Draft - Dissertation Prospectus Snapshot Version 1.0
26 September 2010

Connecting Minds in a Multimedia Episteme:

The Academic Supercomputer Centers and the Construction an Advanced Cognitive Infrastructure for the U.S. Research Community: 1983-1995

Kevin Walsh
Graduate Student
Department of History and
Program in Science Studies
University of California, San Diego
kwalsh@ucsd.edu

Abstract

In the early 1980s, a group of university physicists led a concerted effort to establish national supercomputer centers for use by the U.S. research community. Their initial struggle centered on freeing the supercomputing technology from the national weapons labs, where access was restricted to those who had the necessary security clearance. This was not a credential held by many academics. The majority of the advocates for supercomputing on university campuses were privileged by their Baby Boomer birthrights, their actual use of supercomputers, and by intellectual and family pedigrees connecting them to leading scientists of the Manhattan Project. Their credibility was enhanced by scientific accomplishment and a response to the federal call to research universities to contribute to U.S. economic competitiveness. This historical period was marked by high unemployment, global challenges to U.S industry, especially in the high technology sector, and a dramatic escalation of Cold War tensions. Supercomputing as a technology, was caught in the crosshairs of restrictive government polices and the promise of possibilities in the minds of a generation who came of age during the 1960s and 70s.

Based upon the recommendations of several scientific panel reports and Congressional Hearings, the National Science Foundation funded four national supercomputer centers on university campuses in 1985, at Cornell, Princeton, University of Illinois, Urbana-Champaign, and the University of California, San Diego. A fifth center was added in 1986, the Pittsburgh Supercomputer Center, a

partnership between Carnegie Mellon University and University of Pittsburgh. The centers quickly transformed from computer rooms to interconnected facilities integrating high performance networking, mass storage, and computer visualization. In becoming epistemic centers, the scientists and technical staff developed software tool kits and digital scientific workbenches, which enabled multidisciplinary distributed cognition, and the creation of shared knowledge in time and space gated by the bandwidth of universities' network interfaces.

The level of computational power in the academic centers held parity with the national weapons labs during the period under study, largely do the robust roll out of cutting-edge high performance experimental computer architectures that were welcomed in each of the centers. Students and researchers exercised the machines, seeking to calibrate their imaginations in self-similar silicon. The pervasive growth of microcomputers and personal computer technology was taking place in parallel to supercomputing platforms. The two paths would meet when in a cluster of interconnected microcomputers redefined what it meant be a supercomputer. The "hero" class of supercomputer user grew out of this computational arms race, with research agendas scaling in tandem with each incremental increase of computational power, data storage capacity, network bandwidth, and visualization capability.

The differences between what each center proposed to do initially, and what ensued, have influenced the development of computing and computer communications infrastructures that have extended epistemic boundaries, profoundly changed scientific practice, and coded the culture of a data dependent transnational society. The role of the centers within the High Performance Computing community shows that the focus supercomputing is not a machine, but new ways of thinking, learning, and communicating within a cognitive infrastructure. The users have proven to be the most important parts of this infrastructure.

Introduction – significance

"In an age when the machine is supreme, should a historian be allowed to ignore how machines are designed and modified?....We have no other remedy than to substitute, in place in place of the skills of a single man, the pooling of techniques, practiced by different scholars, but all tending to throw light upon a specific subject. This method presupposes a spirit of teamwork."[1]

- Marc Bloch

The significance of my topic can be measured in three principal ways: 1) the continued impact of technologies and the computer communications infrastructure, that were designed and developed as a result of the emergence of the academic supercomputer centers; 2) the changes in scientific practice as scientific users engaged their research problems within the epistemic settings of the centers; and 3) the range of scholarly fields I will call upon to explain and understand the consequences of this transformation. I seek to describe and analyze how the dynamic between the leaders, builders and users of this high performance computing infrastructure expanded the boundaries of what each

