Jorgenson, Dale W., Mun S. Ho, and Kevin J. Stiroh. 2002. "Information Technology, Education, and the Sources of Economic Growth Across U.S. Industries." Presented at the Brookings Workshop "Services Industry Productivity: New Estimates and New Problems," March 14. Available at *http://www.brook.edu/dybdocroot/es/research/projects/productivity/workshops/20020517.htm*.

Kelejian, Harry H., and Robert V. Nicoletti. c. 1971. "The Rental Price of Computers: An Attribute Approach." Unpublished paper, New York University (no date).

Knight, Kenneth E. 1966. "Changes in Computer Performance: A Historical View." *Datamation* (September):40–54.

Knight, Kenneth E. 1970. "Application of Technological Forecasting to the Computer Industry." In James R. Bright and Milton E.F. Schieman, eds. *A Guide to Practical Technological Forecasting*. Englewood Cliffs, NJ: Prentice-Hall.

Knight, Kenneth E. 1985. "A Functional and Structural Measure of Technology." *Technological Forecasting and Technical Change* 27(May):107–127.

Koskimäki, Timo, and Yrjö Vartia. 2001. "Beyond Matched Pairs and Griliches-Type Hedonic Methods for Controlling Quality Changes in CPI Sub-indices." Presented at Sixth Meeting of the International Working Group on Price Indices, sponsored by the Australian Bureau of Statistics, April.

Levine, Jordan. 2002. "U.S. Producer Price Index for Pre-Packaged Software." Presented at the 17th Voorburg Group Meeting, Nantes, France, September.

Levy, David, and Steve Welzer. 1985. "An Unintended Consequence of Antitrust Policy: The Effect of the IBM Suit on Pricing Policy." Unpublished paper, Rutgers University Department of Economics, December.

Lias, Edward. 1980. "Tacking the Elusive KOPS." *Datamation* (November):99–118.

Lim, Poh Ping, and Richard McKenzie. 2002. "Hedonic Price Analysis for Personal Computers in Australia: An Alternative Approach to Quality Adjustments in the Australian Price Indexes."

Michaels, Robert. 1979. "Hedonic Prices and the Structure of the Digital Computer Industry." *The Journal of Industrial Economics* 27(March):263–275.

Moch, Dietmar. 2001. "Price Indices for Information and Communication Technology Industries: An Application to the German PC Market." Center for European Economic Research (ZEW) Discussion Paper No. 01-20, Mannheim, Germany, August.

Moylan, Carol, 2001. "Estimation of Software in the U.S. National Income and Product Accounts: New Developments." OECD Paper. September. Available at *http://webnet1.oecd.org/doc/M00017000/M00017821.doc*.

Nelson, R. A., T. L. Tanguay, and C. C. Patterson. 1994. "A Quality-adjusted Price Index for Personal Computers." *Journal of Business and Economics Statistics* 12(1):23–31.

Nordhaus, William D. 2002. "The Progress of Computing." Yale University, March 4.

Ohta, Makoto, and Zvi Griliches. 1976. "Automobile Prices Revisited: Extensions of the Hedonic Hypothesis." In Nestor E. Terleckyj, ed. *Household Production and Consumption.* Conference on Research in Income and Wealth, *Studies in Income and Wealth* 40:325–90. New York: National Bureau of Economic Research.

Okamoto, Masato, and Tomohiko Sato. 2001. "Comparison of Hedonic Method and Matched Models Method Using Scanner Data: The Case of PCs, TVs and Digital Cameras." Presented at the Sixth Meeting of the International Working Group on Price Indices, sponsored by the Australian Bureau of Statistics, April.

Pakes, Ariel. 2001. "A Reconsideration of Hedonic Price Indices with an Application to PCs." Harvard University, November.

Parker, Robert P., and Bruce Grimm. 2000. "Recognition of Business and Government Expenditures for Software as Investment: Methodology and Quantitative Impacts, 1959–98." Paper presented to BEA's Advisory Committee, May 5. *http://www.bea.doc.gov/bea/papers/software.pdf.*

PC World Magazine. 2003. "20 Years of Hardware." March.

Patrick, James M. 1969. "Computer Cost/Effectiveness." Unpublished paper summarized in Sharpe (1969, p. 352).

Phister, Montgomery. 1979. *Data Processing Technology and Economics,* Second Edition. Bedford, MA: Santa Monica Company Publishing and Digital Press.

Prud'homme, Marc, and Kam Yu. 2002. "A Price Index for Computer Software Using Scanner Data." Unpublished working paper, Prices Division, Statistics Canada.

Rao, H. Raghaw, and Brian D. Lynch. 1993. "Hedonic Price Analysis of Workstation Attributes." *Communications of the Association for Computing Machinery (ACM)* 36(12):94–103.

Ratchford, Brian T., and Gary T. Ford. 1976. "A Study of Prices and Market Shares in the Computer Mainframe Industry." *The Journal of Business* 49:194–218.

Ratchford, Brian T., and Gary T. Ford. 1979. "Reply." *The Journal of Business* 52:125–134.

Rosch, Winn L. 1994. *The Winn L. Rosch Hardware Bible.* Indianapolis: Sams Publishing.

Serlin, Omri. 1986. "MIPS, Dhrystones, and Other Tables." *Datamation* 32(June 1):112–118.

Sharpe, William F. 1969. *The Economics of the Computer.* New York and London: Columbia University Press.

Statistics Finland. 2000. "Measuring the Price Development of Personal Computers in the Consumer Price Index." Paper for the Meeting of the International Hedonic Price Indexes Project. Paris, France, September 27.

Stoneman, Paul. 1976. *Technological Diffusion and the Computer Revolution: The U.K. Experience.* Cambridge: Cambridge University Press.

Stoneman, Paul. 1978. "Merger and Technological Progressiveness: The Case of the British Computer Industry." *Applied Economics* 10:125–140. Reprinted as chapter 9 in Keith Cowling, Paul Stoneman, John Cubbin, John Cable, Graham Hall, Simon Domberger, and Patricia Dutton, *Mergers and Economic Performance*. Cambridge: Cambridge University Press (1980).

Triplett, Jack E. 1985. "Measuring Technological Change with Characteristics-Space Techniques." *Technological Forecasting and Social Change* 27:283–307.

Triplett, Jack E. 1989. "Price and Technological Change in a Capital Good: A Survey of Research On Computers." In *Technology and Capital Formation*, Dale W. Jorgenson and Ralph Landau, eds. Cambridge, MA: MIT Press.

Triplett, Jack E. 1996. "High-Tech Industry Productivity and Hedonic Price Indices." In *OECD Proceedings: Industry Productivity, International Comparison and Measurement Issues*. Paris: Organisation for Economic Co-operation and Development.

Triplett, Jack E., and Barry Bosworth. 2002. "Baumol's Disease Has Been Cured: IT and Multifactor Productivity in U.S. Service Industries." Presented at the Brookings Workshop "Services Industry Productivity: New Estimates and New Problems," March 14. Available at *http://www.brook.edu/dybdocroot/es/research/projects/productivity/workshops/20020517.htm*.

van Mulligen, Peter Hein. 2002. "Alternative Price Indices for Computers in the Netherlands Using Scanner Data." Prepared for the 27th General Conference of the International Association for Research in Income and Wealth, Djurhamn, Sweden.

VeriTest. 2003. "Business Winstone ™ 2002 Basics." Available at *http://www.veritest.com/benchmarks/bwinstone/wshome.asp*. Accessed February 19, 2003.

Wallace, William E. 1985. "Industrial Policies and the Computer Industry." The Futures Group Working Paper #007. Glastonbury, CT: The Futures Group.

Wyckoff, Andrew W. 1995. "The Impact of Computer Prices on International Comparisons of Labour Productivity." *Economics of Innovation and New Technology* 3:277–293.

ANNEX A Comparison of Variables in Hedonic Functions for Computers

Author	Data Sources	Dependent Variable	Explanatory Variables
Knight (1966, 1970, 1985)	Price: "published" rental prices. Independent variables: own, plus "published" specifications	Monthly rental for "most typical" configuration	1. Computing "power" (operations per second) $C = \frac{10^{12}(M(L-7)Wk)^{a}}{t_1 + t_2} = \frac{10^{12}}{t_1 + t_2} x \ (\text{"}\mathit{memory\ factor}\text{"})$ 2. s = monthly seconds per dollar of monthly rental Definitions: M = memory size (in words) L = word length (in bits) W = "word factor" (dummy variable for memory types) k = scaling constant t_1 = time to perform 1 million operations (in microseconds) t_2 = I/O or other idle time for 1 million operations (in microseconds) a = .05 for scientific, .33 for commercial NB: t_1 and t_2 were calculated (from computer specifications and computer center operations data) as weighted average of five categories of computations.
Chow (1967)	Special government survey; *Computers and Automation*; and IBM	Average monthly rental for specific configurations of computers newly introduced in year *t*	1. Multiplication time (in microseconds) 2. Memory size (words × word length) 3. Memory access time ("average time required to retrieve information from the memory") NB: Also tried addition time, rejected for multicollinearity with multiplication time ("a slightly inferior variable"). Notes other omitted hardware characteristics, which he assumes correlated with included characteristics.
Schneidewind (Sharpe, 1969)	Not specified	Monthly rental	1. Memory size (thousands of characters) 2. Memory cycles per second (words)
Skattum (Sharpe, 1969)	Not specified	Monthly rental	Same as Schneidewind

