Lawrence Livermore National Laboratory
Computation Directorate

Dona L. Crawford
Computation Associate Director
Lawrence Livermore National Laboratory
7000 East Avenue, L-559
Livermore, CA 94550

September 14, 2009

Dear Colleague:

Several years ago, I commissioned Alex R. Larzelere II to research and write a history of the U.S. Department of Energy's Accelerated Strategic Computing Initiative (ASCI) and its evolution into the Advanced Simulation and Computing (ASC) Program. The goal was to document the first 10 years of ASCI: how this integrated and sustained collaborative effort reached its goals, became a base program, and changed the face of high-performance computing in a way that independent, individually funded R&D projects for applications, facilities, infrastructure, and software development never could have achieved.

Mr. Larzelere has combined the documented record with first-hand recollections of prominent leaders into a highly readable, 200-page account of the history of ASCI. The manuscript is a testament to thousands of hours of research and writing and the contributions of dozens of people. It represents, most fittingly, a collaborative effort many years in the making.

I'm pleased to announce that *Delivering Insight: The History of the Accelerated Strategic Computing Initiative (ASCI)* has been approved for unlimited distribution and is available online at https://asc.llnl.gov/asc_history/.

Sincerely,

Dona L. Crawford
Computation Associate Director
Lawrence Livermore National Laboratory

Delivering Insight

The History of the Accelerated Strategic Computing Initiative (ASCI)

Prepared by: Alex R. Larzelere II

For

Lawrence Livermore National Laboratory
Under Sub-Contract B545072

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory in part under Contract W-7405-Eng-48 and in part under Contract DE-AC52-07NA27344.

UCRL-TR-231286

Contents

Contents……………..………..iii

Figures…………….. v

Acknowledgements….. vii

Foreword by Dona Crawford…. ix

Executive Summary 1

Chapter One: New Tools for Scientific Insight...7

Chapter Two: Building the Initiative...19

Chapter Three: Applications – At the Heart of Delivering Insight..33

Performance and Safety Applications – *The Direct Connection to the Weapons*…..41
Materials Modeling – *At the Level of First Principles* 49
Programming Models – *Creating the Applications* 55

Chapter Four: Platforms – Power-Plants for Simulations..65

ASCI Red – *Breaking the TeraFLOP/s Barrier*...71
ASCI Blue Pacific and Mountain – *Keeping the Industry Viable*..79
Linux Clusters – *Providing Cost Effective TeraFLOP/s*...87
BlueGene/L – *Collaborative Innovation at its Best* 95
ASC Purple – *Fulfilling the Promise of Power* 105

Chapter Five – Environments for Simulation Capabilities..111

Parallel Programming Tools – *Enabling a New Approach to High-Performance Applications* 117
Scalable Visualization – *Insights from Data* 125

Chapter Six – Partnering to Deliver Insight 135

Academic Alliances – *Harnessing the Power of Universities* 139
IBM and Lawrence Livermore – *A Sense of Shared Mission* 149
ASCI Organization – *Synergy at the Edge of Chaos* 157

Chapter Seven – Impact and Lessons Learned 165

Chapter Eight – Looking Toward the Future 171

Epilogue – A Future Envisioned by Visionaries 177

Schlesinger Award Citation for Gil Weigand 178
Text of Gil Weigand's Acceptance Speech 179
Schlesinger Award Citation for Vic Ries 181
Text of Vic Ries' Acceptance Speech 182

Appendix A – Chronology of Events 187

Appendix B – ASCI Platform Configurations 191

Appendix C – ASCI Leadership 193

About the Author 195

Glossary 197

Bibliography 205

Index 213

Figures

1-1 Vic Reis 11
1-2 ENIAC 11
1-3 Vignette Timeframes 17
2-1 Gil Weigand 22
2-2 ASCI Organizational Plan 23
3-1 Galileo's Gravity Experiment 34
3-2 Example of a Finite Element Mesh 37
3-3 Example of a Validation Experiment 38
3-4 Comparison of Simulation and Experiment for Validation 46
3-5 Example of 3D Simulation of Rayleigh Taylor Instability 47
3-6 Physics Occurs at Many Scales 51
3-7 A 160 Million Atom Simulation of Copper Undergoing a Shock 53
4-1 The ASCI Red System at Sandia 76
4-2 ASC Red Storm at Sandia 77
4-3 The ASCI "Curve" as it Appeared in the *1996 ASCI Program Plan* 80
4-4 ASCI Blue Mountain at Los Alamos 84
4-5 ASCI Blue Pacific at Livermore 85
4-6 Typical Beowulf Cluster Configuration 90
4-7 Multi-programmatic Capability Resource (MCR) at Livermore 92
4-8 CPlant at Sandia 93
4-9 A Model of the Japanese Earth Simulator 98
4-10 BlueGene/L Architecture 100
4-11 BlueGene/L, with Characteristic Slanted Cabinets 101
4-12 Atomic Level Simulation of Molten Tantalum Solidification 102
4-13 ASC Purple at Livermore 107
5-1 TotalView Screen Shot 122
5-2 Bertha Display 130
5-3 Livermore PowerWall 131
5-4 Sandia VIEWS Corridor 131
5-5 Los Alamos CAVE 132
6-1 University of Illinois Urbana-Champaign Simulation of Solid Rocket Operation 145
6-2 University of Utah Simulation of Explosive Reactions in a Fire 145
6-3 California Institute of Technology Simulation of Energetic Material Reaction to Shockwaves 146
6-4 University of Chicago Cover of Physics Today Showing FLASH Simulation Result 146
6-5 Stanford University Simulation of Turbine Engine Performance 147
6-6 The IBM 701 Computer at Livermore 150
6-7 Practices for Building Large-Scale Advanced Computing Systems 154

Acknowledgements

This history of the Accelerated Strategic Computing Initiative (ASCI) was commissioned by Lawrence Livermore National Laboratory under subcontract B545072. Now operated for the Department of Energy's National Nuclear Security Administration by the Lawrence Livermore National Security, LLC, the Laboratory was operated by the University of California throughout the time covered by this history.

In distilling many thousands of hours of work done by hundreds of people into a brief overview, I have only been able to scratch the surface of the achievements made by the people at the National Laboratories, Defense Programs, universities, and in industry. Their clever innovations, devoted contributions, and extraordinary dedication made the successes of ASCI possible.

Many people helped prepare this history. In particular, I acknowledge the contributions of the Lawrence Livermore Editorial Board who provided me invaluable assistance in sorting through the mass of available material. The members of that board included Tom Adams, Brandt Esser, Randy Christensen, Lynn Kissel, Steve Louis, Mike McCoy, Charlie McMillan, Terri Quinn, Jim Rathkopf, Mark Seager, and Mary Zosel. I would also like to thank the Lawrence Livermore Associate Directors Dona Crawford and Bruce Goodwin for their support in this effort.

Telling the history of ASCI would be impossible without access to the people who were directly involved. Much of the early history of the Initiative was not recorded and would be lost had I not been able to talk with these people. So, my appreciation goes out to people who gave of their time to be interviewed for this book. This includes: Tom Adams, Ed Barsis, Steve Berggren, Bill Camp, Randy Christensen, Dave Cooper, Dona Crawford, Tom D’Agostino, Nick Donofrio, Sid Drell, Bo Ewald, Bruce Goodwin, Art Hale, Thuc Hoang, Paul Hommert, Sid Karin, Lynn Kissel, Steve Louis, Christian Mailhiot, Mike McCoy, Charlie McMillan, Bob Meisner, Paul Messina, John Morrison, Jose Munoz, Dave Nelson, David Nowak, Ed Oliver, James Peery, Terri Quinn, Jim Rathkopf, Bill Reed, Vic Reis, Mark Seager, Hank Shay, Larry Smarr, Mike Sohn, Tom Sterling, Bruce Tarter, John Toole, Mike Vahle, Dick Watson, Gil Weigand, and Mary Zosel.

I thank Alan Altschuler, Stephanie Chu, Karyn Furlong, and Matt Sablan, who provided invaluable assistance by reading drafts and helping me find ways to explain the difficult concepts that were integral to making ASCI both challenging and ultimately successful. I also need to thank Dorothy Whiteford of WhiteMatter Consulting for her editing work. Finally, thanks go to my wife, Mary Ann Larzelere, for her patience as I undertook this project as well as for her valuable inputs reading draft after draft.

-Alex Larzelere

Foreword by Dona Crawford

The Accelerated Strategic Computing Initiative, or ASCI, is an example of government-sponsored science at its best. Launched in response to an urgent national necessity, it was grand in scope and scale, involving many people, organizations, and technologies. It brought together the Department of Energy's National Security Laboratories, industry, and academia in a vast collaborative effort and culminated in extraordinary scientific and technological achievements.

In many ways, ASCI transformed the way the Laboratories interact with each other; it also fostered sea changes in high-performance computing, dramatically affecting the way in which next-generation supercomputers are created, acquired, and employed. ASCI contributed significantly to a nascent revolution in which computational simulation assumed its place as a peer to theory and experiment in a fundamentally new paradigm of science.

First and foremost, ASCI was mission-driven, and it accomplished all the objectives of the mission. In being mission driven, it had to assemble all the components necessary to achieve the goals of that decade. It created capabilities critical to the ongoing success of the nation's Science Based Stockpile Stewardship Program, which has ensured the safety and reliability of the nuclear stockpile since 1994.

This ASCI history was commissioned to document how ASCI reached its goals and to illuminate how this integrated and sustained collaborative effort changed the face of high-performance computing in a way that independent, individually funded R&D projects for applications, facilities, infrastructure, and software development never could have achieved.

As this history goes to press, the impacts of the capabilities enabled and advanced by ASCI are becoming more important and obvious. We face daunting global national security challenges in the coming decades. In global warming and climate change, in conservation and renewable energy, and in antiterrorism and national defense, indeed in a host of areas, we face scientific and technological challenges of unprecedented complexity. In all of these areas, the modeling of a complete physical system to predict its evolution, possibly more so than the modeling of a scientific phenomenon in order to understand its nature, will play the key role both in understanding the problems and in gaining insight into their solutions. This change in emphasis and the development of an integrated set of methodologies to achieve mission goals are the two key contributions of ASCI. This is, indeed, great work in progress.

Alex Larzelere, who undertook to write this history, was a key leader when ASCI was created and maintained an interest in the Initiative throughout its existence. I thank Alex for his effort and perseverance in creating this history: for the long hours of personal interviews that he conducted, for his meticulous collection of project documentation, and for his integration of the documented record with the first-hand recollections of many ASCI principals into a highly readable account of the 10 years of ASCI.

I also thank Van Emden Henson of LLNL for his considerable effort as publishing editor: his work in clarifying the text, as well as in incorporating the many suggestions of the document reviewers, has appreciably strengthened the document.

An editorial committee assisted in oversight and review of the several drafts of the history. Each willingly contributed time, reading the document (in some cases reading multiple drafts), correcting errors, adding details, or making clarifying comments. For their efforts, I thank Tom Adams, Randy Christensen, Brandt Esser, Lynn Kissel, Steve Louis, Charlie McMillan, Mike McCoy, Terri Quinn, Jim Rathkopf, Mark Seager, and Mary Zosel.

Dona Crawford
Associate Director for Computation
Lawrence Livermore National Laboratory
April 2009

Executive Summary

The history of the Accelerated Strategic Computing Initiative (ASCI) tells of the development of computational simulation into a third fundamental piece of the scientific method, on a par with theory and experiment. ASCI did not invent the idea, nor was it alone in bringing it to fruition. But ASCI provided the wherewithal—hardware, software, environment, funding, and, most of all, the urgency—that made it happen.

On October 1, 2005, the Initiative completed its tenth year of funding. The advances made by ASCI over its first decade are truly incredible. Lawrence Livermore, Los Alamos, and Sandia National Laboratories, along with leadership provided by the Department of Energy's Defense Programs Headquarters, fundamentally changed computational simulation and how it is used to enable scientific insight.

To do this, astounding advances were made in simulation applications, computing platforms, and user environments. ASCI dramatically changed existing—and forged new—relationships, both among the Laboratories and with outside partners. By its tenth anniversary, despite daunting challenges, ASCI had accomplished all of the major goals set at its beginning. The history of ASCI is about the vision, leadership, endurance, and partnerships that made these advances possible.

Why ASCI?

ASCI was created out of need. The Initiative was established as a critical part of the Science Based Stockpile Stewardship (SBSS) program, which began in 1994. SBSS was instituted to employ a new method of assessing the reliability, performance, and safety of the U.S. nuclear weapons arsenal. Before SBSS, those assessments had been made using full-scale nuclear testing. With the end of the Cold War in 1992, President George H. W. Bush established a unilateral moratorium on nuclear testing, and President William J. Clinton later made the moratorium permanent with his endorsement of the Comprehensive Test Ban Treaty (CTBT). The end of nuclear testing presented significant challenges to the stewards of the U.S. nuclear arsenal.

At the time testing was banned, Defense Programs, the government organization responsible for the weapons, was in the Department of Energy (DOE), supported in its mission by three National Laboratories—Lawrence Livermore, Los Alamos, and Sandia. The Laboratories provide the people and the scientific resources used to design the nuclear weapons and to certify their performance and safety. Annually, the directors of the three Laboratories give the President a statement expressing whether the weapons continue to be safe and reliable and will perform as expected. In 2000, the National Nuclear Security Administration (NNSA) was created to oversee the nuclear weapons programs. NNSA is a semi-autonomous agency within DOE, and Defense Programs was moved into NNSA.

Prior to SBSS, the Laboratories approached the initial design and annual certification of the nuclear weapons systems leaning heavily on the traditional scientific method. In this approach, scientists use fundamental principles to devise theories about how things work in the physical world. Experiments are then used to confirm, deny, modify, or extend the theories. For thousands of years, scientists have used this interplay between theory and experiment to gain insight into nature. Until 1992, full-scale nuclear tests, at first conducted in the atmosphere and later conducted underground, were the primary means of obtaining information about how, and how well, the weapons worked. Testing formed a vital part of executing the Laboratories' missions. With the end of such testing, SBSS called for a new way of doing science.

Supplementing the traditional approach, high-performance computers were used by the Laboratories almost since the invention of computing. ENIAC, one of the world's first electronic computers, was used to calculate the expected explosive yield of early thermonuclear weapon designs. Ever since then, computers have been used to explore particular theories to calculate properties, project yields, estimate effects, and, to a limited extent, perform simulations of weapon design testing. In recent decades, the Laboratories used some of the world's most powerful computers to understand massive quantities of data generated by full-scale nuclear tests. The computers were employed to analyze those data and use them to extrapolate results to different designs.

After underground nuclear testing ended in 1992, the Laboratories looked to computer simulations to take on a new and even more important role. On August 11, 1995, President Clinton announced his administration's support for the CTBT and put his faith in SBSS. He said, "I am assured by the Secretary of Energy and the Directors of our nuclear weapons labs that we can meet the challenge of maintaining our nuclear deterrent under a CTBT through a Science Based Stockpile Stewardship program without nuclear testing." [1]

SBSS was developed to facilitate a better scientific understanding of nuclear weapons operation. This would provide the Laboratories with the means to evaluate the impact, on weapon safety, reliability, and performance, of the aging of the weapons. It was also intended to assess the effects of replacing aging components with new, remanufactured parts. SBSS laid out aggressive plans to create several new scientific capabilities. The new capabilities included a number of large experimental facilities to test various aspects of weapons science. One of these aggressive plans was for a new approach to using computational simulations. This plan became known as ASCI.