---

[1] Marc Leopold Benjamin Bloch, *The Historian's Craft* (Manchester: Manchester University Press, 1954). 68-69. It is important to note that Bloch wrote this passage in late 1940 to early 1941. He expands on this comment to predict that the future of history as a science will be guided by multidisciplinary efforts.

group believed was possible to know. It is important to keep in mind those boundaries within the ethos of the supercomputing community change with each incremental measure of computing power. Boundaries are potentially redrawn semi-annually with the announcement of the Top 500 fastest supercomputers, which in turns scales existing research agendas.[2] This thought style is intrinsic to the definition of a supercomputer - the fastest class of computer currently available.

Introduction – relation to other fields

My topic relates to questions in several fields and subfields:

- The philosophy of science, especially epistemology – how do we know what we know.
- The sociology of science, especially the social dynamics involving beliefs about how facts should be established and who has the authority, virtue, and credibility to do so across generational periods.
- The history of science, especially in stressing the historical and social contexts of the period under study. History shows that information and theories rejected as false under one set of historical conditions may be accepted under different conditions at a later time.
- The history of technology, especially in assessing the cultural and technical attributes of technology. The history of technology also speaks to the question of what it means to be human when technology has become increasingly pervasive in shaping our way

---

[2] Home | TOP500 Supercomputing Sites," n.d., http://www.top500.org/. The TOP500 project was started in 1993 to provide a reliable basis for tracking and detecting trends in high-performance computing. Twice a year, a list of the sites operating the 500 most powerful computer systems is assembled and released. The best performance on the Linpack benchmark is used as performance measure for ranking the computer systems. The list contains a variety of information including the system specifications and its major application areas.

of life in developed countries, communicating with our own bodies, our own minds, and each other.

- Science studies, which as a multidisciplinary field has taught that there is no proof without persuasion, and that persuasion is a historically contingent process involving, psychological, cultural (theological), political, and methods and style of communication.

- University studies - the university is a special institution and campuses are unique social and political settings. My project with intersect with fundamentals of the university ideal, academic freedom and the global role of higher education in society, projecting idea of service, and the relations linking the university to the international and national scene. The changing relationship between the university and industry is a key theme during the period under study.

- Science Policy – especially at the intersection with economic policy at the level of the National Science Board and sponsored research funding. The definition of National Science Foundation's research priorities during the period under study was greatly influenced by far reaching science policy and legislation that had a great impact in the status of intellectual property in the university setting. Science policy prioritized the funding of the high performance computing infrastructure, and shaped the historical context in which the academic supercomputer centers emerged.

- Lab studies – Labs are unique, often mission driven social and epistemic settings, often centered on an advanced technology, a specific scientific problem, and area of applied research. High performance computing was a key technology in the establishment of "trading zones" in national weapons labs. A rich variety of lab studies will inform my approach to studying the characteristics of the academic supercomputer centers as lab settings.

- Cognitive science and cognitive anthropology – "the central hypothesis of cognitive science is that thinking can best be understood in terms of representational structures in the mind and computational procedures that operate on those structures."[3] Cognitive anthropology concerns how humans think in different cultural settings, across cultures and in particular physical and social environments. As a resident of the high performance computing community, my study has strong leanings toward

---

[3] Paul Thagard, "Cognitive Science," n.d., http://plato.stanford.edu/archives/fall2008/entries/cognitive-science/.

ethnography, which will help explain the social and cognitive systems of the community.

- Computer science – supercomputers existed before the field of computer science was an academic discipline. Early supercomputer users were domain scientists, with physicists the majority. However, at the time academic supercomputer centers where established, computer science as set of academic subdisciplines was in a state of formation and self-definition, seeking acceptance as a field of science. My project is situated in a period where supercomputing and high performance computing architectures were becoming objects of academic study by computer scientists. Thus, the dynamic between domain scientists seeking to use supercomputing as a platform to help answer their own research problems, met with the research aims of computer scientists. Such developments were taking place in parallel with the increasing use of computing in all academic disciplines, giving rise to computational subdisciplines. Communities with their own conferences and publications in computation physics, computational chemistry and computational biology arose during the period under study.