continued

ANNEX A Continued

Author	Data Sources	Dependent Variable	Explanatory Variables
Early, Barro, and Margolis (Sharpe, 1969)	Not specified	Monthly rental	1. Memory size (in bits) 2. Memory cycles per second (presumably in bits) 3. Several others (including additions per second), which were not significant
Patrick (Sharpe, 1969)	*Computer Characteristics Quarterly*; *Computers and Automation*	Monthly rental for “typical” configuration, second-generation computers	1. Space occupied (in square feet) 2. Additions per second (in thousands) 3. Minimum memory (in bits) 4. Maximum memory (in bits) 5. IBM dummy 6. Number of months since first installation 7. Number of machines installed since introduction
Jacob (Sharpe, 1969)	As Patrick	As Patrick, third-generation computers	1. Additions per second (in thousands) 2. Minimum memory (in thousands of bits) 3. Maximum memory (in thousands of bits) 4. Memory cycles per second (thousands of bits) 5. Number of operations codes 6. IBM dummy 7. Number of months since first installation 8. Number of machines installed since introduction
Kelejian and Nicoletti (1971)	*Computers and Automation*; *Computer Characteristics Quarterly*	Minimum monthly rental	1. Add time (in microseconds) 2. Storage cycle time (in microseconds) 3. Minimum memory size (thousands of bits)
Stoneman (1976)	*British Commercial Computer Digest*; *Computers and Automation*; other	Published average price, all installations, all years machine is sold	1. Cycle time in microseconds 2. Maximum storage in thousands of bits 3. Floor area in square feet 4. Year and “generation” dummies NB: Final set of variables selected from a much larger original set by comparing adjusted R^2 for groupings of the original variables. Author comments that owing to multicollinearity floor area proxies for speed.

ANNEX A Continued

Author	Data Sources	Dependent Variable	Explanatory Variables
Ratchford and Ford (1976, 1979)	Auerbach Corp. (two sources), cross sections for 1964, 1967, and 1971	Average monthly rental, computer systems (CPU plus peripherals)	1. Memory size (maximum words in storage available with particular CPU) 2. Add time (in microseconds) 3. Dummies for age of machine and manufacturer NB: 36 variables tested with factor analysis; however, regression based on four variables mentioned by Chow, with two retained.
Stoneman (1978)	*British Commercial Computer Digest*; *Computers and Automation*; other	Prices of newly introduced machines	1. Cycle time in microseconds 2. Maximum storage in thousands of bits 3. Dummies for year of introduction
Archibald and Reece (1979)	*Computer Price Guide*; characteristics from various published sources	Asking price for used IBM machines of specified configuration	1. Add time (in microseconds) 2. Memory size (bits) in configuration 3. Cycle (read) time (in microseconds) 4. Access time (in milliseconds) 5. Number of time share features 6. Number of CPU “intensiveness” features 7. Printer speed (hundreds of lines per minute) 8. Card reader speed (hundreds of cards per minute) 9. Several other characteristics of peripherals NB: Got “incorrect” signs for major variables, which often happens with multicollinearity and many variables
Michaels (1979)	Auerbach Corp. (same as Ratchford and Ford: 264 “configurations” of CPU and peripherals, as of July 1971)	“Basic” monthly rental for specified configuration	1. Add time (in microseconds) 2. Index, memory core size (thousands of bytes), and transfer speed within core (in kilobytes per second) 3. Index, card reader speed, and card punch speed 4. Index, number of tape drives, and maximum read-write speed 5. Storage capacity (millions of bytes) in configuration

continued

ANNEX A Continued

Author	Data Sources	Dependent Variable	Explanatory Variables
			6. Dummies for manufacturer and introduction year (gives price index, relative to earliest machines) NB: Justification for forming indexes based on technical assumptions—e.g., number of tape drive substitutes for speed in achieving same results.
Cale, Gremillion, and McKenney (1979)	Datapro	Price at introduction for a "balanced" system (processor plus peripherals)	1. Memory size in bytes 2. Size (in megabytes) of online direct access storage NB: Addition time and other unspecified speed measures insignificant, partly owing to multicollinearity
Fisher, McGowan and Greenwood (1983)	Government lease price lists	Lease prices to federal government	1. Memory size in thousands of bits 2. Addition time (including access time) 3. Transfer rate (bytes per second)
Wallace (1985)	GML Corp.; International Data Corp; Phister (1979)	List prices of all machines	1. Linear combination of MIPS and KOPS 2. Memory size included or minimum memory size (units not given) 3. Dummy variables for computer size class 4. Dummy variables for manufacturers
Cartwright, Donohoe, and Parker (1985)	Auerbach Corp.; Datapro Corp.; and *Computerworld*	List prices, all machines available	1. Speed (memory cycle time, machine cycle time, or MIPS, depending on period) 2. Memory size (in megabytes) 3. Maximum number of channels
Levy and Welzer (1985)	*Computerworld*	Published (list) prices, all machines from major producers	1. MIPS 2. Average memory size 3. Dummy variables for manufacturer, and for newly introduced
Ein-Dor (1985)	*Computerworld*; other sources	List price, selection of 106 machines	1. MIPS (a number of other performance measures were related to MIPS and to "average computational cost")
Flamm (1987)	Phister (1979)	List price, all machines in source	1. KOPS × 10^{-3} 2. Memory size in megabytes

ANNEX A Continued

Author	Data Sources	Dependent Variable	Explanatory Variables
Gordon (1989) 1954–1979 *regressions*	Phister (1979)	Prices of newly introduced machines	1. Memory cycle time (in microseconds) 2. Memory size (in megabytes) 3. IBM dummy
1977–1984 *regressions*	*Computerworld*	Prices of all machines	1. Machine cycle time (in nanoseconds 2. Memory size (in megabytes) 3. Minimum number of channels 4. Maximum number of channels 5. Cache buffer size (units not given)
Dulberger (this volume)	*Datamation*; *Computerworld*; IBM	List price, IBM and "plug-compatible" machines	1. MIPS 2. Memory size (in megabytes)—maximum and minimum 3. "Technology class" dummy variables NB: Each machine entered twice in the data set, once with maximum memory size available, once with minimum memory size, with the appropriate price for each.

ANNEX B1 Variables in Computer Hedonic Functions, Hardware Components Only

	Cole et al. 1986	Berndt and Griliches 1993	Berndt et al. 1995 (desktops)	Berndt and Rappaport 2002[a]	Chwelos 2003 (laptops)
Processor (CPU)					
Speed	MIPS	MHz	MHz	MHz	MHz * CPU or benchmark scores
Memory	MB (min and max)	KB	KB (installed and maximum)	MB	MB
Cache	no	no	no	no	no
Technology variables	Chip dummies	16- or 32-bit processor chip dummies	8-, 16- or 32-bit processor chip dummies	Processor type; processor type*MHz	Intel dummy
Disk (hard) drive					
Capacity	MB	MB	MB	MB	MB
Speed	Sum of 3	no	no	no	no
Other	no	no	Dummy for no HD	no	no
Displays (terminals, monitors, and keyboards)					
Screen size	Number of characters	no	no	no	Size
Resolution	Dpi	no	no	no	Pixels in maximum resolution
Color	Number	Dummy	no	no	Dummy
Other	Number of function keys	no	no	no	Active or passive matrix LCD dummies
Other hardware features (if yes, see Annex B2)	no	7	6	2	8
Software features (if yes, see Annex B2)	no	no	no	no	no

Nelson et al. 1994	Pakes 2001	Moch 2001 (Germany)	Rao and Lynch 1993 (workstations)	Holdway 2001 (U.S.)	Bourot 1997 (INSEE)
MHz	MHz, MHz^2	Test score	MIPS	MHz	MHz
MB	MB	MB	KB	MB	MB[3]
no	no	KB	no	no	KB
Processor type	Maximum memory; Apple*speed	Architecture dummy	no	Celeron dummy	Chip dummies
MB	GB	MB	MB	MB	MB
no	no	no	no	no	no
no	no	no	no	no	Type dummies
no	no	Size	no	Size dummies	Size
no	no	no	no	Trinitron dummy	dpi
Dummy	no	Dummy	no	no	no
Monochrome monitor dummy	no		Monochrome monitor dummy	no	no
5	7	6	3	9	7
2	no	yes	no	3	yes

continued

ANNEX B1 Continued

	Evans 2002		Barzyk 1999 (StatCan)	Dalen 1989 (Sweden)	Koshimaki and Vartia 2001
	INSEE01	INSEE02			
Processor (CPU)					
Speed	MHz	[b]	Test score	MHz	MHz
Memory	MB	MB	MB	MB	MB
Cache	no	max	KB	no	no
Technology	Memory type; maximum memory	no	no	no	no
Disk (hard) drive					
Capacity	GB	[b]	MB	MB	no
Speed	no	no	no	Access time	no
Other	no	no	Type dummies	no	no
Displays (terminals, monitors, and keyboards)					
Screen size	no	no	no	no	no
Resolution	no	no	no	no	no
Color	no	no	no	no	no
Other	no	no	no	no	no
Other hardware features (if yes, see Annex B2)	7	4	6	no	no
Software features (if yes, see Annex B2)	no	no	no	no	no

---

[a] Includes the same variables as Berndt and Rappaport (2001) plus microprocessor-type dummy variables and interactions between microprocessor type and clock speed.