The 'A' in ASCI stands for 'Accelerated' and was a critical attribute of the program. ASCI had to accelerate the development of the computational simulation capabilities because it was essential that they be validated by comparison with data from actual underground tests; without that validation there could be no confidence in simulation, no matter how sophisticated it appeared. But only a limited number of scientists and engineers had been involved with actual underground nuclear tests. ASCI simulations would have to be validated by these testing veterans, who possessed the knowledge and experience to understand the data from the old tests (so-called legacy data) and compare the stockpile assessments based on legacy data with assessments from advanced simulations. Unfortunately, these people were expected to retire or leave the Laboratories within 10 years. ASCI's challenge, therefore, was to build a new simulation capability on a scale never before attempted and rarely even imagined and to do it on an accelerated basis.

What ASCI Accomplished

For ASCI to deliver on its promises, it simultaneously had to accelerate development of a wide range of technologies, but especially in three areas: simulation applications, computing platforms, and the environments to support them. *Applications* are the computer programs used to simulate the weapons and their operations. Before ASCI, the Laboratories' applications had been written to supplement the full-scale nuclear testing and could be validated readily by comparison to the tests; hence, applications only needed to incorporate a limited set of physics, and one- or two-dimensional models generally sufficed. These relatively simple codes were also necessary because existing computers, even the supercomputers of the 1980s and early 1990s, were limited in power, speed, and memory. Moreover, the lack of computing power limited the number of physical systems and the duration of events that could be studied in a single simulation. To meet the mission and to make SBSS a reality, ASCI's applications would have to capture all three dimensions and as much physics as possible and be able to simulate full systems operating over long periods of time.

But to run simulations using the anticipated sophisticated multiphysics applications, the Laboratories needed significantly larger, faster, and more powerful computers, or *platforms*. A standard unit of computer speed and power is the FLOP/s, that is, one floating-point operation per second. An early calculation done at Lawrence Livermore National Laboratory (LLNL) concluded that more than 100 teraFLOP/s (i.e., 100 trillion FLOP/s) would be required to execute the simulations needed by SBSS to sufficient accuracy. This audacious estimate meant that computing power at LLNL would have to increase by over 7000 fold, since at the time LLNL's most powerful computer provided only 13.7 gigaFLOP/s (13.7 million FLOP/s). Accomplishing this in just the 10 years ASCI allotted implied a technology growth rate many times that given by Moore's Law, then the industry standard for predicting increases in computing power. Somehow, the Initiative would have to accelerate the development of high-performance computing systems—otherwise there would be nothing on which to run the applications needed by SBSS.

To acquire the platforms, ASCI could not simply issue a purchase order. The computers they needed did not exist, and would have to be researched, designed, and constructed. Worse yet, the high-performance computing industry was in disarray at that time. The market was shrinking as its dominant customer, the federal government, began to focus on budget deficit reduction. As a result, there were too many companies selling and too few customers buying. The industry was in upheaval; some companies went bankrupt, others merged, and some got out of the high-performance computer business.

This upheaval was partly due to the burgeoning success of the Personal Computer (PC). For the first time, computers were available for individual desktops; no longer was the individual user restricted to the limited access of the mainframe computers, housed in large computer centers and attended by flocks of specialists. The extraordinary growth in popularity of the PC meant that many companies were suddenly devoted primarily to exploiting this new and seemingly limitless market. While the PC put tremendous pressure on companies, it also provided a new way to deliver computing power. Until the late 1980s, high-end systems were built so that they appeared to the application as a single large processor and a single large bank of memory. The PC looked the same, except with a small processor (called a microprocessor) and small memory. But the introduction of PCs led to a new approach to the architecture of high-performance computers, in which many microprocessors, each with its own bank of memory, would be networked into a much larger system that would run applications using the many microprocessors operating in parallel.

The applications for these computers would have to be split into many operations running simultaneously on separate processors. Information needed by one processor from a program running on another processor would have to be sent across a communications network. This approach to running applications was called parallel processing. One big advantage of parallel processing was that parallel computers could offer vast amounts of computing power, and, by leveraging the large PC market for microprocessors, might do so at a much lower cost.

When ASCI was launched, it was not clear what direction the chaotic computer industry was headed. Conventional serial processing machines were reaching the limit of their potential. While parallel computers looked promising, many thought they would be impossible to program effectively, and with thousands of parts, some would fail so frequently that the computers would be unusable.

ASCI did not have much of a choice. The Laboratories realized that to meet the demands of the expected applications, they would have to move their codes onto parallel computers. ASCI leaders also recognized the importance of enabling computer companies to develop commercially viable technology, to protect the vitality of an industry critical to the success of the Initiative.

To create applications for the new platforms, ASCI would have to provide a robust set of capabilities that are collectively called *user environments*. This environment included, among other things, tools such as compilers, debuggers, data handling and storage controls, operating systems, job schedulers, and communications systems. Collected together, these myriad tools formed the environment, and would be needed to allow the end-users, the weapons scientists, to create models and simulations. Because the simulations would create huge volumes of results, the user environment would need to provide innovative tools to visualize and analyze the data. The environment also had to facilitate inter-Lab sharing of resources and data while providing the highest levels of security.

ASCI Delivers

By 2005, ASCI had fulfilled the goals set out a decade earlier. The applications had been created, allowing nuclear weapons scientists and engineers to gain a better understanding of how the weapons work. Indeed, these applications allowed users to see things that were previously unknown—unrecognized in experiments and not imagined in theories.

By 2005, the ASCI computers had met and surpassed the computing power threshold set in the early LLNL estimate. In October of that year, a ceremony was held to commission the 100 teraFLOP/s Purple system at LLNL. Purple was built by IBM and consisted of more than 12,000 commercial microprocessors tied together with a fast interconnection network. Among the most impressive features was that, despite physically being in Livermore, Purple could be used by all three Laboratories to securely run highly classified applications relevant to weapons issues.

IBM also installed BlueGene/L in 2005. This machine represented a new approach to high-performance architecture and used more than 131,000 low-power processors to deliver a peak of 360 teraFLOP/s of computing power, far exceeding the ASCI target of 100 teraFLOP/s. The system was designed around the idea that using larger numbers of cheaper, simpler, and smaller components could dramatically reduce system cost, size, and energy consumption, all while increasing computing power. As with Purple, BlueGene/L was used almost immediately to shed light on important questions about material properties.

The fact that simulations could be productively run on the Purple and BlueGene/L systems in 2005—as soon as they were delivered—was remarkable and demonstrates ASCI's success accelerating the development of the user environments.

None of the advances fostered by ASCI—in applications, computing platforms, and environments—would have been possible without an extensive program of outreach and partnership. The scope of the work needed to create the required simulation capabilities was well beyond what any one organization could hope to accomplish. ASCI's history is largely about the evolution of relationships among the Initiative's stakeholders. Defense Programs and its Laboratories learned new ways to interact. The Laboratories formed new, exceptionally productive relationships with the computer industry and with other national programs interested in high-performance computing. Finally, ASCI's research focus and scientific requirements inspired the Laboratories to establish deep relationships with universities, a process that included the founding and development of several groundbreaking simulation-science centers at major research universities.

How ASCI Did It

ASCI delivered. It would be impossible to ascribe its success to a single reason, as many factors clearly contributed. In some ways, the timing was simply right. The end of the Cold War meant that Defense Programs, and its Laboratories, had no choice but to change. In a dramatically changed world, they faced a new mission that required unprecedented levels of computing power. At the same time, the expanding field of parallel processing had matured enough to be thought capable of handling nuclear weapons simulations. ASCI began just as simulation technologies were ready to flourish, if given the right push. But making that push required leadership and a lot of patient, sharply focused, hard work.

The factors in ASCI's success generally fell into four categories: vision, leadership, endurance, and partnership. ASCI developed and communicated a clear, mission-based vision: the ban on full-scale nuclear testing meant that SBSS would depend critically on computer simulation capabilities that did not exist at the time, and to create those capabilities, ASCI would have to develop new applications, platforms, and environments on an unprecedented scale. ASCI succeeded in propagating this vision to everyone involved in the Initiative.

That clarity of vision was due to ASCI's leadership, the second factor in ASCI's success. ASCI was led by exceptional people who fostered a "we can do it" culture that permeated the entire Initiative. This culture was promulgated from the beginning by outstanding leaders at Defense Programs, the Laboratories, and in the computer industry. The early leaders, especially Vic Reis (DOE Assistant Secretary for Defense Programs) and Gil Weigand (ASCI's first Director), focused on the ASCI vision and had the will to act as necessary to see the mission done. They recruited like-minded individuals to lead at all levels of the endeavor. In the end, the culture they inspired—of competence, teamwork, and a shared sense of mission—was justified by extraordinary accomplishment.

Getting ASCI started required vision and leadership, but keeping it going through a decade of daunting challenges required endurance, a third crucial factor in ASCI's success. The challenge for ASCI was to maintain a sense of urgency during the long, sustained effort to develop the required simulation capabilities. The cutting-edge research took time—time to develop the world's most complex applications, to invent and build the world's most powerful computers, and to create user environments of unprecedented capabilities. The endurance of the women and men doing that research, patience and encouragement of the leadership, and support from the funding sources made it possible.

The final key to ASCI's success was its development and reliance on partnership. ASCI worked best when it built relationships, both within its organizations and with outside partners. Committees and task forces drew members from all stakeholders. Progress meetings fostered peer review and horizontal communication, encouraging both innovation and facilitating timely progress. External relationships with computer companies, other federal programs, and the academic community helped solve problems faster, deepened the well of ideas, and promoted the scientific utility of computational simulation. Ultimately, the outreach and openness followed in the American tradition of working together to meet an urgent need for the nation.

Telling the History

ASCI's history is complicated. The Initiative involved hundreds of people at the Laboratories and Defense Programs, along with hundreds more in academia and industry. During its first decade, ASCI spent over $5.2 billion, over 80% of which went to support people researching technology advances. Because so many activities took place simultaneously, it is nearly impossible to relate the history in a purely chronological, comprehensive fashion, and limited space means the hard work and individual accomplishments of many people who made important contributions cannot be fully described. Therefore, this history of ASCI is told in the following way. The first two chapters lay a foundation describing the need for the Initiative and describing what was required to get it underway. The next four chapters describe the four major areas in which ASCI produced extraordinary feats by employing fundamentally new approaches: applications, platforms, environments, and partnerships. Each of these themes is illuminated through a number of vignettes highlighting important events or approaches. Two chapters highlighting the impact of ASCI and its future close out the story.

In its first decade, ASCI changed how Defense Programs and the Laboratories assess the performance, safety, and reliability of weapons in the U.S. nuclear stockpile. In doing so, ASCI helped foster a change in the fundamental approach to science by creating tools that enable computer simulation to stand as peer to theory and experiment.

The future impact on the world of science is not yet clear—simulations for science at the ASCI scale have only existed for a few years—but it is an exciting future. New discoveries have already been made using computational simulations and many more are in the pipeline. Just as ASCI had to change the world to build these capabilities, it is clear that having the capabilities will change our understanding of the world.

---

[1] Office of the Press Secretary, *Statement by the President*, 3.

Chapter One
New Tools for Scientific Insight

Insight can be defined as the capacity to obtain an accurate and deep understanding of a person, situation, or thing. A deep, accurate, science-based understanding of nuclear weapons is vital for maintaining a safe, reliable weapons stockpile and the deterrent to war the stockpile provides. The National Nuclear Security Administration (NNSA), the semi-autonomous agency within the Department of Energy (DOE) responsible for the stockpile, is acutely aware of the need for insight about the weapons. During the first 47 years of U.S. nuclear weapons programs, that insight was largely obtained through full-scale nuclear testing. In 1992, nuclear testing was banned, and the Laboratories realized they would have to invent new ways to arrive at scientific insight.

In the decade from 1995 to 2005, the Accelerated Strategic Computing Initiative, or ASCI, delivered remarkable new tools, based on computational simulation, that enabed scientists to gain that insight. These capabilities have provided the United States with new ways to understand the performance and safety of the weapons in its nuclear arsenal. More profoundly, ASCI's development of these new methods has contributed greatly—even led—to a fundamental change in the way the world does science; the use of high-performance simulations in science and engineering is beginning to have significant effects on people's everyday lives.

Theory and Experiment to Develop Insight

Traditionally, scientific insight has been obtained through interplay of *theory* and *experiment* in an iterative process of developing hypotheses (theory), testing them (experiment), analyzing the experimental results, and revising the hypotheses. Theory is not easily manufactured; it arises from the human imagination, employing (and often discovering) fundamental principles to explain how and why things happen the way they do in the physical world. The Theory of Relativity, for example, was developed in the incredible imagination of Albert Einstein and led scientists to discoveries that enabled the invention of the first atomic weapons. Experimentation involves the observation and measurement of actual physical processes, and many experiments produce results that directly illuminate the questions where insight is sought. Sometimes the physical processes of interest are impossible to create exactly or observe directly, and scientists must instead form experimental settings that mimick the desired physical conditions as closely as possible. In these cases, the results must be carefully extrapolated to enable valid analysis.

The vast majority of scientific understanding of the physical universe has been obtained through the interplay between theory and experiment: an observation from an experiment leads to a new theory, which is then tested with a new experiment. For centuries, this process has been the fundamental approach to scientific exploration.

Reliance on Nuclear Testing

During World War II, the United States pursued the Manhattan Project, an intense program to explore the possibility of creating weapons based on the fact that nuclear fission liberates extraordinary amounts of energy from small quantities of nuclear material. The scale of the Project was astounding; it was easily among the largest, most complex, and most expensive of all human endeavors at that time. Its most important component was the secret Laboratory assembled by J. Robert Oppenheimer at Los Alamos, New Mexico, where many of the greatest scientists to ever live were gathered to explore how nuclear fission could be harnessed into an operational military weapon. Throughout the Project, Oppenheimer and his staff employed the scientific method, developing theory and testing it through experiment. Their efforts produced the Little Boy and Fat Man weapons that were used against Japan on Hiroshima and Nagasaki, respectively, leading to the end of the war in 1945.

For the next 47 years, the scientists in the U.S. nuclear weapons complex used theory and experiment as the primary means to develop insight about nuclear weapons. Throughout that period, the United States was embroiled in the Cold War against the Soviet Union and its allies. Early nuclear weapons were large, unwieldy, and difficult to handle; both sides in the Cold War worked intensely to create an arsenal of smaller, lighter, more powerful weapons that could be delivered by aircraft, missiles, or even by howitzer. Oppenheimer's group became the Los Alamos National Laboratory (LANL), and, given the feverish pace of the Cold War, was soon joined by organizations that would eventually become the Sandia and Lawrence Livermore National Laboratories (SNL and LLNL). Ultimately, the United States designed many different weapons types, leading to a stockpile comprising tens of thousands of operational nuclear weapons. Confirmation that the broad governing theory was reasonable was obtained when the weapons worked in full-scale detonations. The United States conducted more than 1,000 full-scale nuclear tests to ensure that its nuclear stockpile was both safe and operational[1], testing originally in the open atmosphere and, beginning in 1963, solely underground.