All of these fields have their own issues and driving questions. Each, in their own way will contribute to my over arching historical question – why the ideas about the transformational potential of supercomputing took place when they did – no sooner, no later. Moreover, I will explore the consequences of place, the major research universities, with unique ties to Cold War science, and to the federal research establishment. Answers to these questions resonate in the present as well as in the future as funding agencies are much more interest in research outcomes and scientific impact as a metric for large grants for big science tools and infrastructure projects. At present, as was the case in the early 1980s, the university was called upon to contribute to U.S. economic competitiveness. The essential tension for research universities is how to do so while maintaining its academic mission amidst economic challenges on several fronts.

Historiography and review of the scholarship

There is little scholarly treatment of the role of supercomputing on science, and even less on academic science. Both of the original directors of SDSC and NCSA co-wrote books, which outline the capabilities of supercomputing, and provide a number of examples across scientific disciplines.

Published in 1987, *The Supercomputer Era* by the Director of SDSC, Sidney Karin and Norris Parker Smith was written with the general reader in mind.[4] The book is a snap shot of the uses of supercomputing during the 1980s in government, industry, and academia. The appendices are valuable for the listing of supercomputing companies at the time, and locations where supercomputing cycles were available to select users, and academics.

Supercomputing and the Transformation of Science, by William J. Kaufmann III and Larry Smarr, NCSA's first director, was published in 1993, and exhibits the a range of experiences of academic scientists in several domains who coupled their science to the recently available supercomputing tool.[5] By the

---

4 Karin, Stanley, and Norris Parker Smith. *The Supercomputer Era*. 1st ed. Boston: Harcourt Brace Jovanovich, 1987.

5 Kaufmann, William J., and Larry L. Smarr. *Supercomputing and the Transformation of Science*. New York: Scientific American Library: Distributed by W.H. Freeman, 1993.

time of publication, the Centers had been in operation for seven years. The technological advances in computing power, networking, and visualization during this period cast a logical pattern of development that could be expanded to support the concept of a national machine room, a virtual computer center which was geographically distributed, but appeared as one to a national research community.

Smarr and Kaufmann clearly place the use of supercomputing in a historical continuum, and boldly so. In their minds, supercomputing had transformed science because it was at the apex of a new methodology for scientific investigation. They outline three modes of science: The first is the experimental/observational mode. Galileo employed this mode in the early 1600s by use of a telescope to observe the movement of the Earth's moon and the moons of Saturn. The second mode of science is the theoretical mode, which is exemplified by Isaac Newton in the mid-1600s, where the regularities and patterns of the physical world are expressed mathematically. The mathematical relationships came to form laws of nature. Albert Einstein's equations of general relativity are perhaps the most well known example of the second mode of science. The third mode of science, according to Smarr and Kaufmann, is the computational mode.[6] Computing, especially supercomputing, would enable scientists to solve mathematical theoretical models at the scale and

---

[6] Ibid. Page 4.

complexity present in large, multi-variable natural events, such as climate change, storms, earthquakes, and disease transmission in whole populations.

The majority of historical scholarship in laboratory science, especially big science, is aimed at the national labs, such as Brookhaven (BNL), Fermi (FNL) Los Alamos (LANL), and Lawrence Livermore (LLNL).[7] I believe there is much