[b] Replaced by external volume measure.

Statistics Finland 2000	Okamoto and Sato 2001	Lim and McKenzie 2002	van Mulligen 2002
MHz	MHz	CPU score	MHz
MB	MB	MB	MB
no	no	KB	no
Type dummy	Processor type	no	Processor type
GB	MB	MB	GB
no	no	no	no
no	no	no	no
Size	Size	17″ dummy	no
no	no	no	no
no	no	no	no
no	No monitor dummy; LCD dummy	no	Dummy variable for presence
no	4	9	3
no	no	no	no

ANNEX B2 Computer Hedonic Functions, Other Hardware and Software Features (for other variables and sources, see Annex B1)

	Berndt and Griliches	Berndt et al.	Berndt and Rappaport	Chwelos	Nelson et al.
ZIP dummy	no	no	no	no	no
CDROM dummy	no	no	yes	no	no
CDROM speed	no	no	no	yes	no
CDRW dummy	no	no	no	no	no
DVD dummy	no	no	no	no	no
Sound card dummy	no	no	no	no	no
Video (MB)	no	no	no	no	no
Network card	no	no	no	no	no
Modem dummy	no	no	no	modem speed	no
Speakers dummy	no	no	no	no	no
Case type dummy	no	no	no	no	no
Warranty dummy	no	no	no	no	no
Seller dummies	yes	yes	major brand	major brand	yes
SCSI control	no	no	no	no	no
Operating system	no	no	no	no	yes
Other software	no	no	no	no	other software utilities
Other	number of floppy drives slots available for expansion board mobile dummy discounted by vendor age extra hardware	two or more floppy drives dummy size weight density age		battery type battery life index density discount price weight	number of floppy drives extended industry standard architecture bus number of slots number of ports

Pakes	Moch	Rao and Lynch	Holdway	Bourot
no	no	no	yes	no
no	yes	no	no	yes
no	no	no	no	yes
yes	no	no	no	no
yes	no	no	yes	no
yes	no	no	no	yes
yes	yes	no	yes	yes
yes	no	no	yes	no
yes	no	no	yes	yes
no	no	no	yes	no
no	yes	no	no	yes
no	no	no	yes	no
Apple	no	yes	yes	no
no	no	yes	no	no
no	yes	no	yes	no
no	number of bundled applications	no	software office suite; MS Office	no
	second floppy dummy bus width	number of graphics standards supported	business market	other cards
	mouse dummy			

continued

ANNEX B2 Continued

	Evans		Barzyk
	INSEE01	INSEE02	
ZIP dummy	no	no	no
CDROM	no	no	yes
CDROM speed	yes	no	no
CDRW	yes	no	no
DVD dummy	no	no	no
Sound card dummy	yes	no	no
Video (MB)	no	yes	no
Network card	yes	no	yes
Modem dummy	yes	no	yes
Speakers dummy	no	no	no
Case-type dummy	yes	no	yes
Warranty dummy	no	yes	no
Seller dummies	no	yes	yes
SCSI control	no	no	yes
Operating system	no	no	no
Other software	no	no	no
Other	number of slots	network location	

Okamoto and Sato	Lim and McKenzie	van Mulligen
no	no	no
no	no	no
no	no	no
no	yes	no
no	no	no
no	no	no
no	yes	no
no	yes	no
yes	no	no
no	yes	no
no	yes	no
no	yes	no
Apple	yes	yes
no	yes	no
no	no	no
no	no	no
TV tuner	expandability	USB port
vintage dummies		workstation

III

APPENDIXES

Appendix A

Biographies of Speakers*

MICHAEL BORRUS

Michael Borrus is a Managing Director of the Petkevich Group, an investment bank focused on the health-care and information technology industries. Before joining the Petkevich Group, Mr. Borrus was a Co-Director of the Berkeley Roundtable on the International Economy (BRIE) at the University of California at Berkeley and Adjunct Professor in the College of Engineering, where he taught management and technology.

He is the author of two books and over 60 chapters, articles, and monographs on a variety of topics including high-technology competition, international trade and investment, and the impact of new technologies on industry and society. For the last decade, he has served as consultant to a variety of governments and firms in the United States, Asia, and Europe on policy and business strategy for international competition in high-technology industries. Mr. Borrus is a graduate of Harvard Law School and a member of the California State Bar.

MARK BREGMAN

As executive vice president of product operations at VERITAS Software Corporation, Mark Bregman oversees VERITAS Software's engineering and

---

*As of February 2003.

product management departments to ensure integrated product delivery. Dr. Bregman is responsible for serving VERITAS Software's existing markets as well as for expanding the company's portfolio of storage software solutions.

Dr. Bregman spent 16 years at IBM where he managed the RS/6000 and Pervasive Computing divisions, IBM Research, and IBM Japan. He was also technical assistant to IBM CEO Lou Gerstner. Most recently, Dr. Bregman was CEO of Airmedia, a wireless Internet firm.

Dr. Bregman holds a bachelor's degree in physics from Harvard College and a master's degree and doctorate in physics from Columbia University. He also serves on the Advisory Board of OptronX, Inc., and on the Board of Overseers of Fermi National Accelerator Lab. He is a member of the Visiting Committee to the Harvard University Libraries, a member of the American Physical Society, and a senior member of IEEE.

CAROL A. CORRADO

Carol Corrado is Chief of the Industrial Output Section of the Division of Research and Statistics at the Federal Reserve's Board of Governors. She currently serves as a member of the Executive Committee of the NBER Conference for Research on Income and Wealth and as a member of the Statistics Committee of the National Association of Business Economists. Prior to joining the Governors Board in 1977, Dr. Corrado was a Research Associate at the Joint Center for Urban Studies of the Massachussetts Institute of Technology and Harvard University and worked as an Instructor at Tufts University. In the spring of 2000, she extended her field research as a Visiting Research Faculty Member in the Department of Economics at the University of Maryland.

Dr. Corrado's areas of interest include macroeconomics, technology change and growth, and data collection and estimation methodology. She received a B.S. in administration management science from Carnegie-Mellon University and her Ph.D. in economics from the University of Pennsylvania in 1976.

Dr. Corrado has published several books and articles including: "Industrial Production and Capacity Utilization: The 1997 Revision,'' *Business Economics* (July 1997); ''Decomposition of Productivity and Costs'' (with Laurence Slifman), *American Economic Review* (May 1999); and *Measuring Capital in a New Economy* (ed. with John Haltiwanger and Dan Sichel), which is forthcoming.

KENNETH FLAMM

Kenneth Flamm is the Dean Rusk Professor of International Affairs at the LBJ School at the University of Texas at Austin. Before this, he worked at the Brookings Institution in Washington, D.C, where he served 11 years as a Senior Fellow in the Foreign Policy Studies Program. He is a 1973 honors graduate of

Stanford University and received a Ph.D. in economics from MIT in 1979. From 1993 to 1995, Dr. Flamm served as Principal Deputy Assistant Secretary of Defense for Economic Security and Special Assistant to the Deputy Secretary of Defense for Dual Use Technology Policy. Defense Secretary William J. Perry awarded him the Department's Distinguished Public Service Medal in 1995.

Dr. Flamm has been a professor of economics at the Instituto Tecnológico de México in Mexico City, the University of Massachusetts, and George Washington University. He has also been an adviser to the Director General of Income Policy in the Mexican Ministry of Finance and a consultant to the Organization for Economic Cooperation and Development, the World Bank, the National Academy of Sciences, the Latin American Economic System, the U.S. Department of Defense, the U.S. Department of Justice, the U.S. Agency for International Development, and the Office of Technology Assessment of the U.S. Congress.

Dr. Flamm has made major contributions to our understanding of the growth of the electronics industry, with a particular focus on the development of the computer and the U.S. semiconductor industry. He is currently working on an analytical study of the post-Cold War defense industrial base and has expert knowledge of international trade and high technology industry issues.

DALE W. JORGENSON

Dale Jorgenson is the Frederic Eaton Abbe Professor of Economics at Harvard University. He has been a Professor in the Department of Economics at Harvard since 1969 and Director of the Program on Technology and Economic Policy at the Kennedy School of Government since 1984. He served as Chairman of the Department of Economics from 1994 to 1997. Dr. Jorgenson received his Ph.D. degree in economics from Harvard in 1959 and his B.A. in economics from Reed College in Portland, Oregon, in 1955.

Dr. Jorgenson was elected to membership in the American Philosophical Society in 1998, the Royal Swedish Academy of Sciences in 1989, the U.S. National Academy of Sciences in 1978, and the American Academy of Arts and Sciences in 1969. He was elected to Fellowship in the American Association for the Advancement of Science in 1982, the American Statistical Association in 1965, and the Econometric Society in 1964. He was awarded honorary doctorates by Uppsala University and the University of Oslo in 1991.

Dr. Jorgenson is President of the American Economic Association. He has been a member of the Board on Science, Technology, and Economic Policy of the National Research Council since 1991 and was appointed to be Chairman of the Board in 1998. He is also Chairman of Section 54, Economic Sciences, of the National Academy of Sciences. He served as President of the Econometric Society in 1987.