The End of the Cold War

By the early 1990s, the nuclear weapons created by scientists at the National Laboratories had become incredibly sophisticated. Weapons had to be powerful, yet compact and lightweight, so that a single inter-continental ballistic missile could deliver multiple warheads. Weapons had to be reliable, so that both the United States and its adversaries could be certain that the weapons would perform as designed if ever they were employed. Weapons also had to be safe from the unexpected: storage location fires, accidents in handling or transport of the weapons, temperature or weather extremes, etc. The Laboratories had to assure the government that these conditions were met, relying heavily on various forms of testing, including full-scale detonations.

The process of providing that assurance became extremely complicated by a series of events in the early 1990s that would shake the nuclear weapons complex to its foundation. With the collapse of the Soviet Union in 1992, the primary adversary for these weapons disappeared seemingly overnight. Later that year, President George H. W. Bush decided to halt the development of new weapon designs. On October 2, 1992, President Bush instituted a unilateral moratorium on underground nuclear testing.[2] In 1993, President William J. Clinton extended that moratorium and, in 1995, he endorsed the Comprehensive Test Ban Treaty[3] (CTBT). (Although never ratified by the Senate, the United States continues to adhere to the terms of the CTBT as of January 2009.)

Suddenly the U.S. nuclear weapons complex was faced with a new world—and a greatly complicated mission. Weapons still had to be maintained, and the government needed to retain confidence in the reliability, performance, and safety of the stockpile. But this would now have to occur without underground nuclear testing, the primary tool that had been used for nearly 50 years to design, test, and certify the weapons. With a ban on new designs and a ban on testing, the U.S. nuclear weapons complex would be forced to reevaluate its fundamental approach of using only theory, experiments, and full-scale nuclear testing to gain insight into the weapons. ASCI was born of this urgent need, of a commitment by the complex to meet the challenges of a post-Cold War, post-CTBT, world. To understand how ASCI succeeded, it will be useful to explore the tradition from which it arose.

Needed: a New Way of Doing Business

The U.S. nuclear weapons complex was established during the Cold War with a mission to design, build, and support the nation's nuclear stockpile. The very existence of the stockpile, together with convincing evidence of its reliability, provided the nuclear deterrence that was vital to national security through the Cold War years.

The complex consists of four major components. One is Defense Programs (DP), a part of the NNSA. DP is responsible for managing the entire nuclear weapons complex. A second component of the complex comprises nuclear weapons design and engineering, carried out at LANL, LLNL, and SNL. LANL and LLNL both focus on the nuclear components of the weapons, while SNL is responsible for certifying the non-nuclear parts, including such items as firing systems, radars, and the parachutes used to retard a weapon's fall. The third part of the complex consists of the plants that produce materials, manufacture parts, and assemble them into weapons for delivery to the military. The final part of the complex is the huge Nevada Test Site (NTS), located in the Nevada desert northwest of Las Vegas, where nuclear weapons testing was conducted from 1951 to 1992 and non-nuclear experimentation continues today.

All four elements have been essential to the safe, effective development, and maintenance of the American nuclear arsenal. A critical activity is the ongoing assessment of the condition of the weapons stockpile, which is the responsibility of the Laboratories. Annually, the three Directors of the Laboratories must certify to the President that the weapons are safe and will perform as expected.

But the weapons complex entered the 1990s faced with several serious challenges of a new era. These challenges included:

- The nuclear arsenal included weapons that were the smallest, lightest, and most powerful for their weight that had ever existed. Speaking in 2005, John Immele, LANL Deputy Director for National Security, described them thusly: "The nuclear weapons of our stockpile are the Ferraris of the nuclear-weapons age; they are small, they are difficult to manufacture, and in some cases their design margins are tight."[4] (By *design margin* we mean the difference between the most extreme conditions a weapon is designed to encounter, survive, and function properly, and the least extreme conditions that would cause it to fail.) Due to the ban on new designs, the weapons he described were the weapons that existed, and that the Laboratories would have to continue to certify, after the Cold War drew to a close in the early 1990s.

- The weapons would remain in the arsenal and be maintained much longer than originally anticipated.
- The moratorium on new weapons design precluded replacing the aging nuclear weapons with new ones.
- The facilities used to manufacture the weapons were changing, along with manufacturing technology; hence, many replacement parts could not be made identical to the original parts.
- Underground nuclear testing, the primary means of understanding the integrated, full-scale performance of the weapons, was prohibited.

Before 1992, all three Laboratories relied heavily on underground nuclear testing, and, although testing was just one of many tools employed by the Laboratories, it alone provided the opportunity to conduct experiments at full-scale and to observe how a weapon system would operate as an integrated whole. The test ban changed all that and left the nation's nuclear weapons scientists scrambling to find new tools to certify the stockpile. By analogy, imagine automobile engineers being told they could never again drive any car but must ensure that every car in their fleet would operate, at any time, as originally designed, and, by the way, win the Indianapolis 500. Nuclear weapons scientists still had theory and "subcritical" experiment (i.e., no self-sustaining nuclear chain reactions) in their tool kit, but without the ability to test the entire system end-to-end, they knew something else was desperately needed.

To complicate matters more, the Laboratories themselves were changing. During the 1980s, with President Ronald Reagan's emphasis on the Strategic Defense Initiative (SDI, also known as Star Wars), the Laboratories saw significant funding increases. But the early 1990s brought the end of the Cold War and an increased emphasis on federal budget deficit reduction, and the Laboratories were facing major budget cuts.

Enter Vic Reis

In August of 1993, Vic Reis became the U.S. Department of Energy Assistant Secretary for Defense Programs and took responsibility for the nuclear weapons complex. In a 2005 interview, Reis recalled, "My mission from Hazel [O'Leary, then the Secretary of Energy] was to downsize the complex and do it in an orderly manner. I also knew the President wanted the Test Ban Treaty. To do this, I realized that maintaining a strong laboratory system was vital." He went on to say, "I felt the strength of the Department of Energy and Defense Programs was in the Laboratories. The first thing I had to do was to save the Laboratories, and the CTBT provided the reason for doing that. To do that, I knew I needed them to get involved in big experiments, which is what the Laboratories do best. I also knew they had to do big computing because if you cannot do [underground nuclear] testing, this was the opportunity to do virtual testing."[5]

Science Based Stockpile Stewardship

In his plan to preserve the Laboratories while downsizing the complex, Reis had foreseen a new way to get insight into nuclear weapons. It would create new methods of understanding how nuclear weapons worked and of predicting and analyzing the factors affecting their performance and safety. The broad new plan became known as Science-Based Stockpile Stewardship, or SBSS.

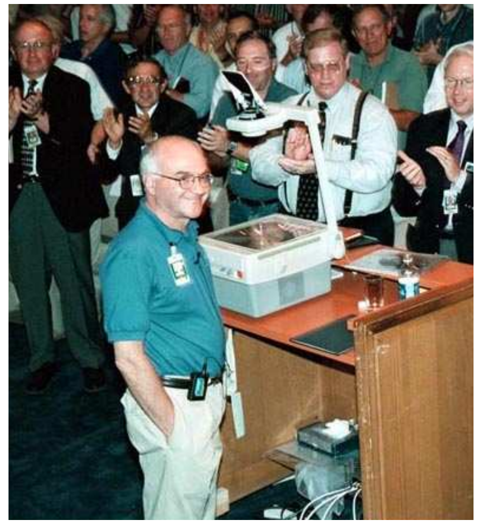

Figure 1-1. Vic Reis.

SBSS would consist of three major elements. The first was the construction of more powerful and precise facilities for conducting scientific experiments and subcritical tests. The second was an enhanced surveillance program to monitor the stockpile and detect changes in the aging weapons. The third element, ASCI, would develop computational simulation capabilities sufficiently fast, robust, accurate, and detailed to replace the data previously available only through testing. ASCI's goal was even more ambitious. It was to generate far more detailed information than testing could and enable scientists to examine the physics, chemistry, metallurgy, and dynamics as never before—sufficient enough, in fact, to enable scientists to achieve scientific insight.

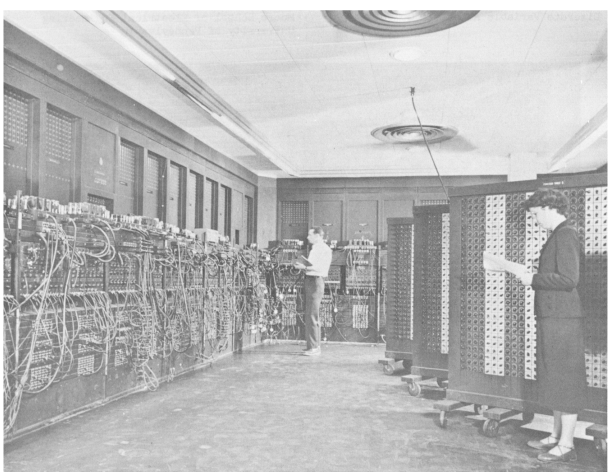

Figure 1-2. The ENIAC computer.

Computer simulations were nothing new. The Laboratories had used them for decades to support the design and development of nuclear weapons. ENIAC, among the world's first electronic computers, was used to test theories about the first hydrogen bomb design and calculate the bomb's expected explosive yield.[6] Throughout the Cold War, the Laboratories routinely acquired and employed each new "world's most powerful computer" to perform calculations that supported and refined scientists' theories. Scientists also used computers as tools to help predict and interpret experimental results.

Beginning in the late 1980s, DP and its Laboratories began to consider computer simulations in an entirely new way. A few visionary scientists, both within and outside DOE, looked at recent breakthroughs in the field of high-performance computing and hypothesized that with a massive effort to develop advanced new hardware and the programming to exploit it, computational simulation could become an equal to theory and experiment in a fundamental new approach to science. Forward-thinking DP leaders, with Vic Reis as their champion, believed that ASCI could amass the funding, leadership, drive, and brainpower to make that vision a reality and use it as a centerpiece of the "science-based" part of stockpile stewardship.

Delivering Computational Simulation Capabilities

ASCI's mission posed significant challenges. The biggest of these was the need to develop balanced *computational systems*. Such systems consist of:

- Simulation **applications**, that is, computer programs (codes) that solve the mathematical equations simulating physical events;
- Computing **platforms**, that is, machines with sufficient processing power, memory, and speed to run the applications and complete them in a reasonable amount of time;
- User **environments**, that is, the operating systems, editors, compilers, debuggers, data manipulation and analysis tools that allow programmers to develop and test the applications, and that enable application users to create and run models and then understand the simulation results.

At the start of ASCI, applications, platforms, and environments all existed at the National Laboratories; however, none of them was sufficiently developed, nor were they sufficiently integrated, to permit computational simulation to play the crucial role envisioned for it in the SBSS.

A New Approach Evolves

Throughout much of the development of computing, the machines had always consisted of a single processing unit connected to a bank of electronic memory and generally having some sort of long-term storage device attached (either tapes or disks). The computations were entirely serial, in the sense that the computer accessed the storage location of each data entry (number) individually, fetched it into memory, performed the mathematical operation, updated the value, and proceeded to the next piece of data. In the 1970s and into the 1980s, a technique known as *vector* or *array* processing was developed. In this approach, also known as *Single Instruction Multiple Data* (SIMD), many data entries were pulled simultaneously into the memory of special-purpose unit where the same operation was performed, very rapidly, on all the data, which was then written back out to the memory device. This capability to perform the same operation on many data entries, essentially at once, revolutionized computing and dominated the high end of computer performance for most of a decade. The Cray-1, introduced in 1976, was the first supercomputer to fully exploit vector processing (although not the first to use it) and led to a succession of generations of top machines throughout the 1980s, machines that became the mainstays of scientific computing both within and outside the Laboratories.

Still, the top computer scientists realized that memory and data bottlenecks, along with limitations on the size of vector processors, effectively limited the performance achievable with vector processing, and ASCI scientists knew that the limits were far below what would be necessary for reaching ASCI's goals. Fortunately, a solution seemed to be arising. Since the mid-1980s, a new approach to computing, known as *parallel processing*, had been evolving, and this approach offered the promise of unprecedented computational power.

Over the years, improvements had allowed computers to operate faster and use additional processor units, which lead to performance gains. While it had become possible to harness a few (on the order of tens) of processors together, the underlying architecture of one (or just a few) processors connected to a single common bank of memory remained standard. Exploiting an explosive growth in silicon microchip technology, the personal computers that emerged in the 1980s introduced much cheaper processing units and memory. Shortly thereafter, scientists began to explore the idea of harnessing these cheaper microprocessors together to create extremely powerful computers. At one point, LLNL's Eugene Brooks labeled this trend the "Attack of the Killer Micros." [7] As it developed, this type of computing became more commonly known as *massively parallel processing* (MPP).

In 1991, LANL's Frank Bobrowicz wrote a white paper entitled, "The Nuclear Weapons High-Performance Computing Paradigm for the Twenty-First Century," in which he described the current state of nuclear weapons simulation and computing at the Laboratories. He observed:

> During the past four decades tremendous progress has been made in our ability to understand many of the important phenomena associated with nuclear weapons. However, *a priori* computer simulations capable of predicting the true behavior of these devices require computing capabilities far in excess of what has previously been available. This is due to the complexity of the theoretical and mathematical descriptions governing the multitude of physical processes involved, the nonlinearity of the interactions between these processes, and the need to deal with all of these as a fully integrated system evolving in both time and space.[8]

Bobrowicz suggested that the time was right in the early 1990s to consider using computer simulations in a new way, writing that "the potential now exists for the Nuclear Weapons Laboratories to make a fundamental technological breakthrough into an era in which predicatively realistic numerical experiments can serve as a major driving force and guiding paradigm."[9]

Another champion of the emerging computing technologies in the early 1990s was LLNL's Randy Christensen. In late 1992, he prepared a presentation entitled, "The Numerical Test Site – A Conceptual Outline." In the presentation, he envisioned a fledgling computational capability growing into a powerful and important set of tools that would be crucial to maintaining the nation's nuclear deterrent. Several important points were made in the presentation:

> Computing and Communications Technology are in the early phases of a revolutionary jump in capability. The goals of the Numerical Test Site (NTS) are to exploit this jump and:
>
> 1. Provide a vastly improved numerical simulation capability that will, when combined with our existing Nuclear Test database and improved Non-Nuclear Experimental Testing facilities, greatly enhance the confidence with which the performance and reliability of our stockpile can be modeled.
>
> 2. Apply, through separate unclassified resources, the technologies and capabilities developed in the NTS to the solution of a wide-ranging array of commercially relevant problems, ranging from automotive design to the design of advanced materials.
>
> 3. Serve as a driving force for the development of High-Performance Computing in the U.S.—a place where the future of computing is shaped and tested.[10]

Christensen envisioned the NTS (deliberately named to use the same acronym as the Nevada Test Site) as something that would not only affect how Livermore conducted its primary mission of maintaining the performance and reliability of the weapons, but also transform the Laboratory as an institution. At one point in the presentation he wrote:

> This makes possible the development of an NTS that will not only permit DS [Defense Systems] to meet its own goals but will also once again establish LLNL as an institution capable of carrying out projects and programs too big to be done by Universities or Industry. It will aid in maintaining LLNL's unique role in the nation.[11]

At about the same time, both SNL and LANL were also considering how parallel processing could be applied to the nuclear weapons problem. The revolution in parallel computing provided real hope that Reis, DP, and the Laboratories could build a computational simulation capability enabling scientists to achieve insight.