---

[7] The following works are representative examples: Badash, Lawrence, Joseph Oakland Hirschfelder, and Herbert P. Broida. *Reminiscences of Los Alamos, 1943-1945*. Dordrecht, Holland ; Boston: Reidel Pub. Co. ; Hingham MA Boston, 1980.; Dennis, Michael Aaron. ""Our First Line of Defense": Two University Laboratories in the Postwar American State." *Isis* 85, no. 3 (1994): 427-55.; Edmondson, Frank K. *Aura and Its Us National Observatories*. Cambridge ; New York: Cambridge University Press, 1997.; Furman, Necah Stewart. *Sandia National Laboratories : The Postwar Decade*. 1st ed. Albuquerque, NM: University of New Mexico Press, 1990.; Goldman, Joanne Abel. "National Science in the Nation's Heartland: The Ames Laboratory and Iowa State University, 1942-1965." *Technology and Culture* 4, no. 3 (2000): 435-59.; Galison, Peter Louis, and Bruce William Hevly. *Big Science : The Growth of Large-Scale Research*. Stanford, Calif.: Stanford University Press, 1992.; Geiger, Roger L. *To Advance Knowledge: The Growth of American Research Universities, 1900-1940*. New York: Oxford University Press, 1986.; Geiger, Roger L. *Research & Relevant Knowledge : American Research Universities since World War Ii*, Transaction Series in Higher Education. New Brunswick, N.J.: Transaction Publishers, 2004.; Heilbron, J. L., eScholarship (Online service), and Robert W. Seidel. *Lawrence and His Laboratory a History of the Lawrence Berkeley Laboratory*, California Studies in the History of Science. Berkeley: University of California Press, 1989.; Hoddeson, Lillian, and Gordon Baym. *Critical Assembly : A Technical History of Los Alamos During the Oppenheimer Years, 1943-1945*. Cambridge England ; New York: Cambridge University Press, 1993.; Kleinman, Daniel Lee. *Politics on the Endless Frontier : Postwar Research Policy in the United States*. Durham: Duke University Press, 1995.; Leslie, Stuart W., and American Council of Learned Societies. *The Cold War and American Science the Military-Industrial-Academic Complex at MIT and Stanford*. New York: Columbia University Press, 1993.; Lowen, Rebecca S., *Creating the Cold War University the Transformation of Stanford*. Berkeley: University of California Press, 1997.; Seidel, Robert W. "A Home for Big Science: The AEC's Laboratory System." *Historical Studies in the Physical and Biological Sciences* 16, no. 1 (1986): 135-75. Truslow, Edith C., and Kasha V. Thayer. *Manhattan District History : Nonscientific Aspects of Los Alamos Project Y, 1942 through 1946*. Los Alamos, N.M.: Los Alamos Historical

to learn from this work that applies to my own project because the supercomputing capabilities that the university academics wanted were in the national labs, and the facility at LLNL was the model that was most often referenced for use by such researchers.   There are some reoccurring themes in the studies of the national labs. One such theme is that the labs espoused to be national, multidisciplinary, and open to university academics throughout the 1950s, 60s, and 70s. However, the Cold War mission of the national labs, and the perennial need to secure funding militated to keep “outsiders” such as university researchers at a distance.

One of the large-scale, core texts in the history of computers is Paul Ceruzzi's *A History of Modern Computing*. The two editions of this book, the first in 1998 and the second in 2003, chronicle developments in computer technologies beginning with the transition from punched-card machines to the first vacuum-tube electronic computers in the late 1940s and early 1950s, to the

---

Society, 1991.; Westwick, Peter J. *The National Labs : Science in an American System, 1947-1974*. Cambridge, Mass.: Harvard University Press, 2003.

development of silicon integrated circuits in the late 1960s and early 1970s, the emergence of the personal computer in the late 1970s, and networked computers in the 1980s and 1990s. In this work, Ceruzzi uses a theoretical framework that draws on STS concepts of social construction and Thomas Hughes's technological systems, in order to explain the historical changes in computer technology and the various groups of people relevant to this narrative.