Dr. Jorgenson is the author of more than 200 articles and the author and editor of 20 books in economics. His collected papers have been published in nine volumes by The MIT Press, beginning in 1995. The most recent volume, *Economics and Producer Behavior*, was published in 2000.

Prior to Dr. Jorgenson's appointment at Harvard, he was Professor of Economics at the University of California, Berkeley, where he taught from 1959 to 1969. He has been Visiting Professor of Economics at Stanford University and the Hebrew University of Jerusalem and Visiting Professor of Statistics at Oxford University. He has also served as Ford Foundation Research Professor of Economics at the University of Chicago.

Forty-two economists have collaborated with Dr. Jorgenson on published research. An important feature of his research program has been collaboration with students in economics at Berkeley and Harvard, mainly through the supervision of doctoral research. This collaboration has often been the outgrowth of a student's dissertation research and has led to subsequent joint publications. Many of his former students are professors at leading academic institutions in the United States and abroad and several occupy endowed chairs.

DALEN E. KEYS

Dalen Keys joined DuPont in 1984 in Wilmington, Delaware, at the Experimental Station as a Research Chemist working on high speed photopolymer imaging systems. While at the Experimental Station from 1984 to 1988, he supported projects on microencapsulation, novel materials for the protection of photomasks (pellicles) and holographics. In 1988, he became a Research Supervisor in Towanda, Pennsylvania responsible for the product development efforts supporting the North American Proofing business. In 1991, he became Lab Head for the DuPont Graphic Arts business in Rochester, New York, and then moved into manufacturing as the Technical Superintendent in Parlin, New Jersey, in 1993.

In 1994, Dr. Keys relocated to the Frankfurt, Germany area for the DuPont Graphic Arts business. While in Germany from 1994 until 1998, he was Worldwide Technology Director, European Business Manager, and finally Neu Isenburg Plant Manager and Technology Director. He was responsible for global development and manufacture of Graphic Arts products and managed about 650 people. He joined DuPont iTechnologies (then known as P&EM) in August 1998 and participated in the initial "founding" of the DuPont Displays business. In December 1998 he became Global Technology Director for iTechnologies, reporting to the Group Vice President and General Manager. Since January 2001, Dr. Keys has focused his full attentions on Displays and is the Chief Technology Officer.

Dr. Keys obtained a B.S. in chemistry from the University of North Alabama in 1980 and a Ph.D. in physical organic chemistry from Rice University in 1984. He holds six patents and numerous publications. He was a member of the

Houghton College President's Advisory Counsel on Excellence. He is currently Chairman of the United States Displays Consortium.

STEVEN LANDEFELD

Steven Landefeld has been Director of the Bureau of Economic Analysis (BEA) since 1995. BEA is the statistical agency within the Department of Commerce responsible for the national, international, regional, and industry accounts—including such estimates as Gross Domestic Product (GDP), personal income, corporate profits, the U.S. balance of payments, State and local area personal income, U.S. capital stocks, input-output estimates, foreign direct investment estimates, and GDP-by-industry.

Prior to becoming Director of BEA, Dr. Landefeld served in a number of other capacities at the Bureau, including Acting Director, Deputy Director, and Associate Director for International Economics. While at BEA, he has led a number of pioneering efforts in statistics, including the introduction of unbiased estimates of real GDP and prices, the development of monthly estimates of trade in goods and services, alternative balance of payments accounts, integrated economic and environmental accounts, and the use of data exchanges with foreign banks to improve international capital estimates.

Dr. Landefeld also has led a number of managerial improvements at the Bureau including the introduction of a performance-based personnel system, the development of "private-sector" financial accounts (BEA was one of the first Bureaus in the Department to receive an unqualified opinion from an outside auditor on its financial statement), and the move from an antiquated mainframe to an integrated micro-computer network (BEA was the first major statistical agency to successfully make such a move).

Before coming to BEA, Dr. Landefeld held a number of positions, including Chief of Staff for Presidents Reagan and Bush's Council of Economic Advisers, Director of the Business Issues Analysis Division at the Department of Commerce, and Research Assistant Professor at Georgetown University. He has authored numerous professional articles and has received numerous awards for his work including the Henri Willem Methorst Medal from the International Statistical Institute; two Abramson Scroll Awards from the National Association of Business Economists; Gold, Silver, and Bronze Awards from the Department of Commerce; and most recently, a Distinguished Executive Award from President Bush. Dr. Landefeld has served on numerous professional committees and working groups including those of the Organization for Economic Co-operation and Development, the International Monetary Fund, the United Nations, and the Conference on Research in Income and Wealth.

CHRIS A. MALACHOWSKY

Chris Malachowsky co-founded NVIDIA in April 1993 and has been Vice President of Hardware Engineering for the Company since that time. From 1987 until April 1993, Mr. Malachowsky was a Senior Staff Engineer for Sun Microsystems, Inc., a supplier of enterprise network computing products. From 1980 to 1986, Mr. Malachowsky was a manufacturing design engineer at Hewlett-Packard Company. Mr. Malachowsky was a co-inventor of Sun Microsystems' GX graphics architecture and has authored 39 patents, most of which relate to graphics.

Mr. Malachowsky holds a B.S.E.E. degree from the University of Florida and an M.S.C.S. degree from Santa Clara University.

MARILYN E. MANSER

Marilyn Manser is Associate Commissioner, Office of Productivity and Technology, at the U. S. Department of Labor, Bureau of Labor Statistics (BLS). From 1995 through 1999, she was Assistant Commissioner, Office of Employment Research and Program Development at BLS. She held other positions at BLS from 1984 to 1995. Previously, she was Assistant Director and Senior Economist at Mathematica Policy Research, 1978–1983, and Assistant Professor, Department of Economics, State University of New York at Buffalo, 1973–1978.

Dr. Manser has published a number of research papers, primarily in the areas of price measurement, consumer demand, and labor economics. She has coedited two National Bureau of Economic Research Studies in Income and Wealth conference volumes: *Labor Statistics Measurement Issues* (1998) and *Price Measurements and Their Uses* (1993).

Dr. Manser holds a Ph.D. in economics from the University of Wisconsin at Madison (1974), an M.A. in economics from Duke University (1968), and a B.S. in mathematics from the University of Maryland (1967).

DAVID F. MCQUEENEY

David McQueeney leads the IBM Global Services Intellectual Property and Asset Commercialization team. Dave is responsible for developing and deploying the business and technical strategies that maximize the codification and reuse of a broad range of intellectual capital for IBM's customers.

In his previous role, Dr. McQueeney launched the Emerging Business team at IBM Research. Bringing research innovations more quickly to the marketplace, and bringing marketplace forces closer to our researchers, the Emerging Business team was designed to explore several new approaches to further extend the impact of IBM's research investment on IBM's customers and partners.

Dr. McQueeney has also held other significant positions in IBM Research, including Director of the IBM Zurich Research Laboratory, Vice President of

Communication Technology, and Vice President of Technical Strategy and Worldwide Operations.

As General Manager of Global Solutions in IBM's Government Industry team, Dr. McQueeney led a team of technical and subject-matter experts offering a portfolio of integrated solutions to governments around the world and was also responsible for the IBM Global Services system integration team dedicated to the U.S. federal government.

David McQueeney joined IBM Research in 1988, after receiving a Ph.D. in low-temperature physics from Cornell University.

WILLIAM J. RADUCHEL

William Raduchel served as the Executive Vice President and Chief Technology Officer of AOL Time Warner, Inc. He assumed that position in 2001 from a similar role at America Online, Inc. He joined AOL in 1999 from Sun Microsystems, Inc., where he was chief strategy officer and a member of its executive committee. In his 11 years at Sun he was also chief information officer, chief financial officer, acting vice president of human resources and vice president of corporate planning and development. Prior to that he had senior executive roles at Xerox Corporation and McGraw-Hill, Inc.

Receiving his undergraduate degree in economics from Michigan State University, he earned A.M. and Ph.D. degrees in economics from Harvard University. He became a member of the National Research Council Board on Science, Technology, and Economic Policy in 2000. He is also a member of the National Resarch Council Committee on Internet Navigation and Domain Name Services and of the National Advisory Board of the Salvation Army and has served as a director of several public companies.

WILLIAM T. SIEGLE

William Siegle is Senior Vice President, Technology Operations at Advanced Micro Devices (AMD). He is also AMD's Chief Scientist. He is directly responsible for overseeing the operations of AMD's Technology Development Group, and the development of technology strategies that support AMD's business units.

Dr. Siegle joined AMD in 1990. Prior to that he was with IBM, most recently as director of the Advanced Technology Center. He has served as Chairman of the board of the Semiconductor Research Corporation (SRC) and the SRC Education Alliance. Siegle has also served on the Microelectronics Industrial Advisory Committees at MIT, Stanford CIS and Rensselaer, and on the boards of the ETEC Corporation and the Microelectronics and Computer Technology Corporation (MCC). He is currently a member of the board of International SEMATECH, and the SIA Technology Strategy Committee.

Dr. Siegle has a bachelor's degree, master's degree and doctorate in electrical engineering, all from Rensselaer Polytechnic Institute. He is a member of the IEEE, Tau Beta Pi, Eta Kappa Nu, and Sigma Xi.