Starting an Initiative

Actually achieving this goal meant that an initiative had to be organized. That task fell in the summer of 1994 to Air Force Lieutenant Colonel Steve Berggren, who was temporarily assigned to DP where he was managing part of the high-performance computing (HPC) activities at the nuclear weapons Laboratories. Berggren was instructed to develop a strategic plan for what would become ASCI. His report was published on July 8, 1994, and concisely summarized the challenges ASCI would face:

> Current state-of-the-art computational capability is not adequate to meet the needs of SBSS. Computer memory sizes and processor speeds do not allow the fine spatial and temporal resolution needed for high-fidelity predictions. In addition, existing computational physics models are not detailed enough to make reliable predictions without validation through nuclear testing. The improved models will require additional computer memory and processor speed. In all, the improvement in memory size and processor speed by factors of 1,000 or more will be required by the first half of the next decade. [12]

From the outset, Berggren recognized that success for ASCI would require, beyond revolutionary platforms and advanced applications, an integrated, balanced system including rich environments that would empower developers and scientists to fully realize the promise of the initiative:

> ASCI's principal goal is to assure that appropriate computational-based predictive capability will be available to meet the needs of SBSS. It will develop a fully functioning computing environment for design evaluation and agile manufacture. This includes the necessary computers, operating systems, compilers, debuggers, data management and analysis tools, data storage systems, and applications software.[13]

Early planning for ASCI also recognized that a fourth element was required: the Initiative needed to build partnerships with universities, other government HPC programs, and with industrial developers and producers of high-end systems. The ASCI Plan reflected this:

> Even though ASCI will provide a strong incentive to U.S. industry to accelerate the development of High-End Supercomputing, it can't do the job by itself. A key ASCI policy is the formation of strategic partnerships with other elements of the U.S. HPC community to accelerate the development of the needed technology. Only by working together can the various elements of the HPC community surmount the technology and economic barriers that stand in the way of the development of this incredibly powerful capability. [14]

The first step in establishing these partnerships was an outreach effort by Berggren and others on the DP staff. They made a series of visits to the manufacturers of the high-end systems and solicited white papers proposing ways to meet ASCI's goals.

A Strong, Shared Mission

The initial outreach effort culminated in a two-day workshop held in Santa Fe, New Mexico, in late September of 1994. Vic Reis opened the workshop with a presentation introducing the idea of SBSS and describing the key role that ASCI would play. A summary of Reis' remarks, prepared by LANL staff, reported several significant points about partnerships:

- ASCI is a joint venture; it can't be defined by one Laboratory, by Headquarters, or by industry.
- It is important that industry tell DOE and the Laboratories what they think are their future directions.
- The high end of the supercomputing industry was built to respond to defense needs and with support from the DOE weapon Laboratories, but with resources shrinking, action is necessary or the nation will be out of the supercomputing business.[15]

Reis also emphasized the importance to the United States of having a great Laboratory system, suggesting that a "great Laboratory" must embody three crucial characteristics:

- It has to work on an important national or world problem.
- It must have a real mission and a science and technology challenge.
- It must deliver a product and not be just a general science laboratory.[16]

Reis observed that all three characteristics were in ample evidence through the history of the Laboratories. The important world problem had been the threat of a nuclear war between the United States and the Soviet Union, and deterring that threat became the clear mission. There were significant scientific and technical challenges in designing nuclear weapons, particularly for light weight and high yield, and the Laboratories were always under pressure to deliver the actual product, the weapons themselves.

Noting the impact the end of the Cold War would have on the Laboratories, Reis suggested that in the future the DP Laboratories would still share those characteristics as they face challenges to reduce the nuclear danger (the world problem) by maintaining a stockpile (the mission) of safe, secure, reliable weapons (the product), all to be carried out without new designs or production, and without nuclear testing (the science and technology challenge).[17] Reis said, "All these factors point towards a strong push by Defense Programs in advanced computing. There needs to be a strong, shared vision of ASCI that all parties agree is important; then we can make progress together and achieve our goals."[18]

Reis then issued a challenge to his audience of DOE employees, Laboratory scientists, and computer company representatives. He said, "We have a 10-year window; if we do not have sufficient computer simulation capabilities by then, we will need to go back to testing and that will probably not be an option. We must succeed. The Laboratories will need to change to being experiment- and computer-driven, rather than test-driven as in the past."[19] He concluded thusly: "We want ASCI to be so good that others recognize that they should be aware of it and what we are doing and support it."[20]

ASCI would prove to be as good as Reis demanded. During the following decade, fueled by a sense of urgency felt throughout the DOE, fed by the budgetary largess of a convinced Congress, and powered by the energy, talent, and efforts of thousands of people across the nuclear weapons Laboratories, DP, other National Laboratories, universities, and industry, the program would achieve milestone after milestone. In addition to serving the

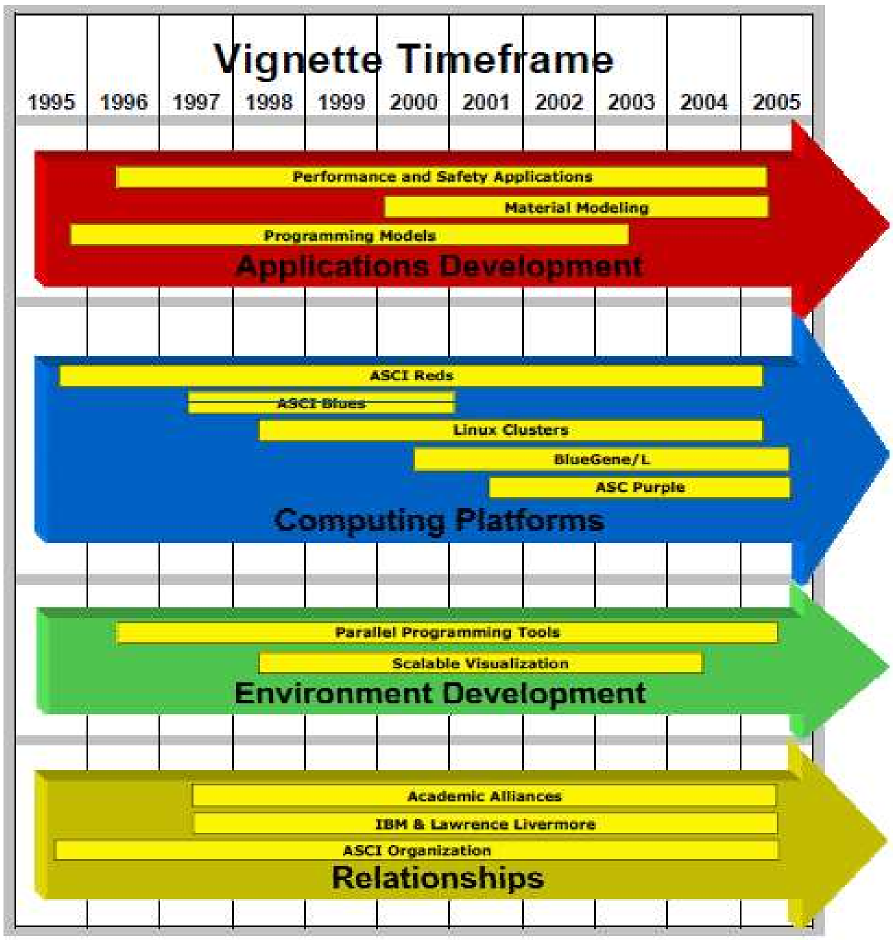

Figure 1-3. The vignettes and timeframes.

needs of SBSS, these successes marked a path leading inexorably to the realization that computational simulation can stand shoulder-to-shoulder as peer to theory and experiment in a new paradigm for scientific investigation.

Telling the Story

This history of ASCI is told in 13 vignettes describing ASCI's path to success. The Initiative was required to push technologies simultaneously along many different paths, engendering revolutionary approaches in all areas of applications, platforms, environments. This makes telling the story cleanly a difficult proposition, and as a result, even though that each vignette is narroly focused, it is well to recall that the vignette illustrates but a small piece of a huge effort. Figure 1-3 will help the reader understand the time covered by each vignette and what other technology development efforts were happening simultaneously.

The stories illuminate the wide array of technologies required make the case for predictive simulation as a major tool in stockpile stewardship, which was the primary goal of ASCI. But much more was accomplished, and these stories may appeal to those interested in using computational simulation to help illuminate the fundamental nature of many physical processes—especially to those who see predictive computational simulation as a new branch of science.

---

[1] U.S. Dept. of Energy, *United States Nuclear Tests July 1945 through September 1992*, DOE/NV-209-REV 15, xi.
[2] U.S. Dept. of Energy Defense Programs, *Above Ground Testing Capabilities Needed in the Absence of Underground Nuclear Testing*, 2.
[3] The White House, Office of the Press Secretary, *Statement by the President*, August 11, 1995, 3.
[4] Tollefson, "Part 1 of 3: Stockpile Stalemate," *Santa Fe New Mexican*, August 1, 2005, ,http://www.freenewmexican.com/news/2515.html.
[5] Vic Reis, interviews by author, telephone and at the Forestall Building, Washington, DC, April 21, and June 20, 2005.
[6] Rhodes, *Dark Sun: The Making of the Hydrogen Bomb*, 251.
[7] Brooks, E. D., and Warren, K. H., Eds, *The 1991 MPCI Yearly Report: The Attack of the Killer Micros*, Lawrence Livermore National Laboratory Report UCRL-JC-107022, 1991.
[8] Bobrowicz, *Numerical Experiments: The Nuclear Weapons High-Performance Computing for the Twenty-First Century*, 5.
[9] Ibid.
[10] Christensen, *The Numerical Test Site*, 2.
[11] Ibid., 5.
[12] Berggren. *Accelerated Strategic Computing Initiative: Strategic Planning Document*. July 8,1994.1.
[13] Ibid., 1.
[14] Ibid., 4.
[15] U.S. Dept. of Energy Los Alamos National Laboratory, *Summary of ASCI Workshop*, 1.
[16] Ibid.
[17] Ibid., 2.
[18] Ibid., 3.
[19] Ibid.
[20] Ibid.

Chapter Two
Building the Initiative

Initiative is defined as "the power or ability to begin or to follow through energetically with a plan or task."[1] Generating an initiative to build a predictive simulation capability proved to be a complicated endeavor, requiring major technological advances in numerous disciplines; moreover, ASCI's need to accomplish it in a decade meant that dramatic improvements in diverse technologies had to be made simultaneously. Moreover, to fulfill the promise of ASCI it would be necessary to integrating these improved technologies into a useable operational system for the stewards of the nuclear weapons.

Needing a Team

First and foremost, ASCI needed a team capable of bringing the concept to fruition. The most powerful computers in the world were neither big nor fast enough at the onset of ASCI. Even had they been, however, merely having the platforms would not suffice for ASCI unless suitable applications also existed. Similarly, having both platforms and applications would be of limited value without a means of understanding the results. ASCI would need a wide variety of specialists, including nuclear weapons scientists, weapons engineers, mathematicians, material scientists, computer architects, computer scientists, networking engineers, security specialists, system administrators, and procurement specialists, among others. As the program progressed and the size and power requirements of the computers became better understood, building architects, electrical engineers, and ventilation and air conditioning technicians would augment the list of required specialties. All these specialists would need to work very closely together, inducing ASCI to place tremendous emphasis on having truly interdisciplinary working teams. This approach was not unknown, but was then a bit unusual in science. Since cutting edge research, whether at Laboratories, universities, or industry required world-class understanding of the research topic, scientists tended to focus on and to organize around a specific topic area. At about the time ASCI was beginning, the nation's science Laboratories and academia were beginning to see the value of organizing projects into interdisciplinary teams rather than into disjoint, technically deep units, each taking care only of its piece of the project. ASCI planners, seeing the exceptional complexity and tight deadlines of the initiative, became early leaders in promulgating interdisciplinary science as an organizing paradigm.

Collegial Competition and External Collaboration

The traditional approach to conducting science would provide another challenge for organizing ASCI. Simply put, research institutions are competitive. This is natural and, at least usually, a good thing. For instance, institutionalized competition among the Laboratories provided the opportunity for peer review, an essential and common method of checking the validity and value of scientific hypotheses. The idea behind peer review is that scientists with equally deep understandings of particular areas review each other's work. These reviews often involve much discussion and study and the replication of experiments to check results. Peer review of nuclear weapons science presents a special challenge because the science is largely classified and not available for review by the general scientific community.

Lawrence Livermore National Laboratory was founded in 1952, during the early days of the Cold War, spurred by an urgent need to speed the development of the hydrogen bomb. The idea was that a healthy competition between Los Alamos and Livermore would accelerate the design process. The two Laboratories would also provide each other the independent peer review otherwise lacking within the classified environment.

LANL and LLNL were encouraged to compete with each other and have done so throughout their history. When new weapons were to be produced, design competitions were held between the two Laboratories. A third weapons laboratory was also founded (originally an offshoot of LANL’s Z Division) and was operated independent of the two design laboratories. Named Sandia National Laboratories (SNL), it specialized in the engineering aspects of the weapons, supporting the designs of both LLNL and LANL.

The competitive spirit among the Laboratories had served the nation well, but the technical challenges ASCI faced were more than any single institution could surmount. ASCI needed a cooperative team and needed to overcome the legacy of independence and competition.

Not only would ASCI have to change how the Laboratories interacted, the Initiative would have to involve outside organizations. This was atypical, since most DP activities were kept “inside the fence,” for very good reason. Over the years, the entire complex had been developed to control the design, manufacturing, maintenance, and the non-explosive disposal of the weapons. Within the complex, information about nuclear weapons was so sensitive that great effort was expended, and much infrastructure was built, to preserve its secrecy. Although there was some interaction with outside suppliers, it was kept to a minimum. Similarly, while there was some collaboration between the Laboratories and academia, it was generally rather limited and carefully constrained to protect secrets and sensitive information.

Speaking at the September 1994 workshop in Santa Fe, Vic Reis pointed out that ASCI would have to be much more open. Because the success of the Initiative was highly dependent on the computer manufacturers to deliver the necessary systems, ASCI would depend critically on establishing and maintaining close working relationships with computer manufacturers. Many of the technologies ASCI needed would have broad application (and indeed could be developed) outside the fence. Much of the science that would be critical in the simulations was neither classified nor sensitive and could be provided by partnerships with academics. Eventually, it was decided that ASCI would rely heavily on commercial sources, other National Laboratories, and universities.

Timing

Reis’s sense of urgency was well justified. He knew that DP would need the new simulation capabilities in about decade. This timing was based on two critical factors. First, the nuclear weapons in the stockpile were aging. Many had already exceeded the lifetime envisioned by their designers. All materials change somewhat with age, but nuclear weapons contain highly radioactive matter. How that radioactivity would affect the components of the weapon over long time spans was not well understood. Without testing, it would fall to simulation to provide the information scientists would need to develop insight into how those changes would impact the performance, safety, and reliability of the weapons. The second factor was that not only were the weapons aging, so were the scientists and engineers who possessed underground nuclear testing experience. These people were crucial to validate that ASCI simulations could be used as an alternative to underground testing, a fundamental tenet of SBSS. But many of these people had retired, and most of those who hadn't would do so soon. It was imperative that ASCI be ready for validation before they all left. Hence, when ASCI was started in 1995, a deadline of 2005 was set for development of an initial full-scale simulation capability with enough resolution, sufficiently accurate modeling of the physics, adequate speed, and the reliability necessary to enable true scientific enlightenment.