Two other influential works are Paul Edwards's 1996 book, *The Closed World: Computers and the Politics of Disclosure in Cold War America*, and Janet Abbate's 1999 book, *Inventing the Internet.* Edwards analyzes the computer in both its technical form and its importance as a cultural metaphor during the Cold War. He uses this understanding to argue that the relationship between computers and American culture is one of co-construction: both have shaped each other, rather than one category determining the other. Abbate examines the history of the Internet, describing the creation of pre-internet networking, and the evolution of the military ARPANET into the public World Wide Web. She focuses on both the details of the technology involved and the individual people and organizations that played a role in this history.

In his recent book, from 2006, *From Counterculture to Cyberculture: Stewart Brand, the Whole Earth Network, and the Rise of Digital Utopianism*, Fred Turner discusses the influence of the 1960s counterculture on the development of the personal computer. Turner traces direct links between major figures in countercultural movements, primarily Stewart Brand, and individuals involved in the creation of early personal computers. He also examines the place of certain countercultural values in the design in use of computers that have continued into aspects of computer culture today.

Two influential books in the history of computing focus on the contributions of single individuals to fundamental changes to computer technologies. William Aspray, in his 1990 book, *John von Neumann and the Origins of Modern Computing* establishes the institutional background von Neumann worked in where he developed theories of computer organization that brought about the shift from punched-cards to programs stored in memory. Moving forward to another major shift in computer thought, Thierry Bardini's 2000 book, *Bootstrapping: Douglas Engelbart, Coevolution, and the Origins of Personal Computing* traces Engelbart's pioneering research on computers as a means to augment human intelligence and his influence on the development of the personal computer. According to Bardini, Engelbart viewed the computer as something like a prosthesis, one that would develop with its user over time to lead to new abilities for people. Both the computer and the user were required to learn from each other, in order to develop and provide people with better ways to live in the world.

Much in the history of computing has been written as popular or industry sponsored histories, instead of traditional academic works. Possibly the most important popular history is Steven Levy's 1984 book, *Hackers: Heroes of the Computer Revolution*. Here, Levy explains how early computer hackers, following an almost utopian view of technological progress, played an essential role in the development of major computer innovations. A number of well-researched popular histories are written by *New York Times* journalist John Markoff. A recent book of his, from 2005, *What the Dormouse Said: How the 60s Counterculture Shaped the Personal Computer Industry*, covers historical subject matter similar to Turner's work. However, Markoff provides an autobiographical perspective on this history, as he has been working as a journalist on these subjects during this time period and writes from his own involvement in this history.

My dissertation project is also somewhat autobiographical because I experienced the period under study as a high performance computer network specialist at a major research university, which is also a host site to one of the original national academic supercomputer centers. In contrast to the works outlined above, my focus is on universities and supercomputer centers as epistemic settings, and more specifically, on first generation of scientists that came to make up the high performance computing community. As indicated by my larger bibliography, I will situate my approach in a larger written discourse on the capabilities of computers as thinking machines, and on the human-computer interface.

Core question of dissertation project

How did the creation and development of the academic supercomputer centers contribute to the emergence of an advanced cognitive infrastructure that expanded epistemic boundaries and changed scientific practice?

This question is important to address as a historical project because it will focus on the first generation of scientists who embraced high performance computing, global computer networks, large scientific data sets, and the development of computer visualization tools to make meaning of computational outputs of unprecedented size and scale. In the early 1980s, the scientists and the computer creative class that I follow in my project were the leading edge Baby Boomers, coming into positions of responsibility on university campuses and national labs. Their response to their times was conditioned by their formative years during the 1960s and 70s. In 2010, almost 30 years later, many of the same historical actors are witness to a similar set of national economic conditions in which the federal government is calling upon the scientific community to contribute to society. However, this time it is different.[8]

I want to know what were they thinking then – in the early 1980s and through the mid-1990s. I want to know what are they thinking now, and how the continuum of their life experience can help explain the past through the present. It is in their thought style to make proposals, to prognosticate, to offer “what if” scenarios, and assert their expected research outcomes. The differences between what they proposed in the early 1980s and what occurred subsequently is the body of evidence for my project. Their scientific lives are exercises in cognitive kinesis in relation to the movement the interface between computational and human capabilities.