WILLIAM J. SPENCER

William Spencer was named Chairman Emeritus of the International SEMATECH Board in November 2000 after serving as Chairman of the SEMATECH and International SEMATECH Boards since July 1996. He came to SEMATECH in October 1990 as President and Chief Executive Officer. He continued to serve as President until January 1997 and as CEO until November 1997. During this time, SEMATECH became totally privately funded and expanded to include non-U.S. members. Many gave SEMATECH part of the credit for the U.S. semiconductor turnaround in the 1990s.

Dr. Spencer has held key research positions at Xerox Corporation, Bell Laboratories, and Sandia National Laboratories. Before joining SEMATECH he was Group Vice President and Senior Technical Officer at Xerox Corporation in Stamford, Connecticut, from 1986 to 1990. He established new research centers in Europe and developed a plan for Xerox retaining ownership in spin-out companies from research.

Prior to joining the Xerox Palo Alto Research Center (PARC) as manager of the Integrated Circuit Laboratory in 1981 and as the Center Manager of PARC in 1982 to 1986, Dr. Spencer served as Director of Systems Development from 1978 to 1981 at Sandia National Laboratories in Livermore, and as Director of Microelectronics at Sandia National Laboratories in Albuquerque from 1973 to 1978, where he developed a silicon processing facility for Department of Energy needs. He began his career in 1959 at Bell Laboratories.

Dr. Spencer received the Regents Meritorious Service Medal from the University of New Mexico in 1981; the C. B. Sawyer Award for contribution to "The Theory and Development of Piezoelectric Devices" in 1972; and a Citation for Achievement from William Jewell College in 1969, where he also received an Doctor of Science degree in 1990. He is a member of the National Academy of Engineering and a Fellow of IEEE, and serves on numerous advisory groups and boards. He was the Regents Professor at the University of California in the spring of 1998. He has been a visiting professor at the University of California at Berkeley School of Engineering and the Haas School of Business since the fall of 1998. He is a Research Professor of Medicine at the University of New Mexico.

William Spencer received an A.B. degree from William Jewell College in Liberty, Missouri, and an M.S. degree in mathematics and a Ph.D. in physics from Kansas State University.

HOWARD TAUB

Howard Taub is Vice President and Director of the Printing and Imaging Research Center at HP Labs. He is responsible for research programs in digital imaging, commercial printing and publishing, photo finishing, printer system architecture, color science, network interconnect technologies, displays and breakthrough personal storage devices.

The center has had a long history of contribution to HP's printing and storage businesses. Some recent contributions include playing a central role in driving HP's move into commercial printing, providing core technology that motivated and enabled HP to enter the digital photography business and inventing the encoding scheme that is used in all DVD+RW optical drives.

Dr. Taub has worked at HP Labs for 22 years and, before that, at the IBM Watson Research Center and for Dataproducts Corporation. One of his early HP assignments was to manage the research project that invented and did fundamental research on HP's thermal ink-jet technology.

Dr. Taub earned his Ph.D. in solid state physics from the Polytechnic Institute of Brooklyn in New York City. He is an inventor on more than 20 patents in the printing and imaging area and is a member of the American Physical Society and the Society for Imaging Science and Technology.

JACK E. TRIPLETT

Jack Triplett is a Visiting Fellow at the Brookings Institution, in Washington, D.C. His current research at Brookings concerns productivity in health, finance and other services industries, with a focus on developing improved measures of output for these notably difficult to measure sectors of the economy. He serves as a consultant to international organizations and to the statistical agencies of a number of countries on issues of economic measurement and economic statistics.

From 1985 to 1997, he was Chief Economist, U.S. Bureau of Economic Analysis (on leave in 1996–1997 to the National Bureau of Economic Research). From 1971 to 1985, Mr. Triplett held positions at the U.S. Bureau of Labor Statistics, including Associate Commissioner for Research and Evaluation, and Chief of the Price Research Division. In 1979, he was Assistant Director for Price Monitoring at the Council on Wage and Price Stability. Before his government positions, he taught economics at Washington University (St. Louis) and the University of Oregon, where he was also Assistant Director of the Institute of Labor and Industrial Relations.

Dr. Triplett has written extensively on problems of economic measurement, including price indexes, national accounts, capital stock and labor input, and productivity and technical change. He is the editor of *Fifty Years of Economic Measurement* (with Ernst R. Berndt) and *The Measurement of Labor Cost*, both for the National Bureau of Economic Research, and of *Measuring the Prices of Medical Treatments*, published by The Brookings Institution. He is an elected fellow of the American Statistical Association, and is the 1997 winner of the Julius Shiskin Award for Economic Statistics, which is awarded jointly by the National Association of Business Economists and the Washington Statistical Society. He was born in Portland, Oregon, and attended Lewis and Clark College. He holds A.B., M.A., and Ph.D. degrees from University of California, Berkeley.

KENNETH E. WALKER

Kenneth Walker has held titles like Application Hunter, Evangelist and most recently, Vice President of Technology Strategy at Philips Electronics. He works as a bridge between the technology and strategy groups and helps define and promote devices for the digital future. Mr. Walker's primary focus is on new product and market applications, new business models and assisting with key third-party relationships. He utilizes a rich base of knowledge of the industry to form his technology trend analysis.

During the past four years at Philips, Mr. Walker was responsible for the establishment of new product lines and the expansion of Philips' business through emerging markets and digital product categories. He has held specific responsibility for the development of leading electronics, ASICs and integrated solutions to create intelligent devices involving displays, storage and integrated communications. He led the effort at Philips to create the industry's first "smart panel" LCD and was a leader in the development of LCD television. As Philips' representative to the USDC Technical Council, he advised on the development of U.S. competence in display research and he was a key technology liaison for the LG.Philips LCD joint venture. This effort enabled two highly competitive technology groups to come together for the benefit of all parties and to create the world's leader in flat panel displays.

Mr. Walker joined Philips after more than 10 years in Silicon Valley working for a variety of companies, from the innovative, content management start-up Verano, to industry visionaries such as Apple Computer. Prior to his time in Silicon Valley, Kenn received a B.S. in Information and Computer Science and an M.S. in Management from the Georgia Institute of Technology.

ROBERT WHITMORE

Robert Whitmore is the Senior Vice President of Product Development Engineering at Seagate, Inc., where he is manages all disk drive product development-engineering organizations from various design centers. His production applications at Seagate range from Consumer Electronics to PC computing to Enterprise Storage. In his tenure with Seagate, Mr. Whitmore has held several positions including Vice President of Enterprise Store Design, Vice President of the Twin Cities Manufacturing Operations, and several first- and second-level management positions within the Design Engineering field.

Mr. Whitmore currently lives in Eden Prairie, Minnesota. He received his B.S. in mechanical engineering from the University of Iowa and an M.S. in mechanical engineering from the University of Minnesota.

Appendix B

Participants List*

Ana Aizcorbe
Federal Reserve Board

Kimberly Bayard
Federal Reserve Board

Tabitha Benney
The National Academies

Michael Borrus
The Petkevich Group, LLC

Mark Bregman
Veritas Software Corporation

Alphonse Buccino
Contemporary Communication, Inc.

David Byrne
Federal Reserve Board

Greg Cardinale
Sandia National Labratories

Paul Chwelos
University of British Columbia

Michael Ciesinski
United States Display Consortium

Iain M. Cockburn
Boston University

Alessandra Colecchia
Organization for Economic
Cooperation and Development

E. William Colglazier
The National Academies

Hugh Conway
National Defense University

---

*Speakers in italics

Carol A. Corrado
Federal Reserve Board

David Dierksheide
The National Academies

Thomas Donahue
U.S. Government

Eric Emch
University of Chicago Law School

Kevin Finneran
Issues in Science and Technology Magazine

Kenneth Flamm
University of Texas at Austin

Christopher Forman
Carnegie Mellon University

Douglas A. Galbi
Federal Communications Commission

John Gardenier
Centers for Disease Control and Prevention

Shane M. Greenstein
Northwestern University

John Grosh
Department of Defense

Meg Hardon
The Dutko Group

Dietmar Harhoff
Ludwig Maximilians University

Chris Hayter
The National Academies

Cray Henry
Department of Defense

Robert Hershey
Engineering and Management Consulting

Jim Hock
Veritas Software Corporation

Michael Holdway
Bureau of Labor Statistics

Kenneth Jacobson
Consultant

David Scott Johnson
Bureau of Labor Statistics

Dale W. Jorgenson
Harvard University

Dalen E. Keys
Dupont Displays

Kathleen Kingscott
International Business Machines

Aaron Kirtley
National Institute of Standards and Technology

Elizabeth Kiser
Federal Reserve Board

Steven Landefeld
Bureau of Economic Analysis

William Long
Business Performance Research Associates

Johannes Loschnigg
Office of Senator Lieberman

Chris A. Malachowsky
NVIDIA Corporation

Marilyn E. Manser
Bureau of Labor Statistics

Victor R. McCrary
National Institute of Standards and Technology

David F. McQueeney
International Business Machines

Stephen Merrill
The National Academies

Dietmar Moch
Centre for European Economic Research (ZEW)

Kazuyuki Motohashi
Research Institute of Economy, Trade and Industry (RIETI)

David Y. Peyton
National Association of Manufacturers

William J. Raduchel

Robert H. Reuss
DARPA

Dimitri Sanga
Statistics Canada

Paul Schreyer
Organization for Economic Cooperation and Development

Sujai Shivakumar
The National Academies

William T. Siegle
Advanced Micro Devices

Mick Silver
Cardiff Business School

William J. Spencer
International SEMATECH, retired

Marc Stanley
National Institute of Standards and Technology

John Stevens
Federal Reserve Board

Howard Taub
Hewlett-Packard Company

Jack E. Triplett
The Brookings Institution

Kenneth E. Walker

David B. Wasshausen
Bureau of Economic Analysis

Philip Webre
Congressional Budget Office

Charles Wessner
The National Academies

Robert Whitmore
Seagate

Kam Yu
Lakehead University

Appendix C

Bibliography

Aizcorbe, Ana, Kenneth Flamm, and Anjum Khurshid. 2002. "The Role of Semiconductor Inputs in IT Hardware Price Decline: Computers vs. Communications." Federal Reserve Board Finance and Economics Series Discussion Paper 2002-37. August.