Even as this decade-long goal was being set, a much shorter-term deadline threatened the beginnings of ASCI. The budget of the federal government is developed and enacted into law through a process requiring about two years. It begins with the formulation of the president's budget request to Congress. In September 1994, as the Santa Fe Workshop took place, the White House Office of Management and Budget was in the process of preparing the request for Fiscal Year 1996. It was critical for ASCI to enter into that process, or face the possibility of funding being delayed another year.

Challenges of Starting ASCI

As he gave his ASCI talk at the Santa Fe Workshop, Reis was facing a host of issues in starting the Initiative. These included:

- The need for a cross-disciplinary approach to creating detailed, accurate computational simulations, despite a tradition of "stovepipe" research at the National Laboratories;
- The need to get the National Laboratories to work together, despite the legacy of operating independently and competitively;
- The need to get industry and academia involved, despite a tradition for the weapons programs to operate primarily inside their fences;
- The need to develop the simulation capabilities quickly to address issues with aging weapons and aging scientists and engineers; and
- The need to act quickly to get ASCI into the federal budgeting process.

To meet these organizational challenges, Reis would rely on a man then sitting with his legs propped up in the back of the darkened Santa Fe conference room: Gil Weigand.

Gil Weigand

Weigand received his Ph.D. in engineering from Purdue University in 1978 and spent his first professional years as a researcher for Westinghouse and General Motors. He then moved to Albuquerque to work at SNL and began working on massively parallel computing systems. In the late 1980s, Weigand was assigned to work at the Defense Advanced Research Projects Agency (DARPA), and it was at DARPA that Weigand met

Reis. DARPA had a budget of billions of dollars and a mission to support advanced technology research for the military. While at DARPA, Weigand supervised several projects focused on high-performance computing. In the early 1990s, Weigand left DARPA to work briefly at the White House and then moved back to SNL.[2] It was there, in late 1994, that Reis asked him to take on ASCI's leadership.

Figure 2-1. Gil Weigand.

From their experiences at the National Laboratories, DARPA, and the White House, Reis and Weigand knew that if ASCI was to be successful, innovative management would be essential, and ASCI would have to reach out to four primary communities: the internal management of DP and the scientists and engineers at the weapons Laboratories, the companies comprising the HPC industry, the non-nuclear weapon research community (other U.S. National Laboratories and universities), and, with great care, the United States Congress.

Organizational Structure

Initial planning for ASCI was underway well before Vic Reis spoke at the Santa Fe workshop and outlined the challenges ASCI would face. All three nuclear weapons Laboratories would be critical participants in ASCI; hence, all were involved in the early planning. It became clear that many of the technologies needed could be used across the nuclear weapons complex, and it was also clear that the mission would be beyond the capabilities of any one Laboratory. By early summer of 1994, Steve Berggren had solicited Laboratory input on how ASCI should be managed and had incorporated their ideas into the *ASCI Strategic Planning Document*. This document, released on July 8, 1994, described the basic philosophy of how ASCI would be managed over the next decade, saying:

> The goal of the ASCI management process is to assure a focused and coordinated effort. This will be particularly difficult since all three Laboratories and various hardware and software vendors will be developing the essential pieces of ASCI that must then work in concert to produce the desired computational capability. Under these circumstances, effective working level communications among the projects is essential. Strongly focused guidance and leadership is also essential.[3]

The planning document also described the initial basic organization for the Initiative. Individual ASCI projects would be executed by "virtual" teams with members from all three Laboratories and from other agencies and industry as needed. The projects would be integrated and coordinated by three working groups, focused respectively on Platforms, on Applications, and on Infrastructure. A "Board of Directors," (later named the

Executive Committee) would include representatives from each Laboratory, industry, DOE Headquarters, and other government agencies, and would be chaired by a DP representative. This board was to provide direction and oversight for the entire enterprise. The management approach described in the original *ASCI Strategic Planning Document* would become known as "One Program – Three Laboratories," and was formalized in the 1996 *ASCI Program Plan* with the following organizational chart:

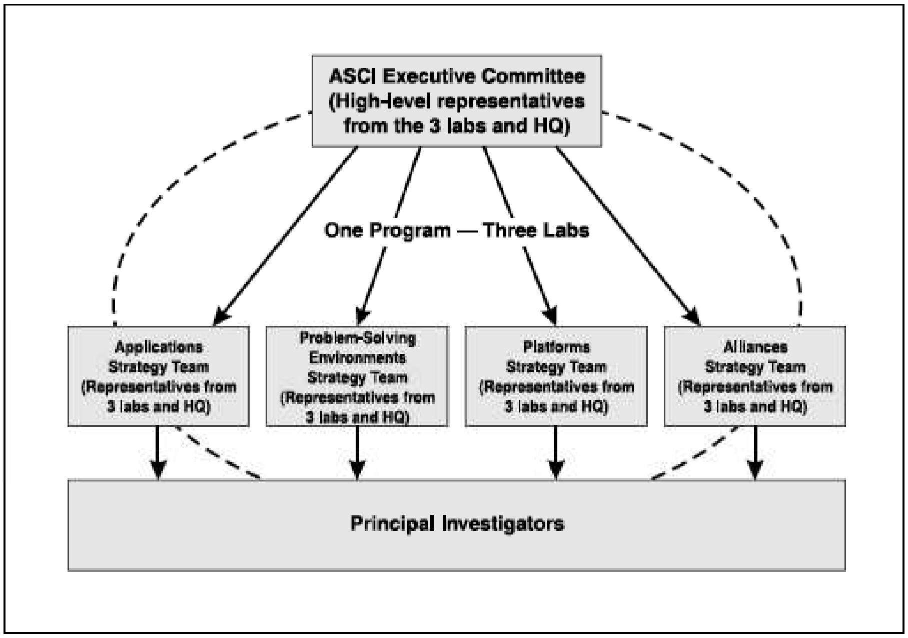

Figure 2-2. ASCI organization plan.

This structure accomplished several things for the Initiative. It put the three Laboratories at the same table to plan the direction and budgets for ASCI, which dampened the sense of competition between them. While competition would not entirely disappear, ASCI's management structure provided an avenue for *healthy* competition and fostered a sense of shared mission. The ASCI management structure also brought DP Headquarters staff into the discussion, meaning that Headquarters staff would not be involved only through their traditional budgeting and oversight activities; they would be full participants in selecting, planning, and managing the technical activities.

The unified management approach also created a solid focal point for the Initiative. One of the first members of the ASCI Executive Committee was LLNL's David Nowak, who remarked in 2005, "I think [the "One Program – Three Laboratories" strategy] was one of the real successes for the program. The committee structure was very effective and created a single vision and voice in front of Congress."[4]

When Weigand joined ASCI in late 1994, he set out an aggressive schedule that would cement the Tri-Lab team together. Preparations for one of the first ASCI presentations made to Reis, in February 1995, required a series of major strategy meetings spaced over just a few weeks. These meetings were held in Santa Fe, NM; Crystal City, VA; Berkeley, CA; and Washington, D.C., and accomplished three important things.

First, they established the initial budget plan for the program. Second, amid very vigorous discussions, the technical approach for the initiative was refined. Third, this period of intense activity forged a tightly knit team that would be able to tackle the challenges ahead.

Committing the Laboratories

Among the most daunting obstacles would be getting the Laboratories themselves to accept that ASCI could create a new means of developing scientific insight. Representative Floyd Spence, then chairman of the House National Security Committee, issued a report in October 1996, entitled *The Clinton Administration and Nuclear Stockpile Stewardship: Erosion by Design*, in which two of the Laboratory Directors are quoted describing the challenges ASCI faced:

> The Director of Sandia National Laboratories, C. Paul Robinson, testified before the House National Security Committee on May 12, 1996, that, "the commercially available and laboratory technologies are inadequate for the stockpile stewardship tasks we will face in the future. Another hundred-to-thousand-fold increase in capability from hardware and software combined will be required." Furthermore, "Some aspects of nuclear explosive design are still not understood at the level of physical principles," he stated. This statement alone raises questions about whether it is possible to simulate these particular phenomena through computer models.[5]

> The Director of LANL, Siegfried Hecker, testified on March 12, 1996, that, "In general, future stockpile assessments will require three-dimensional calculations, which in turn need 1,000 times the computing memory and would take 100 years to perform on current machines." Hecker said that implementing the Clinton plan requires developing "computers and their supporting software ten-thousand-fold more powerful than the largest machines readily available today." The program to develop such capabilities, known as the Accelerated Strategic Computing Initiative (ASCI), must be accomplished in one decade rather than the three decades, which would normally be required.[6]

The comments by Robinson and Hecker reflect the general skepticism weapons scientists and engineers felt about using computational simulations in an entirely new way.

It was essential to get the Laboratory Directors and their technical staffs to believe in predictive computational simulation and commit to the ASCI mission. A crucial task would be demonstrating that ASCI's managers understood the deep scientific and engineering issues facing Laboratory weapons designers and analysts, and an important step in that direction occurred when Nowak joined the Executive Committee. Nowak was a weapons designer; he had the credentials to talk with weapons scientists in detail about their work. Earlier in his career he had acted as the representative between one of the LLNL nuclear weapons divisions and the computational physics department. That experience provided him with an appreciation of the limitations, and the potential, of the ASCI approach. Shortly after Nowak came on board, John Hopson of LANL also joined the ASCI Executive Committee, bringing similar credentials. Nowak and Hopson helped overcome a great deal of skepticism about the program from the Laboratory Directors and the weapons scientists and engineers.

Turning to his DARPA experience, Weigand introduced into ASCI the "Principal Investigator (PI) meeting," an element of ASCI that would do much to win over the skeptics. Over the years, DARPA had developed the concept of PI meetings to encourage communication among its researchers; the meetings provided an opportunity for everyone to see, learn from, and critique what the other DARPA-funded researchers were doing. It also allowed DARPA project managers to evaluate performance on the spot and make technical adjustments as necessary.

With its "One Program – Three Laboratories" structure combined with DP participation, the PI meetings fit ASCI extremely well, offering a forum where ASCI researchers could share results and discuss future directions. Equally important, given the uncharted terrain ASCI was navigating, the PI meetings afforded ASCI scientists and engineers an opportunity to voice their concerns about the Initiative.

The ASCI PI meetings were initially held biannually and generally occurred in interesting locales, chosen to attract ASCI scientists and potential end-users to the discussions. For example, ASCI PI meetings were held in Las Vegas, Nevada; at Snowbird ski resort in Utah; and at the Kennedy Space Center at Cape Canaveral, Florida. One meeting, in Bangor, Washington, featured a tour of a Trident submarine; the meeting at Cape Canaveral featured the launching of a satellite into space, while a PI meeting held at Vandenberg Air Force Base provided the opportunity to witness an MX missile launch.

The ASCI structure, the inclusion of knowledgeable scientists in ASCI management, and the free discussions of the PI meetings all helped, but ultimately, it was the actual delivery of capabilities that convinced the Laboratory Directors, weapons scientists, and engineers that ASCI's new approach to understanding nuclear weapons could work. The key to ASCI's success would be if it could simultaneously develop all the components of the complete system: the massive hardware, the applications software, and the computing environment, that would enable scientists to arrive at new insights.

Bruce Tarter, Director of Lawrence Livermore National Laboratory at the start of ASCI, was asked in 2005 what his initial thoughts about the Initiative had been. "I was very skeptical," he said. "It was a totally new approach to assessing the stockpile. Also, Livermore had not had very good experiences with their recent computer procurements." Tarter was describing an early program known as the Massively Parallel Computing Initiative (MPCI). This program had purchased the BBN Butterfly and the Meiko Massively Parallel Processor, two early versions of MPP technology. Both were considered to be partial successes. While neither quite lived up to its full potential, both offered early and valuable experience with MPP platforms. This experience became the foundation for many of LLNL's later achievements under ASCI. Asked what he thought about ASCI's impact on LLNL, Tarter said, "It has been great. The program has had a tremendous impact on the Laboratory. It has changed the way people coming to the Lab think about doing science. The effects of ASCI are very broad." [7]

Partnering with Industry

The fledgling ASCI had its organizational structure in place and the Laboratories and DP were coming on board. Next, ASCI needed to engage the HPC industry. This community was critical because they would eventually have to create and deliver new, advanced computers with capabilities orders of magnitude beyond any computers then in existence. While the Laboratories had, since their inception, generally purchased and used the most modern and powerful computers in existence, they had acquired whatever the vendors were developing at the time. ASCI would have to change that process dramatically, to a new paradigm in which the Laboratories would work intimately with the vendors, not only specifying the capabilities of the machines, but participating directly in all phases of the design. Berggren began efforts to engage the computer industry during the summer of 1994, asking them for the white papers that were the focus of the September 1994 Santa Fe workshop.

Of serious concern to ASCI was that the computer industry was very unstable in 1994. Representatives of nine companies presented papers at the Santa Fe meeting: Cambridge Parallel Processing, Convex Computer Corp., Cray Computer Corp., Cray Research Inc., IBM, Intel SSD, Meiko Scientific Inc., nCUBE, and Thinking Machines Corp. By the following March, virtually all had gone out of business, were in financial trouble, or were receiving the majority of their revenue from selling products other than high-performance computers.

To get the computer companies committed, Reis, Weigand, and ASCI's staff conducted many high-level visits to the vendors, including stops at Intel, IBM, Convex Computing, and Cray Research. Reis mostly met face-to-face with CEOs, explaining the importance of SBSS and the function of ASCI. He explained the crucial role computer companies would have in making ASCI successful. The CEOs listened. Lou Gerstner of IBM, for instance, heard Reis's presentation and then acknowledged the national importance of the stockpile stewardship problem, and of ASCI. Gerstner stated that IBM would support, in whatever way it could, the development of the required computing platforms. Reis and Weigand received similar endorsements from the other computer manufacturers.

Ties to the Broader Research Community

ASCI would also need to forge strong ties to researchers in academia, industry, and National Laboratories outside of DP. While the DP Laboratories held essentially all the knowledge that existed in the United States regarding nuclear weapons physics, accurate simulation of that physics would require deep understanding more fundamental high-energy physics. Much of that deep expertise of more general atomic physics was held in academia or industry, and it was essential that ASCI reach out to the communities holding that expertise. The Initiative did this in several ways.

One of the first steps was to have ASCI reviewed by the JASON, an august body of mostly university-based scientists hired by the federal government that meets each summer to study scientific issues. The JASON (a term that denotes either the entire body or an individual member of it) held the review of ASCI in the summer of 1996. Representatives of the DP Laboratories presented the technical challenges of developing computational simulation as a new means of conducting scientific investigation. Following their standard procedure, the JASON considered the issues at the review and throughout their summer term, eventually preparing a report of their findings. The overall conclusion was that ASCI was on a good path to develop the required simulation capabilities. The JASON also offered several recommendations regarding the technical execution of the work.[8]

ASCI planners understood that it would be extremely difficult to engage the universities and non-DP National Laboratories without making funding available. Hence, the design of the Initiative included the "Alliances Strategy." This Strategy would fund a number of universities to conduct research and development of technologies needed to create a predictive simulation capability. The technologies addresses in the Alliances Strategy would be developed for specific problems or topics that were unclassified and not directly related to nuclear weapons.