---

[8] Carmen M Reinhart, *This Time Is Different: Eight Centuries of Financial Folly* (Princeton: Princeton University Press, 2009). Reinhart’s organization of historical evidence, at the very least, is an exemplar that shows that history can show what not to do. I suggest that the Baby Boomers in the high performance community are now the senior advisors, and may have some valuable guidance. One of the aims of project is to collect and organize their memories.

Data, evidence, and method

Primary sources

Proposals are the basis of my dissertation project. The range of proposals that I will analyze will include the proposals to create the academic supercomputer centers, and the proposals from potential users who wanted to use the supercomputers. Potential users had to request an allocation of supercomputer time, justify their request, and convince reviewers how their research problem would benefit from supercomputer time. The proposals indicated what outcomes were expected, the credentials of the researchers, the anticipated impact of the research, and a register of supporters, who often wrote letters of support. I then connect what was proposed to what actually occurred, as explained in the published papers that describe research findings which could be attributed to supercomputer use. I will analyze how research problems changed over time as long-term users engaged with supercomputing, and what collaboration networks may have emerged. More importantly, I want to know what the users learned and how their thinking changed over time.

Other sources of institutional evidence are the newsletters of the centers, their annual reports, and their annual reviews carried out by the NSF. The newsletters chronicle what resources each center was developing, and how users were using those resources. The newsletters also report the work being done by the technical support staff, subject matter experts, and consultants who were not the research scientists, but the human interface between the scientific users and the supercomputer center resources. The annual reports indicate what the centers chose to promote as center successes, additions to infrastructure, and user success stories, that were often accompanied by user testimonials. The annual reviews from the NSF reveal a very nuanced assessment of the centers' performance, and indicate shifts in research priorities at the national level.

Unpublished sources also present a wealth of data. Supercomputer center staff and users became adept at creating computer presentations and demonstrations. I have uncovered several presentations on plastic foils, files of computer presentations in Ashton Tate Persuasion, a software application used before the ubiquity of PowerPoint. Posters and videos are other forms of unpublished evidence that I have access to. They are important sources because most represent forms of communication within social settings such as workshops, panel discussions, and academic conferences. I will explore what such media reveals about the thinking of the time, as forms of persuasion, and patterns presentation content sharing.

Oral histories are also an important data source. Many of my subjects remain active in the high performance computing community, and have a body of work that I can study, and pose informed questions to in person. My approach is aimed at answering the question, "What did you think was possible at the time?"

Preliminary fieldwork and research done

To date I have carried out archival research at the University of Illinois, Urbana-Champaign and the University of California, San Diego. I have assembled over 500 unpublished reports, presentations and digital objects, which I am placing in a digital archive. I have carried out a longitudinal study of Gordon Bell Prize winners for a set of posters that were on display at the 20th Anniversary of the annual supercomputing conference in 2008 – SC2008. I also was a key member of the SC 20th anniversary, and organized an exhibit which contained artifacts from all years of the conference, and featured the first Cray-1 as part of the public exhibit.

I have also carried out oral history interviews with over 50 principal participants in the creation of the academic supercomputer centers, as well as key users. Many of the interviews were video recorded, and have been placed in the digital archive, which I created, that is now supported by the IEEE Computer Society.

I continue as a lead in the SC history effort at the annual supercomputing conference, and am assisting in the production of another history exhibit for SC' 2010 in November. The exhibit will also be complemented by an oral history video recording center, as well as outreach to students who will help in the gathering supercomputing history objects from exhibitors and attendees.

Preliminary outline and proposed chapter summaries

My dissertation can be divided into three sections. The first is the "Access" section, which I cover in chapters 1, 2, and 3. The section will treat the supercomputer as a technology, its significance as a measure of nation's technological capabilities, and as a technology subject to controlled access. It will focus on how a group physicists, most from leading research universities led the effort to not only expand access to supercomputing, but to do so on university campuses. Their success meant they had to go about building supercomputer centers.