Archibald, Robert B., and William S. Reece. 1979. "Partial Subindexes of Input Prices: The Case of Computer Services." *Southern Economic* Journal 46(October):528–540.

Baily, M. N., and R. Z. Lawrence. 2001. "Do We Have an E-conomy?" NBER Working Paper 8243. April 23.

Bapco. 2002. "SYSmark ® 2002: An Overview of SYSmark 2002 Business Applications Performance Corporation." Available at *http://www.bapco.com/SYSmark2002Methodology.pdf*. Accessed February 19, 2003.

Bard, Yonathan, and Charles H. Sauer. 1981. "IBM Contributions to Computer Performance Modeling." *IBM Journal of Research and Development* 25:562–570.

Barzyk, Fred. 1999. "Updating the Hedonic Equations for the Price of Computers." Working Paper of Statistics Canada, Prices Division. November 2.

Bell, C. Gordon. 1986. "RISC: Back to the Future?" *Datamation* 32(June):96–108.

Benkard, C. Lanier. 2001. *A Dynamic Analysis of the Market for Wide Bodied Commercial Aircraft.* Stanford: Graduate School of Business, Stanford University. June.

Berndt, Ernst R., and Zvi Griliches. 1993. "Price Indexes for Microcomputers: An Exploratory Study." In Murray F. Foss, Marilyn Manser, and Allan H. Young, eds. *Price Measurements and Their Uses.* Studies in Income and Wealth 57:63–93. Chicago: University of Chicago Press for the National Bureau of Economic Research.

Berndt, Ernst R., Zvi Griliches, and Neal Rappaport. 1995. "Econometric Estimates of Prices in Indexes for Personal Computers in the 1990s." *Journal of Econometrics* 68:243–268.

Berndt, Ernst R., and Neal J. Rappaport. 2001. "Price and Quality of Desktop and Mobile Personal Computers: A Quarter-Century Historical Overview." *American Economic Review* 91(2): 268–273.

Berndt, Ernst R., and Neal J. Rappaport. 2002. "Hedonics for Personal Computers: A Reexamination of Selected Econometric Issues." Unpublished paper.

Bloch, Erich, and Dom Galage. 1978. "Component Progress: Its Effect on High-Speed Computer Architecture and Machine Organization." *Computer* 11(April):64–75.

Bourot, Laurent. 1997. "Indice de Prix des Micro-ordinateurs et des Imprimantes: Bilan d'une rénovation." Working paper of the Institut National De La Statistique Et Des Etudes Economiques (INSEE), Paris, France, March 12.

Bureau of Economic Analysis. 2001. "A Guide to the NIPAs." In *National Income and Product Accounts of the United States, 1929–97.* Washington, D.C.: Government Printing Office. Also available at *http://www.bea.doc.gov/bea/an/nipaguid.pdf*.

Butler Group. 2001. "Is Clock Speed the Best Gauge for Processor Performance?" *Server World Magazine* (September). Available at *http://www.serverworldmagazine.com/opinionw/2001/09/06_clockspeed.shtml*. Accessed February 7, 2003.

Brinkman, W. F. 1986. *Physics Through the 1990s.* Washington, D.C.: National Academy Press.

Bromley, D. Alan. 1972. *Physics in Perspective*. Washington, D.C.: National Academy Press.

Browning, L. D., and J. C. Shetler. 2000. *SEMATECH, Saving the U.S. Semiconductor Industry*. College Station: Texas A&M Press.

Cahners In-Stat Group. 1999. "Is China's Semiconductor Industry Market Worth the Risk for Multinationals? Definitely!" March 29.

Cale, E. G., L. L. Gremillion, and J. L. McKenney. 1979. "Price/Performance Patterns of U.S. Computer Systems." *Communications of the Association for Computing Machinery (ACM)* 22 (April):225–233.

Cartwright, David W. 1986. "Improved Deflation of Purchases of Computers." *Survey of Current Business* 66(3):7–9.

Cartwright, David W., Gerald F. Donahoe, and Robert P. Parker. 1985. "Improved Deflation of Computers in the Gross National Product of the United States." Bureau of Economic Analysis Working Paper 4. Washington, D.C.: U.S. Department of Commerce. December.

Cholewa, Rainier. 1996. "16M DRAM Manufacturing Cooperation IBM/SIEMENS in Corbeil Essonnes in France," *Proceedings of the 1996 IEEE/SEMI Advanced Semiconductor Manufacturing Conference*.

Chow, Gregory C. 1967. "Technological Change and the Demand for Computers." *American Economic Review* 57(December):1117–1130.

Chwelos, Paul. 2003. "Approaches to Performance Measurement in Hedonic Analysis: Price Indexes for Laptop Computers in the 1990s." *Economics of Innovation and New Technology* 12(3):199–224.

Cohen, Wesley M., and John Walsh. 2002. "Public Research, Patents and Implications for Industrial R&D in the Drug, Biotechnology, Semiconductor and Computer Industries" in National Research Council, *Capitalizing on New Needs and New Opportunities: Government-Industry Partnerships in Biotechnology and Information Technologies*. Washington, D.C.: National Academy Press.

Cole, Rosanne, Y. C. Chen, Joan A. Barquin-Stolleman, Ellen Dulberger, Nurhan Helvacian, and James H. Hodge. 1986. "Quality-Adjusted Price Indexes for Computer Processors and Selected Peripheral Equipment." *Survey of Current Business* 66(1):41–50.

Colecchia, Alessandra, and Paul Schreyer. 2002. "ICT investment and economic growth in the 1990s: Is the United States a unique case? A comparative study of nine OECD countries." *Review of Economic Dynamics* 5(2):408–442.

Cunningham, Carl, Denis Fandel, Paul Landler, and Robert Wright. 2000. *Silicon Productivity Trends*. International SEMATECH Technology Transfer #00013875A-ENG. February 29.

Dalén, Jorgen. 1989. "Using Hedonic Regression for Computer Equipment in the Producer Price Index." R&D Report, Statistics Sweden, Research-Methods-Development, Vol. 25.

Dulberger, Ellen R. 1989. "The Application of a Hedonic Model to a Quality Adjusted Price Index for Computer Processors." In Dale W. Jorgenson and Ralph Landau, eds. *Technology and Capital Formation*, pp. 37–75. Cambridge: MIT Press.

Dulberger, Ellen. 1993. "Sources of Price Decline in Computer Processors: Selected Electronic Components." In Murray Foss, Marilyn Manser, and Allan Young, eds., *Price Measurements and Their Uses.* Chicago: University of Chicago Press for the National Bureau of Economic Research.

The Economist. 2000. "A Thinker's Guide." March 30.

The Economist. 2001. "The Great Chip Glut." August 11.

Ein-Dor, Phillip. 1985. "Grosh's Law Re-visited: CPU Power and the Cost of Computation." *Communications of the Association for Computing Machinery (ACM)* 28(February):142–151.

Ericson, R., and A. Pakes. 1995. "Markov-Perfect Industry Dynamics: A Framework for Empirical Work." *Review of Economic Studies* 62:53–82.

Evans, Richard. 2002. "INSEE's Adoption of Market Intelligence Data for Its Hedonic Computer Manufacturing Price Index." Presented at the Symposium on Hedonics at Statistics Netherlands, October 25.

Fershtman, C., and A. Pakes. 2000. "A Dynamic Game with Collusion and Price Wars." *RAND Journal of Economics* 31(2):207–236.

Fisher, Franklin M., John J. McGowan, and Joen E. Greenwood. 1983. *Folded, Spindled, and Multiplied: Economic Analysis and U.S. v. IBM.* Cambridge, MA: MIT Press.

Flamm, Kenneth. 1987. *Targeting the Computer.* Washington, D.C.: The Brookings Institution.

Flamm, Kenneth. 1988. *Creating the Computer.* Washington, D.C.: The Brookings Institution.

Flamm, Kenneth. 1989. "Technological Advance and Costs: Computers vs. Communications." In Robert C. Crandall and Kenneth Flamm, eds., *Changing the Rules: Technological Change, International Competition, and Regulation in Communications*. Washington, D.C.: The Brookings Institution.