The other National Laboratories that ASCI engaged were funded by diverse agencies including the Department of Energy Office of Energy Research, the Department of Defense, the National Aeronautics and Space Administration (NASA), and the Department of Commerce. ASCI used several means to engage these non-nuclear Laboratories in the Initiative planning and technology developments. Representatives of the non-DP laboratories were invited to PI meetings, for instance. ASCI also co-sponsored workshops, such as the PetaFlops workshops, that drew attendees from the non-DP laboratories. In some cases, ASCI directed funding to these other National Laboratories to pay for technology developments.

Obtaining Funding

Fulfilling ASCI's mission would not be cheap. One of the toughest challenges in launching ASCI was getting support for funding from either the president's administration or from Congress as those entities thrashed out the federal budget. It was especially difficult in early 1995. At the midterm elections the previous November, the Democratic party had lost the majority in the House of Representatives, in large part because the Republicans took a strong stance against deficit budgets. President Clinton responded on February 6, 1995, submitting a very lean budget to Congress. The President's Science Advisor, John Gibbons, summarized the cuts by saying, "There is no question that this is a very tough budget, and is based on a clear sense of priorities. The President's FY 1996 budget will terminate 130 programs, consolidate 271 others and cut back on many more, dramatically restructuring agencies to achieve total cuts of $144 billion."[9] Amid the carnage of budget cuts, there was a little noticed line for something new in the President's request. It was $45 million to fund ASCI.

Through the unsettled summer and fall of 1995, as for many other government programs, ASCI's budget waxed and waned. At the time, the ASCI staff at DP Headquarters prepared an e-mail newsletter entitled the "ASCI Team Report" that described how the various bills were progressing through the congressional process. One of those newsletters, published June 30, 1995, provided a nice summary of the turmoil of the budgeting process:

THE GOOD, THE BAD AND THE UGLY BUDGET NEWS

There was considerable action on the DOE and DP budget over the last week. Some of the news was good for DOE, but most of it was at least not great for ASCI. The good news centered on the Republican compromise on the budget resolution. As you will recall, the House voted to kill the Departments of Energy, Commerce and Education, while the Senate only wanted to kill the Department of Commerce. The President submitted a budget proposal that did not kill any of them. As a result of the President's budget "re-"submission, Senator Dole and Speaker Gingrich decided that quick action was required on the budget resolution. They pulled the Senate and House conferees together and forced a compromise that killed the Commerce Department but kept Education and Energy. Other good news on the future of DOE front also includes a very lukewarm reception of the Republican freshmen's plan to kill DOE.

Now for the bad news. We have now received three of the four mark-ups on the DP authorization and appropriation bills. As you will recall the House National Security committee mark gave ASCI an additional $40M and cut the TTI [Technology Transfer Initiative] budget to $25M. That authorization bill has passed the full House with no changes to the ASCI or TTI budgets. (that was also good news, now for the bad). The House DP appropriation committee marked up their bill and did not give any additional funds for ASCI and also cut TTI to $25M. The Senate Armed Services committee met on the Defense Authorization bill. The result of their mark-up is still unclear, but it appears that it did not provide any additional funds for ASCI. The mark-up also cut the TTI budget to zero. The Armed Services committee also strongly questioned the DOE science based approach to stockpile stewardship. Initial information indicates that the "Committee believes the only near-term alternatives to this are underground testing and hydronuclear experiments."

The only remaining mark-up on the ASCI program is the Senate appropriations committee which is chaired by Senator Pete Domenici. They are scheduled to meet on the week of July 10th and with any luck we will have better news from them. Once the mark-ups are made and the authorization and appropriation bills pass, we will have to see how the bills are reconciled. I expect however, we will have a much better understanding of where we are going in the budget after the Senate Appropriation mark-up.[10]

After a bruising budget battle that included a brief shutdown of the federal government, an Omnibus budget was eventually passed. ASCI had received $85 million for FY 1996, nearly twice the amount the President's budget had requested. This became a trend; for several years ASCI received more from Congress than was requested by the President. ASCI's budget rose steadily, from $85 million in FY 1996 to $747 million in 2001. One reason for this was that each year ASCI offered Congress a strong, clear vision of what it would accomplish, how it would be done, and why the work was important to national security and the nuclear weapons stockpile. The "One Program – Three Laboratories" management approach was instrumental in creating that vision and communicating it effectively.

Flexibility

The "One Program – Three Laboratories" management provided flexibility to adjust the ASCI approach as need arose. This was particularly evident in the spring of 1997, when three deficiencies in the original ASCI planning became apparent; their solutions highlight the adaptability afforded by the ASCI management approach. The original planning for ASCI had underestimated the technologies that would be needed to interpret the huge amounts of data that would be generated in simulation applications running on powerful computer platforms. In addition, when the Initiative was started, ASCI did not fully understand the challenges of supporting remote users of the MPP computers and applications. Also, as a result of communications with the weapons scientists and engineers, the Defense Program and Lab managers of ASCI realized they needed to devote more funding to the "verification and validation" (V&V) of the simulations. ASCI's flexibility in managing its budgets, manpower, and programs allowed each of these shortcomings to be addressed in a timely fashion.

Creating the tools to store, visualize and understand simulation results would become the province of an ASCI Strategy designated NEWS, for Numerical Environment for Weapons Simulation. The title was later changed to Visual Interactive Environment for Weapons Simulation, or VIEWS. The issues surrounding the use of distributed computers and remote use by ASCI collaborators gave rise to another Strategy, denoted DISCOM2, for Distributed Computing and Distributed Communications. DISCOM2 was particularly important because ASCI originally planned to procure only one of each "latest generation" computing platform at a time, and its location would be rotated between the Laboratories. That meant that scientists at two of the Laboratories would have to access the platform remotely (as would any academic or industry partners, should they be given access to the machine). Finally, a Verification and Validation (V&V) Strategy was added. Simply put, *verification* of a simulation seeks to assure the scientist that the code is solving the underlying mathematical physics equations correctly, while *validation* seeks to assure the user that the right equations are being used. It is all too easy to create a simulation that incorrectly solves the right problem or correctly solves the wrong problem—but verifying and validating the simulation is not easy at all. ASCI learned early on that it needed to mount an effort to assure V&V.

Adding these Strategies became part of a "re-baselining" of the program during the summer of 1997, an activity in which each program within ASCI justified its budget starting from the ground up (a zero-based budget).

With the completion of the 1997 re-baselining, the elements for ASCI were in place to create an unprecedented computational simulations capability, which would eventually enable ground-breaking scientific discovery. The technology ASCI would need was to be created by following the original ASCI Strategies:

- Apps (focusing on advanced Applications);
- Platforms (designed to produce the MPP high-end computing machines);
- PSE (creating an effective infrastructure and problem-solving environment);
- Alliances (engaging academia, non-DP Laboratories, and industry).

These were later supplemented with three Strategies from the 1997 re-baselining:

- VIEWS (for manipulation, display, and analysis of the massive data sets);
- DISCOM2 (addressing the problems of distributed and remote computing);
- V&V (insuring that ASCI was solving the right problems, and solving them correctly).

Keys to Success

In addition to the technical elements of the Initiative, ASCI fielded a strong team of DP and Laboratory managers to drive the Initiative. There was another set of strategies that went beyond technical development of the required capabilities to express to express the organizational culture and philosophy, and to guide how the Initiative would operate. These strategies included:

- One Program – Three Laboratories;
- Openness.

These strategies, and the characteristics they engendered, were critical because over the next decade ASCI would be faced with significant challenges on the technological, organizational, and political fronts. ASCI's success was keyed by strong vision and leadership, enabling the Initiative to deal with issues on all fronts simultaneously.

To fully appreciate the history of ASCI, it must be understood that the story depends on special people—the ASCI management team and the many people who worked on the Initiative—and on special traits and characteristics these people embodied. Among these were:

- **Leadership —** ASCI's success came because many people demonstrated exceptional leadership in many areas. Leading from the very top, Vic Reis and Gil Weigand stand out in particular. There is general agreement that without their guiding vision and fortitude, ASCI would not have accomplished what it has.

- **Mission —** ASCI's unfailing focus on supporting the nuclear weapons mission stands out as a major asset. This mission clarity always provided a set of criteria to help guide decisions about the direction of the program. Leaders were continually asking, "How does this support stockpile stewardship?" As a result, ASCI personnel at all levels remained conscious of, and focused on, the ASCI mission. A crucial part of this is that the ASCI workers *believed* in the importance and necessity of the mission.
- **Partnership —** The story of ASCI is a story of partnership on many levels. The history shows that the program worked particularly well when partnerships went well.
- **Endurance** — ASCI's people exhibited incredible endurance. Many hundreds of people toiled long and hard to support all manner of technical and programmatic projects. Endurance was also exemplified by those responsible for funding and supporting the Initiative. Beginning with the President, supported by Congress, and ultimately by executed by DOE, NNSA, and DP, ASCI remained a high national priority as it developed and became successful.

The vignettes that follow illustrate how the men and women of the Initiative embodied all of these traits. Thanks to them, ASCI was not only a great technological success, but also a great programmatic success. The road to the creation of a predictive simulation capability was a road fraught with myriad technical, organizational, budgetary, and political pitfalls. The history of ASCI is the story of how that road was navigated.

---

[1] American Heritage Dictionary of the English Language, Fourth Edition, 2000
[2] U.S. Dept. of Energy, "Bio for Gil Weigand".
[3] Berggren, *Accelerated Strategic Computing Initiative: Strategic Planning Document*, 16, attached to Marshall M. Sluyter, memorandum to Vic Reis, July 8, 1994.
[4] David Nowak, telephone interview by author, June 17, 2005.
[5] Spence, *Clinton Administration and Nuclear Stockpile Stewardship: Erosion by Design*, 1996.
[6] Ibid.
[7] Bruce Tarter, interview by author, Building 111, LLNL, August 3, 2005.
[8] JASON, *Simulation for Stewardship*, JSR 96-315, 42-46.
[9] Gibbons, "FY 1996 Science and Technology Budget: Press Briefing"..
[10] Larzelere, e-mail to ASCI, July 1, 1995.

Chapter Three
Applications – At the Heart of Delivering Insight

The application code is at the heart of predictive computational simulation. Also referred to as the app, the code, or simply the program, the application performs the calculations solving the equations of the physics used to describe the behavior of the physical world. ASCI's applications codes have steadily grown more sophisticated, detailed, and complete. While the applications at the start of ASCI were limited in size, scope, and detail, the steady improvement in code accuracy, resolution, and refinement, coupled with the enormous increases in capability of the big ASCI computers, have enabled ASCI scientists to make astonishing discoveries. ASCI applications have provided users with information so detailed and accurate that fundamental scientific discoveries have been made, on a par with the revelations of theory and experiment. In some cases computational simulation offers significant advantages, supplementing the traditional theory-and-experiment approach with relatively inexpensive simulations that can be used to interpret, guide the design of, or replace expensive or even forbidden experiments.

A well-built simulation sometimes exposes complex physical phenomena not anticipated by theory or captured by experiment. A particular advantage of simulations is that they produce readily accessible data; simulations can provide data about very small (at the atomic level) or very large (on the surface of a star) physical events that experimental diagnostic instrumentation cannot capture. One example is the solidification of molten metals. This process cannot readily be observed in detail because the metals are so hot, so dense, the process so rapid, and the important features are so small. But a well-built, powerful simulation code enables scientists to study the details of solidification.

The applications vignettes in the following sections relate how the people and organizations at the National Laboratories succeeded in building these powerful scientific tools. To put their accomplishments in proper context, it may be useful to describe what is involved in developing simulation applications.

Elements of a Simulation Application

Regardless of what is being simulated, four important elements of a simulation application are necessary to accurately represent how physical objects and systems react under a variety of conditions. The conditions being simulated could include both normal operating conditions and extreme conditions, such as might occur during accidents. The four necessary elements are:

- A detailed understanding of the underlying physical processes;
- A means to represent those processes numerically;
- The ability to accurately capture the numerical representations in a computer program; and
- The ability to validate that the simulation accurately represents the physical world.

The physical systems ASCI had to study were nuclear weapons, with the study conditions ranging from the passive aging of materials over decades, to the workings of the electromechanical devices, to the behavior of the nuclear materials that create conditions similar to those found inside a star.

As ASCI began, the Laboratories faced significant challenges in creating applications to capture these complex physical processes. At the time, it was easy to see that the world's most powerful computers would not complete the simulations in a reasonable time (that is, within days or even weeks, rather than years). This meant that to succeed the simulations would have to be built for computers that did not yet exist. These would have to be MPP systems, as no single-processor or even a parallel machine with just a few tens or hundreds of processors could handle the size problems that would be necessary. Thus the applications would have to be split into many hundreds or thousands of smaller pieces to be run on computers using many distributed processors. Finally, the simulations would have to be checked to ensure that they were operating correctly and that the results actually reflected the physical world.

At the beginning of the Initiative, many had serious doubts as to whether accomplishing ASCI's goals would be possible.

A Simple Example

While the simulations created for ASCI are among the most complex in the world, they can be illustrated with a much simpler simulation that uses the same four elements. Consider, for example, Galileo's gravity experiment, in which he purportedly dropped cannonballs of differing sizes from the Leaning Tower of Pisa to see which would hit the ground first. (Galileo thought the experiment through, developing a theory, and although most historians doubt that he ever actually performed the experiment, it has been done many times by many scientists, and has become a legendary example of the classical interplay between theory and experiment. We will bow to popular lore and attribute both the theory and experiment to Galileo.)

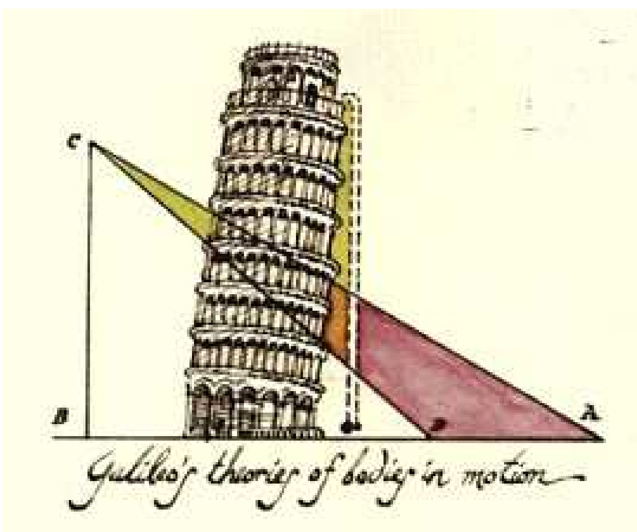

Figure 3-1. Galileo's gravity experiment.

The first element required for this simulation is an understanding of the underlying physical process. Galileo's theory postulated that the mass of a dropped object is not a factor determining the time it takes to reach the ground. The experiments appear to confirm this, to the accuracy that could have been measured at the time.

Second, the simulation requires an equation that describes the physical process. Through theory and experiments it was later determined that the force of gravity on Earth accelerates all objects at a rate of 9.8 meters per second per second, a constant generally represented by *g*. With this information, and letting *t* represent the elapsed time since the cannonballs were released, the height of a falling cannonball above the ground can be computed as $h=H-y$, where *h* is the height of the cannonball, *H* is the height at which the balls were released and $y=\frac{1}{2}gt^2$ gives the distance that cannonball has fallen.