Section II is the infrastructure phase. I cover networking in chapter 3, the range of supercomputing platforms installed in the centers in chapters 4 and 5, mass storage platforms in chapter 6, and scientific visualization in chapter 7. It is important to keep in mind that there was very little commercial software for supercomputers. The operating systems, computer job schedulers, network protocol stacks, network equipment, and software applications were created by the academic centers. The end result was an interconnected system – an infrastructure – that connected users to high performance computing resources.

Section III will focus on scientific users of maturing infrastructure.

Section I

Chapter 1 – Weapons of Mass Computation and the Peace of Open Science

This chapter will cover the origins of the academic supercomputer centers, focusing on how the university created the consensus and credibility to obtain funding and support to establish the centers. I will explore the series of "Blue Ribbon" committee reports, Congressional hearings calling for expanded access to supercomputing. There were many points of contention; largely do the view that supercomputing was a technology vital to military defense. I will focus on what the predictions of stakeholders as to the benefits to the country and to the research community, and how they framed their arguments in the face of opposition from the national weapons labs.

Chapter 2 – The Sons of Manhattan and the Winning of Supercomputer Access

This chapter will cover the generation of physicists who led the efforts to lead the efforts to create the academic supercomputer centers.

Chapter 3 – The Political Economy of the Post-Cold War Research University

This chapter covers the impact of the Mansfield Amendment and the Bayh Dole Act on university research, including the cooperative agreement that would allow the academic supercomputer centers to "sell" 10% of the NSF funded infrastructure.

Section II

Chapter 4 – Cognitive Bandwidth: Connecting a Million Minds in 80 Milliseconds

While the supercomputer centers were under construction, a national academic computer network backbone was implemented to interconnect the centers, and more importantly to connect the center to users.   I give considerable attention to the NSFnet project, a unique, though competitive collaboration between the National Science Foundation, industry and the university. The High Performance Computing Act of 1991 and was a precursor to legislation that led to the creation of or the Information Super Highway and the transition of the NSFnet to the commodity Internet.

I am less interested on the sequence of events that lead to the transition of the NSFnet to the commercial Internet, but the role of the academic supercomputer centers in creating the actual network technology, and tools to connect users to the center resources. Thus, my focus is on the users of the network and how it changed what they thought was possible.

For example, the centers developed TCP/IP drivers for PCs because at the time, most PC networking solutions did not use TCP/IP. The supercomputer center in Illinois, the National Center for Supercomputer Applications (NCSA), developed a tool for PCs and Macs, called NCSA Telnet. It gave PCs the ability to connect to five computers concurrently on the NSFnet or any university network that connected to the NSFnet. The tool also allowed the PC and act as a FTP server for the sharing of files on an individual PCs local storage drive (floppy and hard disk). This is a simple utility extended the capabilities the users personal computer. I want explore how this ability to connect to other computers in, what was truly a matter of milliseconds, change what user thought possible in carrying out their research.

Chapter 5 – The Explosion of High Performance Computing Architectures

With the network pathways constructed, the centers became home to several supercomputers in the late 1980s and early 1990. Many of the machines had very novel architectures. The centers installed and supported several models of supercomputers from Cray – the X-MP/48, the Y-MP8, C90, TD, TE

T90, SCS, Thinking Machines, NCube, the Intel IPSC and Paragon, and the Tera MTA. My focus in on how the machines were made usable for the scientific user base. Each machine has distinguishing characteristics – personalities that presented technical and intellectual challenges to the center support staff and users.

With advances in microprocessors, new processor interconnect mechanisms came to challenge more specialized supercomputers.