Flamm, Kenneth. 1993. "Measurement of DRAM Prices: Technology and Market Structure," in Murray F. Foss, Marilyn E. Manser, and Allan H. Young, eds., *Price Measurements and Their Uses*. Chicago: University of Chicago Press.

Flamm, Kenneth. 1996. "Japan's New Semiconductor Technology Programs," *Asia Technology Information Program Report No. ATIP 96.091*. Tokyo. November.

Flamm, Kenneth. 1996. *Mismanaged Trade? Strategic Policy and the Semiconductor Industry.* Washington, D.C.: The Brookings Institution.

Flamm, Kenneth. 1997. *More for Less: The Economic Impact of Semiconductors.* San Jose, CA: Semiconductor Industry Association. December.

Fransman, M. 1992. *The Market and Beyond: Cooperation and Competition in Information Technology Development in the Japanese System.* Cambridge, MA: Cambridge University Press.

Gordon, Robert J. 1989. "The Postwar Evolution of Computer Prices." In Dale W. Jorgenson and Ralph Landau, eds., *Technology and Capital Formation*, pp. 37–75. Cambridge, MA: MIT Press.

Gandal, Neil. 1994. "Hedonic Price Indexes for Spreadsheets and an Empirical Test for Network Externalities." *RAND Journal of Economics* 25.

Griffith, P. 1993. "Science and the Public Interest." *The Bridge* (Fall):16. Washington, D.C.: National Academy of Engineering.

Grindley, P., D. C. Mowery, and B. Silverman. 1994. "SEMATECH and Collaborative Research: Lessons in the Design of a High-Technology Consortia." *Journal of Policy Analysis and Management* 13.

Grossman, Gene, and Elhannan Helpman. 1993. *Innovation and Growth in the Global Economy*. Cambridge, MA: MIT Press.

Gowrisankaran, G. 1998. "Issues and Prospects for Payment System Deregulation." Working paper, University of Minnesota.

Handler, Philip. 1970. *Biology and the Future of Man.* London: Oxford University Press.

Harhoff, Dietmar, and Dietmar Moch. 1997. "Price Indexes for PC Database Software and the Value of Code Compatibility." *Research Policy* 24(4-5):509–520.

Holdway, Michael. 2001. "Quality-Adjusting Computer Prices in the Producer Price Index: An Overview." Bureau of Labor Statistics, October 16.

Horrigan, John Brendan. 1996. "Cooperation Among Competitors in Research Consortia." Unpublished doctoral dissertation, University of Texas at Austin.

Howell, Thomas. 2003. "Competing Programs: Government Support for Microelectronics," in National Research Council, *Securing the Future: Regional and National Programs to Support the Semiconductor Industry*, C. Wessner, ed. Washington, D.C.: The National Academies Press.

Ishida, Haruhisa. 1972. "On the Origin of the Gibson Mix." *Journal of the Information Processing Society of Japan* 13(May):333–334 (in Japanese).

Jorgenson, Dale W. 2001. "Information Technology and the U.S. Economy." *American Economic Review* 91(1).

Jorgenson, Dale W., Mun S. Ho, and Kevin J. Stiroh. 2002. "Information Technology, Education, and the Sources of Economic Growth Across U.S. Industries." Presented at the Brookings Workshop "Services Industry Productivity: New Estimates and New Problems," March 14. Available at *http://www.brook.edu/dybdocroot/es/research/projects/productivity/ workshops/ 20020517.htm*.

Jorgenson, Dale W., and Kevin J. Stiroh. 1999. "Productivity Growth: Current Recovery and Longer-Term Trends." *American Economic Review* 89(2).

Jorgenson, Dale W., and Kevin J. Stiroh. 2002. "Raising the Speed Limit: U.S. Economic Growth in the Information Age" in National Research Council. *Measuring and Sustaining the New Economy.* Dale W. Jorgenson and Charles W. Wessner, eds. Washington, D.C.: National Academy Press.

Kelejian, Harry H., and Robert V. Nicoletti. c. 1971. "The Rental Price of Computers: An Attribute Approach." Unpublished paper, New York University (no date).

Knight, Kenneth E. 1966. "Changes in Computer Performance: A Historical View." *Datamation* (September):40–54.

Knight, Kenneth E. 1970. "Application of Technological Forecasting to the Computer Industry." In James R. Bright and Milton E. F. Schieman, eds., *A Guide to Practical Technological Forecasting.* Englewood Cliffs, NJ: Prentice-Hall.

Knight, Kenneth E. 1985. "A Functional and Structural Measure of Technology." *Technological Forecasting and Technical Change* 27(May):107–127.

Koskimäki, Timo, and Yrjö Vartia. 2001. "Beyond Matched Pairs and Griliches-Type Hedonic Methods for Controlling Quality Changes in CPI Sub-indices." Presented at Sixth Meeting of the International Working Group on Price Indices, sponsored by the Australian Bureau of Statistics, April.

Levine, Jordan. 2002. "U.S. Producer Price Index for Pre-packaged Software." Presented at the 17th Voorburg Group Meeting, Nantes, France, September.

Levy, David, and Steve Welzer. 1985. "An Unintended Consequence of Antitrust Policy: The Effect of the IBM Suit on Pricing Policy." Unpublished paper, Rutgers University, Department of Economics, December.

Lias, Edward. 1980. "Tacking the Elusive KOPS." *Datamation* (November):99–118.

Lim, Poh Ping, and Richard McKenzie. 2002. "Hedonic Price Analysis for Personal Computers in Australia: An Alternative Approach to Quality Adjustments in the Australian Price Indexes."

Macher, Jeffrey T., David C. Mowery, and David A. Hodges. 1999. "Semiconductors." In National Research Council, *U.S. Industry in 2000: Studies in Competitive Performance.* David C. Mowery, ed. Washington, D.C.: National Academy Press.

Martin, Brookes, and Zaki Wahhaj. 2000. "The Shocking Economic Impact of B2B." *Global Economic Paper*, *37*. Goldman Sachs. February 3.

McKinsey Global Institute. 2001. *U.S. Productivity Growth 1995–2000, Understanding the Contribution of Information Technology Relative to Other Factors.* Washington, D.C.: McKinsey & Co.

Michaels, Robert. 1979. "Hedonic Prices and the Structure of the Digital Computer Industry." *The Journal of Industrial Economics* 27(March):263–275.

Moch, Dietmar. 2001. "Price Indices for Information and Communication Technology Industries: An Application to the German PC Market." Center for European Economic Research (ZEW) Discussion Paper No. 01-20, Mannheim, Germany. August.

Moore, Gordon E. 1965. "Cramming More Components onto Integrated Circuits." *Electronics* 38(8).

Moore, Gordon E. 1975. "Progress in Digital Integrated Circuits," *Proceedings of the 1975 International Electron Devices Meeting,* pp. 11–13.

Moore, Gordon E. 1997. "The Continuing Silicon Technology Evolution Inside the PC Platform." *Intel Developer Update.* (2, October 15).

Moylan, Carol, 2001. "Estimation of Software in the U.S. National Income and Product Accounts: New Developments." OECD Paper. September. Available at *http://webnet1.oecd.org/doc/M00017000/M00017821.doc*.

National Academy of Sciences, National Academy of Engineering, Institute of Medicine. 1993. *Science, Technology and the Federal Government: National Goals for a New Era.* Washington, D.C.: National Academy Press.

National Advisory Committee on Semiconductors. 1989. A *Strategic Industry at Risk.* Washington, D.C.: National Advisory Committee on Semiconductors.

National Advisory Committee on Semiconductors. 1990. *Capital Investment in Semiconductors: The Lifeblood of the U.S. Semiconductor Industry.* Washington, D.C.: National Advisory Committee on Semiconductors.

National Advisory Committee on Semiconductors. 1990. *Preserving the Vital Base: America's Semiconductor Materials and Equipment Industry.* Washington, D.C.: National Advisory Committee on Semiconductors.

National Advisory Committee on Semiconductors. 1991. *MICROTECH 2000 Workshop Report.* Washington, D.C.: National Advisory Committee on Semiconductors.

National Advisory Committee on Semiconductors. 1991. *Toward a National Semiconductor Strategy.* Vols. 1 and 2. Washington, D.C.: National Advisory Committee on Semiconductors.

National Advisory Committee on Semiconductors. 1992. *Competing in Semiconductors*. Washington, D.C.: National Advisory Committee on Semiconductors.

National Advisory Committee on Semiconductors. 1992. *A National Strategy for Semiconductors: An Agenda for the President, the Congress, and the Industry.* Washington, D.C.: National Advisory Committee on Semiconductors.

National Research Council. 1996. *Conflict and Cooperation in National Competition for High-Technology Industry.* Washington, D.C.: National Academy Press.

National Research Council. 1999. *The Advanced Technology Program: Challenges and Opportunities,* Charles W. Wessner, ed. Washington, D.C.: National Academy Press.

National Research Council. 1999. *Industry-Laboratory Partnerships: A Review of the Sandia Science and Technology Park Initiative,* Charles W. Wessner, ed. Washington, D.C.: National Academy Press.

National Research Council. 1999. *New Vistas in Transatlantic Science and Technology Cooperation,* Charles W. Wessner, ed. Washington, D.C.: National Academy Press.