Given the equation for *y* and the underlying theory, the simulation could be completed with a simple hand calculator. Letting Δt be a small time interval (known as the "time step"), the height of the falling cannonball is computed at time $t_k = k\Delta t$ using the formula $h_k = H - \frac{1}{2}gt_k^2$, for $k=1,2,3,\ldots N$ (the total amount of time simulated is $N\Delta t$). An experiment observing the actual position of the dropped cannonball could be used to validate the simulation. Because this simulation calculates only the single dimension *h*, the height of the cannonball, it is called a one-dimensional, or 1D, simulation.

If Galileo had thrown the cannonballs horizontally from the tower rather than just dropping them, a model of his experiment would require a second dimension and a slightly more complicated simulation. A new variable, *x*, representing the horizontal distance of the cannonball from the tower is added to the model (and a second equation governing that variable is added). Just as the acceleration of gravity must be known to determine *y* it is now essential to know the horizontal velocity, *v*, imparted by Galileo throwing the cannonball away from the tower. Once known, the horizontal distance of the cannonball from the tower is given by $x = vt$, where *t* is the same variable for time as is used in calculating *y*.

The simulation then proceeds very simply. For $k=1,2,3,\ldots N$, compute

$$t_k = k\Delta t$$

$$x_k = vt_k$$

$$h_k = H - \tfrac{1}{2}gt_k^2$$

where at each time step the location of the cannonball is now calculated in two dimensions.

Once again, a hand calculator could be used to compute the location of the cannonball at each time step, and observing the behavior of a real cannonball hurled from the tower could provide the data needed to validate the computational results. This simulation is two-dimensional, or 2D, since it results in a depiction of the cannonball's behavior in terms of both its height (*h*) and its distance (*x*) from the tower.

One and Two Dimension Simulations

It is important to note that, although very simple, both the 1D and 2D simulations of the cannonball might be quite useful, depending on what insight is sought. Neither simulation captures all of the physics involved in the actual experiment (such as the effect of the air resistance on the flight of the cannonball, or the interaction between the cannonball and the ground), but the scientist might not be interested in the small effects of those phenomena, or may wish to add them into a more detailed study later. These simple simulations can be completed very quickly and might be considered accurate enough for a specific purpose.

An important characteristic of a simulation is the resolution it yields resulting from the selected time sampling interval Δt. The effect of this choice can be seen by considering the 2D simulation of the cannonball being thrown from the tower. Assume the simulation runs until the cannonball reaches the ground. If the simulation uses only a few time steps, a plot drawn of *x* and *y* resembles a staircase with only a few steps. While this does not provide a complete picture of the cannonball's smooth flight from the tower, it does give a rough picture of the flight of the cannonball, which might be all that is needed. Using shorter time-steps (equating to smaller steps on the stairs) would more accurately portray the ball's flight; given sufficiently small time steps the plot of (x, y) locations resembles the smooth flight of the ball to arbitrary accuracy. In many simulations very short time steps are also necessary to capture events that happen very quickly. In some problems, unlike the simple cannonball example, the final solution cannot be directly calculated. For those cases the equations are solved over and over again, each time using inputs from the previous solution. This method continuously refines the answer until it settles on a steady (converged) result. Again, the length of the time steps can dramatically affect the quality of the solution.

Using smaller time-steps may provide more accurate simulations, but it also requires more computations, taking more time. This is one of the trade-offs that scientists using computational simulations are constantly required to make: the precision of a result versus the time it takes to produce that result. The challenge for ASCI was to improve precision while maintaining—or even improving—the quality of the answers, while dramatically reducing the time required to obtain those results.

As ASCI began, the National Laboratories used primarily 1D or 2D simulation applications. It was assumed that the phenomenon being simulated could be approximated on a straight line, a plane, a cylinder, or a sphere, with variables in one or two dimensions and the values in the third dimension (or in the second and third dimensions) assumed "by symmetry." A few 3D codes did exist, but the limited computing power then in existence meant that a large 3D simulation would take years to complete. Prior to ASCI, simulations in 1D and 2D were acceptable, largely because the results could be verified by the results of underground testing.

3D Simulations

The difficulty facing ASCI and the National Laboratories was that the potential effects of aging on nuclear weapons, such as corrosion, metals becoming brittle and cracking, and so forth, are fundamentally three-dimensional in nature. Accurate simulations would have to be 3D and yet would have to be very precise to capture the all aspects of a weapon's condition.

Why does a 3D simulation demand so much computing power? Galileo's simple falling cannonball can be used to describe how complex a simulation can become. Since air resistance, variations in the gravitational or magnetic fields, and other effects are miniscule, 1D or 2D simulations are generally adequate to follow the cannonball through the air. But what happens when it hits the ground? At that point, things would happen in ways that were not easily predictable. The impact would send shock waves through the cannonball, deforming it, and causing it to respond in complex ways, especially if the internal construction of the cannonball is not uniform. Shocks also propagate through the earth near where the cannonball lands, which could be important in determining the full effect of the impact. A full 3D simulation could be used to capture the shocks involved with this impact and trace them through the cannonball or earth. But this simulation is obviously *much* more complicated than the simple 1D or 2D versions, and requires the computation of many more unknowns than just the (x, y) location of the cannonball. It is no wonder that the computing power required is much greater for a full 3D simulation including effects.

Finite Element Meshes

There are several ways to build simulation codes, but a common one is known as *finite element* analysis, which can be used for 1-, 2-, or 3D models. In this type of simulation, a mesh of interconnected points represents the physical object, dividing it into many 1-, 2-, or 3D subdomains, each called a finite element. Thousands or even millions of points (and hence elements) may be used to represent the object. What is happening within and on the surface of the object is captured by computing the values of the unknowns (e.g., pressure, temperature, etc.) within and on the surface of each of the elements. The finite elements can be used to represent many different physical phenomena, including the physical displacement of the materials, temperature distribution, expansion or contraction due to heating or cooling, changes in density, evolving chemical composition, and many other physical phenomena. An equation-of-state, which is related to the state of a material (gas, liquid, plasma, or solid) given its temperature, pressure, and density, can be used to represent this information.

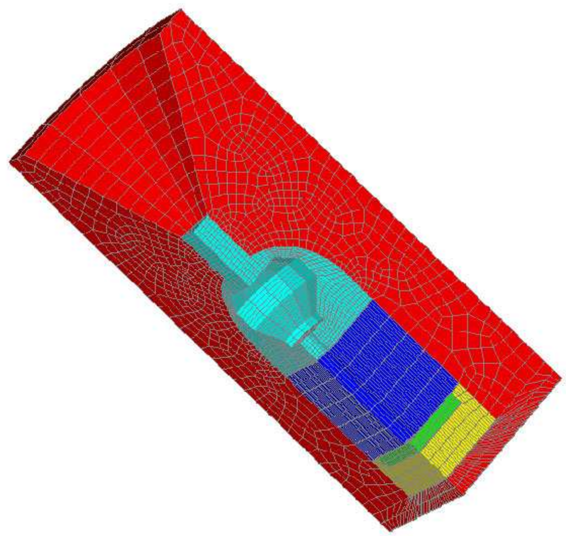

Figure 3-2. Example of a finite element mesh.

These types of simulations can be particularly useful in representing complex physical systems consisting of many different parts, because each of the components can be represented with its own set of finite elements and the computation can proceed by tracking all the components individually and paying special attention to what transpires at the interfaces where the components meet. With enough computing power, one could simulate extraordinarily complex items—a Ford Taurus, a Boeing 737, or perhaps a nuclear weapon.

Running the Simulation

Armed with a set of governing equations and a mesh, the simulation can be run. While it is possible to solve 3D calculations with a pocket calculator, as is easily done with the 1- and 2D cannonball simulations, computing a 3D simulation with that pocket calculator, including all the physical and chemical effects, might take hundreds or even thousands of years. Scientists and engineers need to get results somewhat sooner—at least before they forget the problem they were trying to solve. Using sufficiently high-performance computers allows results to be available in a reasonable amount of time, such as a few weeks. However, as simulations at the National Laboratories became increasingly detailed and complicated, the limits of the existing computers rapidly were reached. Larger, faster, and more powerful computers were needed. This was the situation that faced Vic Reis, Gil Weigand, and the Laboratories at the beginning of the Initiative. ASCI would need a *lot* more computing power.

Complications of Using Parallel Computers

The good news in 1994 was that parallel computing had reached a level of maturity whereby it was realistic to expect it to support the needs of the weapons program. All three National Laboratories had demonstrated the viability of parallel computing for physics-based simulations. LLNL had used the BBN TC2000 and the Meiko while LANL had employed a Thinking Machine CM-5 and SNL had run simulations on an Intel Paragon. While these early efforts showed that parallel computing could work for ASCI, they also demonstrated the complications that application programmers would face in dealing with models that comprising thousands or millions of finite elements distributed across hundreds or even thousands of computer nodes. In many ways, the developers programming the computers in the early days of ASCI had to invent entirely new approaches to simulation. It took the vision of a Gil Weigand or a Vic Reis to see that by the time ASCI came to an end the problems would have grown to include billions of unknowns and be run on systems with more than a hundred thousand processors.

Verification and Validation

Verification is the process of ensuring that a program will operate as its developers intended. That is, given a set of governing equations, verification ensures that the application code solves those equations correctly. It is a big part of the challenge of writing simulation programs, particularly massively parallel ones. With applications consisting of hundreds of thousands of lines of code and using hundreds to hundreds of thousands of processors interacting with each other, many opportunities exist for errors, or bugs, to be introduced. To deal with these and other potential problems, the National

Figure 3-3. A validation experiment.

Laboratories had to enhance their approach to software development. Among the approaches they took was to adopt and customize many of the disciplined practices recommended by the Software Engineering Institute (SEI) and used by commercial software companies.

In addition to verification, a simulation must produce results that actually reflect what occurs in the physical world. Methods of assuring that this is the case are known as *validation*. The most common approach to validation is to use the application to simulate an actual event that has been studied closely, often a previously conducted experiment. The simulation data is then compared to the real world data, which will confirm whether the simulation accurately reflects nature.

Figure 3-3 illustrates how validation can be performed. The photograph on the left illustrates a physical experiment in which a suspended device can be damaged by being dropped or swung forcefully against a hard target. The graphic on the right is the result of a simulation of the damage that would result. The details of the simulation can be compared with readouts from instruments that record the actual results of the experiment.

The following three vignettes provide examples of the challenges facing the National Laboratories as they developed applications that helped put simulation on a par with theory and experiment and establish a new paradigm for science.

Performance and Safety Applications – *The Direct Connection to the Weapons*

ASCI performance and safety applications provide a direct link between simulations and nuclear weapons, capturing incredibly complex interactions that occur over a wide range of time spans. Building these applications caused the Laboratories to change the way they operated and demanded new approaches to verifying and validating the results. ASCI's success in supporting Stockpile Stewardship is directly related to successes in creating these applications.

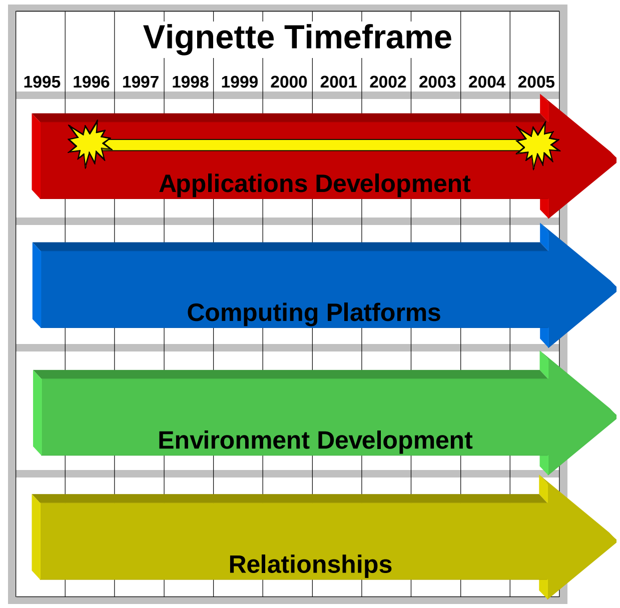

The ASCI performance and safety applications make a direct connection between science and nuclear weapons in the SBSS program. Stewardship involved the creation of huge experimental devices like the National Ignition Facility (NIF) and the Dual Axis Radiographic Hydrodynamic Test (DAHRT) facility, and while these experimental devices provide important scientific information, they were not designed to test or emulate full-scale, fully integrated weapons. The ASCI performance and safety applications were where the science of SBSS was applied to actual weapon designs. Due to the complexity of these applications and because they would have to operate on the world's most powerful, parallel computers, the nuclear National Laboratories adopted new and innovative approaches to writing the applications.

Shift from Test-based to Simulation-based Stewardship

The use of simulation in a predictive way was a significant shift for the U.S. nuclear weapons program. With President Bush imposing the unilateral test moratorium in 1992 and President Clinton continuing it in 1993, DP and the Laboratories needed to find a new way of doing business.[1] The urgency of this was underscored in 1995 when President Clinton announced his support for the Comprehensive Test Ban Treaty. From a written announcement prepared by Robert Bell, then Senior Director at the National Security Council for Defense Policy and Arms Control, the President said:

> One of my Administration's highest priorities is to negotiate a Comprehensive Test Ban Treaty (CTBT) to reduce the danger posed by nuclear weapons proliferation. To advance that goal and secure the strongest possible treaty, I am announcing today my decision to seek a 'zero' yield CTBT. A zero yield CTBT would ban any nuclear weapon explosion or any other nuclear explosion immediately upon entry into force. I hope it will lead to an early consensus among all states at the negotiating table.

The President also confidently addressed stockpile safety and performance:

> I am assured by the Secretary of Energy and the Directors of our nuclear weapons Laboratories that we can meet the challenge of maintaining our nuclear deterrent under a CTBT through a Science Based Stockpile Stewardship program without nuclear testing.[2]

From then on, DP and the Laboratories would have to assess the performance and safety of weapons in the U.S. nuclear arsenal without nuclear testing.

Simulation Challenges

Building simulation applications that could replace full-scale nuclear testing was a daunting challenge. In a 2004 article, Livermore's Randy Christensen recalled the state of simulation codes at the start of ASCI:

> Sophisticated weapon simulation codes existed before the ASC[I] and Stockpile Stewardship programs. However, because of the limited computer power available, those codes were never expected to simulate all the fine points of an exploding nuclear weapon. When the results of these simulations didn't match the results of the underground tests, numerical 'knobs' were tweaked to make the simulation results better match the experiments. When underground nuclear testing was halted in 1992, we could no longer rely so heavily on tweaking those knobs. [3]

Building simulations with predictive capabilities about nuclear weapon performance and safety is challenging, largely due to how the weapons operate. A nuclear weapon is a system with three elements: the engineering components, the "primary," and the "secondary." Engineering components provide the electro-mechanical actions necessary to ensure that the weapon cannot operate inadvertently and that it will operate when required. The primary generates the initial nuclear explosion. Energy from the primary is transferred to the secondary, where thermonuclear fuel is ignited, generating a much more powerful explosion. The engineering components are designed and maintained by SNL; for any given weapon, the primary and secondary are designed by either LANL or LLNL.