Chapter 6 – Hide the Cray: The Invasions of the "Killer Micros"

The impact of high performance scientific desktop workstations such as Sun and Apollo workstations became more significant as their power and graphics increased, while their purchase price was falling. Personal computing technologies, with such rapidly growing market, were also experiencing performance gains. In face of the cost of supercomputers, the "attack of the killer micros" was announced at the annual Supercomputing conference in 1989. Networks of Workstations (NOW) were constructed by the technical staff at the centers, and demonstrated a price/performance ratio that was very compelling. Shortly thereafter, clusters of PCs were built that presented other challenges to big iron supercomputing. My questions of these developments include: What was the impact on the users, and their scientific computing requirements? What were the affects of the apparent reclassification of supercomputing? Who were the stakeholders? How were they different in background and training from earlier supercomputer users?

Chapter 7 – Seeing is believing: My Mind's Eye is Your Mind's Eye

This chapter the development of scientific visualization tools at the academic supercomputer centers, including the first 3D computer graphics. A megaflop computer produces a megabyte data storage challenge. Thus, the centers had to construct data repositories and schemas to organize scientific data sets. The shear size of the data outputs motivated the in-house development of scientific visualization codes and tools in order to make meaning from multidimensional computations.

I will focus primarily on the work carried out by visual artists and scientific visualization experts at NCSA and SDSC, including Donna Cox, Tom De Fanti, and Maxine Brown. Consistent with my focus on the user, I will explore the dynamic between what the user sees and how they think, especially when using tools to interact with data sets to draw out relationships in the underlying data.

Chapter 8 – Minds Modulo Data Storage: Towards Petascale Memory Practices in the Sciences

This chapter will cover how each increase in computer power creates larger and larger scientific data sets. I will focus on the creation of scheduler tools and the development of the Storage Resource Broker (SRB), and the work of the Data Intensive Computer Environment (DICE). One of the key questions is what does the user do with all the data?

It is more than a simple matter of storage. The driver was to have user-friendly online access to large data collections distributed across multiple sites. The centers had to create the standard interfaces so that collections could be shared between applications on the fly.

Section III

Chapter 9 – Community Tools and the Building of Interoperable Computational Selves

This chapter will cover the variety of discipline specific computer tools that were developed in the community, and how the shared use of common tools enabled distributed collaboration. One example I will focus on is the Biology Workbench. It was launched in 1989, with the aim of assembling databases and analysis tools molecular or structural biologists would want from a single access platform – a researchers PC. The biologist did not install the software on his or her PC but accessed a remote platform at one of the supercomputer centers.

Chapter 10 – The Making of the Supercomputer User Class

This chapter will cover the "hero class" user of supercomputers that have research agendas that require more and more infrastructure, faster machines, storage, visualization, network bandwidth, etc. I will focus on exemplar users in weather prediction (Kelvin Droegemeier), earthquake simulation (Jacobo Bielak and Kim Olsen), and protein folding (Peter Kollman).

These scientific areas were identified in the NSF's Grand Challenge Program for the High Performance Computing Community. I will focus on how the users and the center's subject area experts greeted the challenges, what they thought was possible, and why, and connect the initial goals with the actual findings.

Chapter 11 – The Genesis and Development of a Computational Fact

This chapter will covers how the computer, specifically the supercomputer came to gain credibility as a tool to produce "facts." I am currently exploring the role played by the supercomputer centers in assessing ozone depletion. Sherwood Rowland, who was one of the recipients of the Nobel Prize in chemistry in 1995, was a user of the centers resources.

Chapter 12 – Cognitive Infrastructures and the New Mosaic of Scientific Practices

This chapter will cover how computation, visualization, networking, mass storage evolved into a high performance cognitive infrastructure that has the transformed scientific practice. I will focus on the creation of the NCSA Mosaic Web browser and its companion Web server. I will explore how the technologies developed at the centers coalesced to provide the foundation for Web enabled science, in which data banks, community models, high performance computers and multimedia tools (voice, video, graphics) connect our minds.

Bibliography

See separate document.