National Research Council. 1999. *The Small Business Innovation Research Program: Challenges and Opportunities,* Charles W. Wessner, ed. Washington, D.C.: National Academy Press.

National Research Council. 2001. *The Advanced Technology Program: Assessing Outcomes,* Charles W. Wessner, ed. Washington D.C.: National Academy Press.

National Research Council. 2001. *Capitalizing on New Needs and New Opportunities: Government-Industry Partnerships in Biotechnology and Information Technologies,* Charles W. Wessner, ed. Washington, D.C.: National Academy Press.

National Research Council. 2001. *A Review of the New Initiatives at the NASA Ames Research Center,* Charles W. Wessner, ed. Washington, D.C.: National Academy Press.

National Research Council. 2001. *Trends in Federal Support of Research and Graduate Education.* Washington, D.C.: National Academy Press.

National Research Council. 2002. *Government-Industry Partnerships for the Development of New Technologies: Summary Report,* Charles W. Wessner, ed. Washington, D.C.: National Academy Press.

National Research Council. 2002. *Measuring and Sustaining the New Economy,* Dale W. Jorgenson and Charles W. Wessner, eds. Washington, D.C.: National Academy Press.

National Research Council. 2003. *Securing the Future: Regional and National Programs to Support the Semiconductor Industry,* Charles W. Wessner, ed. Washington, D.C.: The National Academies Press.

Nelson, Richard, ed. 1993. *National Innovation Systems.* New York: Oxford University Press.

Nelson, R. A., T. L. Tanguay, and C. C. Patterson. 1994. "A Quality-Adjusted Price Index for Personal Computers." *Journal of Business and Economics Statistics* 12(1):23–31.

Nordhaus, William D. 2002. "The Progress of Computing." Yale University, March 4.

Nikkei Microdevices. 2001. "From Stagnation to Growth: The Push to Strengthen Design." January.

Nikkei Microdevices. 2001. "Three Major European LSI Makers Show Stable Growth Through Large Investments." January.

Ohta, Makoto, and Zvi Griliches. 1976. "Automobile Prices Revisited: Extensions of the Hedonic Hypothesis." In Nestor E. Terleckyj, ed., *Household Production and Consumption.* Conference on Research in Income and Wealth, *Studies in Income and Wealth* 40:325–390. New York: National Bureau of Economic Research.

Okamoto, Masato, and Tomohiko Sato. 2001. "Comparison of Hedonic Method and Matched Models Method Using Scanner Data: The Case of PCs, TVs and Digital Cameras." Presented at Sixth Meeting of the International Working Group on Price Indices, sponsored by the Australian Bureau of Statistics, April.

Oliner, Stephen, and Daniel Sichel. 2000. "The Resurgence of Growth in the Late 1990's: Is Information Technology the Story?" *Journal of Economic Perspectives* 14(4).

Organization for Economic Cooperation and Development. 2000. *Is There a New Economy? A First Report on the OECD Growth Project.* Paris: Organisation for Economic Co-operation and Development. June.

PC World. 2003. "20 Years of Hardware." March.

Pake, G. E. 1966. *Physics Survey and Outlook.* Washington, D.C.: National Academy Press.

Pakes, Ariel. 2001. "A Reconsideration of Hedonic Price Indices with an Application to PCs." Harvard University, November.

Parker, Robert P., and Bruce Grimm. 2000. "Recognition of Business and Government Expenditures for Software as Investment: Methodology and Quantitative Impacts, 1959–98." Paper presented to BEA's Advisory Committee, May 5. *http://www.bea.doc.gov/bea/papers/software.pdf*.

Patrick, James M. 1969. "Computer Cost/Effectiveness." Unpublished paper summarized in Sharpe, 1969, p. 352.

Phister, Montgomery. 1979. *Data Processing Technology and Economics,* Second Edition. Bedford, MA: Santa Monica Company Publishing and Digital Press.

Prince, Betty. 1991. *Semiconductor Memories: A Handbook of Design, Manufacture and Application.* 2nd Edition. Chichester, UK: John Wiley and Sons.

Prud'homme, Marc, and Kam Yu. 2002. "A Price Index for Computer Software Using Scanner Data." Unpublished working paper, Prices Division, Statistics Canada.

Rao, H. Raghaw, and Brian D. Lynch. 1993. "Hedonic Price Analysis of Workstation Attributes." *Communications of the Association for Computing Machinery (ACM)* 36(12):94–103.

Ratchford, Brian T., and Gary T. Ford. 1976. "A Study of Prices and Market Shares in the Computer Mainframe Industry." *The Journal of Business* 49:194–218.

Ratchford, Brian T., and Gary T. Ford. 1979. "Reply." *The Journal of Business* 52:125–134.

Robertson, Jack. 1998. "Die Shrinks Now Causing Logic Chip Glut." *Semiconductor Business News.* October 15.

Rosch, Winn L. 1994. *The Winn L. Rosch Hardware Bible.* Indianapolis: Sams Publishing.

Schaller, Robert R. 1999. "Technology Roadmaps: Implications for Innovation, Strategy, and Policy." Ph.D. dissertation proposal, Institute for Public Policy, George Mason University.

Schaller, Robert R. 2002. "Moore's Law: Past, Present, and Future." Available at *http://www.njtu.edu.cn/depart/xydzxx/ec/spectrum/moore/mlaw.html*. Accessed July 2002.

Serlin, Omri. 1986. "MIPS, Dhrystones, and Other Tables." *Datamation* 32(June 1):112–118.

Semiconductor Industry Association. 1993. *Semiconductor Technology Workshop Conclusions.* San Jose, CA: Semiconductor Industry Association.

Semiconductor Industry Association. 1993. *Semiconductor Technology Workshop Working Group Reports.* San Jose, CA: Semiconductor Industry Association.

Semiconductor Industry Association. 1994. *The National Technology Roadmap for Semiconductors, 1994.* San Jose, CA: Semiconductor Industry Association.

Semiconductor Industry Association. 1997. *The National Technology Roadmap for Semiconductors, Technology Needs.* San Jose, CA: Semiconductor Industry Association.

Sharpe, William F. 1969. *The Economics of the Computer.* New York and London: Columbia University Press.

Sigurdson, J. 1986. *Industry and State Partnership in Japan: The Very Large Scale Integrated Circuits (VLSI) Project.* Lund, Sweden: Research Policy Institute.

Solow, Robert M., Michael Dertouzos, and Richard Lester. 1989. *Made in America*, Cambridge, MA: MIT Press.

Spencer, W. J., and P. Grindley. 1993. "SEMATECH After Five Years: High Technology Consortia and U.S. Competitiveness." *California Management Review* 35.

Stapper, Charles H., and Raymond J. Rosner. 1995. "Integrated Circuit Yield Management and Yield Analysis: Development and Implementation." *IEEE Transactions on Semiconductor Manufacturing* 8(2).

Statistics Finland. 2000. "Measuring the Price Development of Personal Computers in the Consumer Price Index." Paper for the meeting of the International Hedonic Price Indexes Project. Paris, France, September 27.

Stoneman, Paul. 1976. *Technological Diffusion and the Computer Revolution: The U.K. Experience.* Cambridge: Cambridge University Press.

Stoneman, Paul. 1978. "Merger and Technological Progressiveness: The Case of the British Computer Industry." *Applied Economics* 10:125–140.

Triplett, Jack E. 1985. "Measuring Technological Change with Characteristics-Space Techniques." *Technological Forecasting and Social Change* 27:283–307.

Triplett, Jack E. 1989. "Price and Technological Change in a Capital Good: A Survey of Research on Computers." In *Technology and Capital Formation*, Dale W. Jorgenson and Ralph Landau, eds. Cambridge, MA: MIT Press.

Triplett, Jack E. 1996. "High-Tech Productivity and Hedonic Price Indexes." In Organization for Economic Cooperation and Development, *Industry Productivity*. Paris: Organization for Economic Cooperation and Development.

Triplett, Jack E., and Barry Bosworth. 2002. "Baumol's Disease Has Been Cured: IT and Multifactor Productivity in U.S. Service Industries." Presented at the Brookings Workshop "Services Industry Productivity: New Estimates and New Problems," March 14. Available at *http://www.brook.edu/dybdocroot/es/research/projects/productivity/workshops/20020517.htm.*

van Mulligen, Peter Hein. 2002. "Alternative Price Indices for Computers in the Netherlands Using Scanner Data." Prepared for the 27th General Conference of the International Association for Research in Income and Wealth, Djurhamn, Sweden.

VeriTest. 2003. "Business Winstone™ 2002 Basics." Available at *http://www.veritest.com/benchmarks/bwinstone/wshome.asp*. Accessed February 19, 2003.

Wallace, William E. 1985. "Industrial Policies and the Computer Industry." The Futures Group Working Paper #007. Glastonbury, CT: The Futures Group (March).

Wolff, Alan Wm., Thomas R. Howell, Brent L. Bartlett, and R. Michael Gadbaw, eds. 1992. *Conflict Among Nations: Trade Policies in the 1990s.* San Francisco: Westview Press.

World Semiconductor Trade Statistics. 2000. *Annual Consumption Survey.*

Wyckoff, Andrew W. 1995. "The Impact of Computer Prices on International Comparisons of Labour Productivity." *Economics of Innovation and New Technology* 3:277–293.