Simulating the performance of this three-part system, including the safety and reliability concerns, posed significant challenges for the ASCI. The events involved could span a very broad range of times and physical dimensions. Some of these events might take place over hours or days (e.g., the stockpile-to-target sequence, or STS), or even years (aging of the weapons and components). Others, such as stages in the detonation, occur on a scale of nanoseconds (billionths of a second). Furthermore, some effects (e.g., aging, extreme conditions, accidents) required representation in three-dimensions. Consequently, the codes Christensen described that were originally used to design the weapons were not adequate to simulate the performance and safety of the weapons under a broad range of conditions.

New, more accurate simulations of the weapons would also require representing system synergy, how the different physical processes were linked together. When underground test data could be used to tweak simulation results, it was sufficient each element in the weapon system to be simulated separately, passing the results to another simulation examine the action of the next element. This process resulted in inaccuracies, but the data were corrected by calibration with actual test date. The loss of testing demanded that simulations of the elements be coupled so that a single application could represent the behavior of entire systems. Such myriad requirements for improved applications prompted Weigand's demand for ASCI to enable "full-system, full-physics, 3D simulations."[4]

In some ways the Laboratories were almost starting from scratch. Until ASCI, only limited work had been done to develop these types of applications or to put them onto parallel computers. At the time, most codes could only simulate one or two dimensions, and in fact most calculations were done in 1D because of the time and complexity required to complete a 2D calculation.

Furthermore, the fundamental architecture for the computers running these applications was changing from serial to parallel processing. This was a huge change in the landscape of computing and required completely new algorithms be devised for every aspect of the computation. The breadth of this change and the speed with which it transpired was a shock to the scientists building the applications.

Code Teams

To meet these challenges, the Laboratories realized that traditional methods of writing applications would have to change. Up until the early 1990s, teams of just a few people, often with only one primary author, crafted the application codes used to simulate the performance and safety aspects of nuclear weapons, resulting in what were known as "hero-codes."

New organizations and new approaches were needed at all three Laboratories. The new simulation codes had to provide increased dimensionality, improved physics representation, linked system effects, and had to run on computers with hundreds, or even thousands, of individual processors. It also had to happen quickly: ASCI leadership estimated they had a 10-year window in which scientists with nuclear test experience would be available to validate the simulations.

In response, the Laboratories implemented a *code team* approach to developing ASCI applications. Teams were created, often of 20 to 30 people working together for spans of several years. While this might seem quite small compared to commercial software development efforts, it was a significant shift for the Laboratories. A cadre at the center of each team provided overall direction and concentrated on how the physics would be represented in the simulations. Other team members focused on particular aspects of applications, such as specification of the computational mesh, an application module that adapts the finite element mesh to improve the accuracy and computational efficiency of the code. Other team members implemented software engineering and quality practices, including capturing and tracking the fulfillment of requirements. All three Laboratories implemented a rigorous program of *regression testing* in all parts of the codes nightly to ensure that if newly modified portions introduced errors, they were immediately discovered and corrected.

People – The Most Important Resource

As ASCI Executive at LANL, James Peery was responsible for management of the Initiative there, and for oversight of code development. Peery was asked in 2005, "What went right with ASCI?" He singled out the formation of large code teams as a critical improvement in the way the Laboratories built the applications. He also cited the institution of software quality practices as an important, significant change in their approach.[5] Peery was in a good position to know; before joining LANL in 2002, he had spent the previous 12 years at SNL, where, among other things, he had led ASCI code development efforts.

ASCI's development of large, diverse teams focused on simulating the nuclear weapons might be its most important legacy. Mike McCoy, ASCI Executive for LLNL, once said, "The computers come and, after a few years, they go, but the codes and code teams endure."[6]

Enduring Physics-Based Codes

Applications simulate systems operating according to the laws of physics. Because these laws are fixed (more accurately, because our understanding of them evolves very slowly), these applications often endure for many years. These codes also take a long time to develop; years will pass as users validate the codes and grow to accept that the code accurately reflects the physical world. When the code developers do recognize that changes are needed, they tend to make incremental improvements, rather than beginning anew. In this way, an application known by a certain name might easily be used over the course of decades even though the underlying functionality changes significantly.

LLNL's DYNA3D application, first released in the early 1980s, is an excellent example. The application was developed to simulate structural dynamics for physical systems. Over the last 20 years, it has been constantly adapted and improved. Today, LS-DYNA, a version of the application commercially available from Livermore Software Technology Corp., is used by industry to simulate everything from the strength of beer cans to the survivability of automobiles in accidents.[7] As McCoy suggested, codes evolve, persist, and ultimately succeed because of the talent and dedication of code teams. In response to ASCI's mission, the Laboratories transformed how codes were developed by better leveraging their greatest resource—people.

Milestone Driven

In addition to the transformation from hero-codes to code teams, ASCI also implemented the use of milestones in application development. ASCI needed a way to measure its progress. But the evolving nature of the codes, as demonstrated by DYNA3D, made it difficult to pinpoint the "completion" of a new application; therefore, that was an ineffective rallying point. Instead, ASCI milestones were focused on the addition of new capabilities to simulation applications, something much easier to measure.

Milestones became important to setting objectives for the Initiative and provided a sound means to measure progress. An advantage of using the milestone for code development was that they could be tied to progress in other areas of ASCI, such as increases in platform performance. The application milestones the Initiative set were technically quite aggressive, but that was necessary to achieve the goals of the Initiative in the short time allotted. The 2005 ASC Program Plan (ASCI changed its name to the Advanced Simulation and Computing program in 2002, though the "*I*" was not dropped from the acronym until 2005) provided an excellent summary of the application milestone that were achieved, including:

- In FY 2000, ASCI successfully demonstrated the first-ever 3D simulation of a nuclear weapon primary explosion and the visualization capability to analyze the results.
- In FY 2001, ASCI successfully demonstrated simulation of a 3D nuclear weapon secondary explosion, and ASCI completed the 3D analysis for a stockpile-to-target sequence for normal environments.
- In FY 2002, ASCI demonstrated 3D system simulation of a full-system (primary and secondary) thermonuclear weapon explosion, and ASCI completed the 3D analysis for an STS abnormal-environment crash-and-burn accident involving a nuclear weapon.
- In FY 2003, ASCI delivered a nuclear safety simulation of a complex, abnormal, explosive initiation scenario and ASCI demonstrated the capability of computing electrical responses of a weapons system in a hostile (nuclear) environment.
- In FY 2004, ASC provided simulation codes with focused model validation to support the annual certification of the stockpile and to assess manufacturing options.[8]

In June 2004, Tom Adams was quoted in LLNL's *Science and Technology Review*; he stressed the importance of the milestone approach to the ASCI applications development thus:

> We accomplished major objectives on time—with the early milestones demonstrating first-of-a-kind proof-of-principle capabilities. Achieving these milestones was the result of an intense effort by the code teams, who were assisted by dedicated teams from across the Laboratory. ASC milestones have now transitioned from these early demonstrations to milestones focused on improving the physics fidelity of the simulations

and supporting stockpile stewardship activities. We are now in the position of delivering directly to [SBSS].[9]

Ensuring Accuracy

Armed with innovative code teams to develop the applications and milestones to measure the progress, ASCI still faced the challenges of verifying and validating the results. The importance of these activities was recognized with additional funding as part of the Initiative's 1997 re-baselining. In a January 2005 article in *Physics Today*, Douglass Post of LANL and Lawrence Votta of Sun Microsystems provide a good list of the techniques available to verify and validate ASCI applications.

For verification these included:

- Comparing code results to a related problem having an exact answer.
- Establishing that the convergence rate of the truncation error with changing grid spacing is consistent with expectations.
- Comparing calculated results with expected results for a problem specially manufactured to test the code.
- Monitoring conserved quantities and parameters, preservation of symmetry properties, and other easily predictable outcomes.
- Benchmarking—that is, comparing results with those from existing codes that can calculate similar problems.

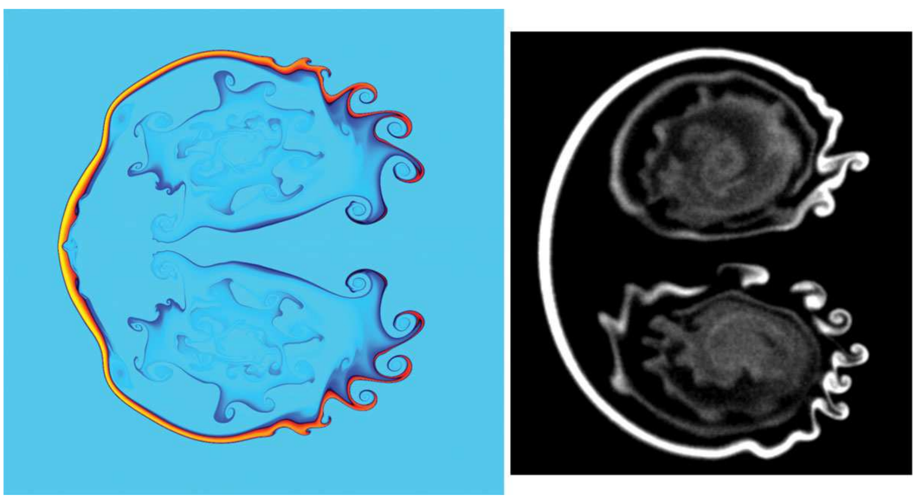

Figure 3-4. Comparison of a simulation (left) with an experiment, useful for validation.

For validation:

- Passive observations of physical events—for example, weather or supernovae.

- Controlled experiments designed to investigate specific physics or engineering principles—for example, nuclear reactions or spectroscopy.
- Experiments designed to certify the performance of a physical component or system—for example, full-scale wind tunnels [tests].
- Experiments specifically designed to validate code calculations—for example, [tests conducted in] laser-fusion facilities.[10]

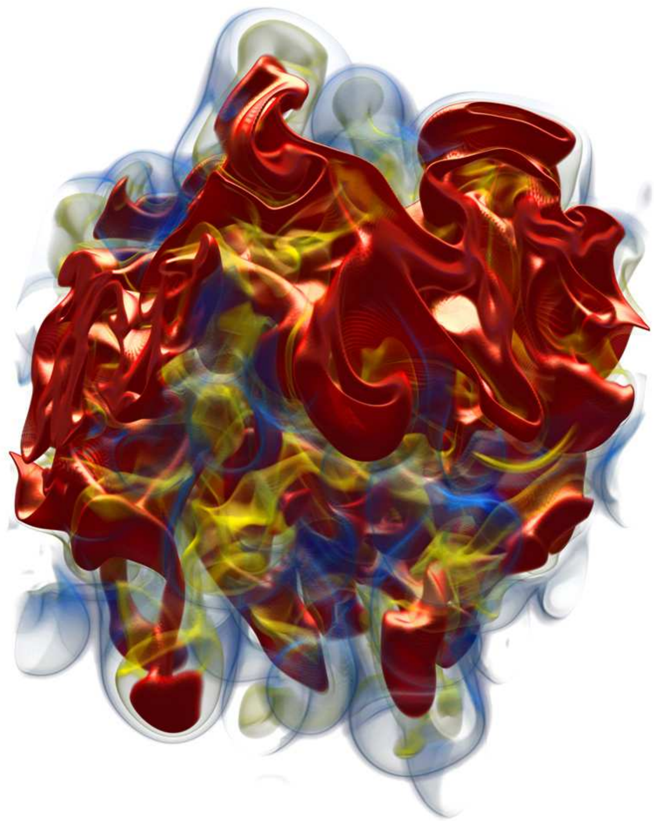

Figure 3-5. Example of 3D simulation of a Rayleigh-Taylor instability.

Post and Votta also discuss the challenges of building application codes that could be used to accurately predict the behavior of physical systems. Their article explores various ASCI code development efforts and assesses the progress made by each. As they point out, some code team efforts were not successful.

Getting the Code Teams Right

Post conducted, with Richard Kendall, a separate, detailed analysis published by LANL, which focused on application development efforts at both LANL and LLNL. Post and Kendall compiled a list of attributes common to the successful code projects. While there is not universal agreement with all their conclusions, the following list provides excellent insight into how ASCI got it right with application development:

- Assembling good teams of highly competent staff is essential. All other considerations are secondary.
- Schedules and resource levels must be determined by project requirements. Setting them independently will wreck the code projects.
- The schedule and resource estimates should be based on the institution's code development experience and history.
- The code development effort should operate as an organized project.
- It is important to identify the risks and provide mitigation. Setting clear requirements, obtaining management and stakeholder support, and getting contingency schedule and resource allocation are all essential.
- Maintenance of customer focus is essential for success.
- The most important product is a code with better physics representation.
- Risks should be minimized by using modern but proven computer science; computer science research is not the goal.
- It is critical to invest in the staff, providing training and support.
- "Best Practices" should be emphasized, rather than "Processes."
- It is essential to develop and execute a verification and validation program.[11]

Opening Up Thought Processes

ASCI and the Laboratories successfully overcame tremendous technical and organizational challenges to develop the simulation applications needed to meet the Initiative's primary objectives. Bruce Goodwin, Associate Director for Defense and Nuclear Technologies at LLNL, was asked in 2005 about ASCI's impact. He replied, "The impact has been great. We are now at a point where we are getting a new scientific understanding about how the weapons work." Reflecting on the applications, platforms, and environments that support the weapon designers and scientists that work for him, he said, "It has really changed the way the designers approach their work. I think the biggest change that occurred due to ASCI is the fact that it has opened up the thought process of the designers."[12]

Materials Modeling – *At the Level of First Principles*

Understanding how materials behave under various and sometimes extreme conditions is a vital part of ensuring the ASCI simulations correctly represent the physical world. ASCI's material modeling, executed on the world's most powerful computers, enabled insight into material properties that were impossible to completely understand through theory or with experiments.

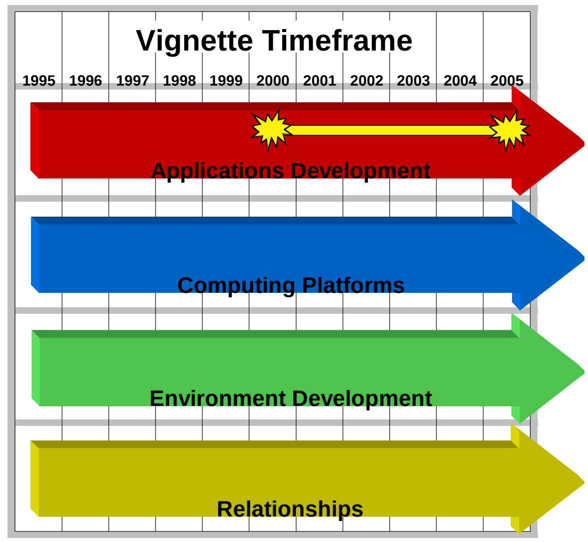

As the United States entered an era without underground nuclear testing, the scientists and engineers responsible for the nation's nuclear stockpile sought new techniques to build a thorough understanding of the weapons. Constructing a good simulation demands a detailed understanding of how materials behave. The 2002 ASCI Program Plan concisely states the challenge:

> In the past, physical properties of materials significant to the nuclear weapons program were often inferred from integral test data. Now, without the ability to conduct such integral tests, there is a premium on the development of advanced capabilities—experimental, theoretical and computational—necessary to predict the physical properties of materials under conditions found in nuclear explosions and stockpile-to-target sequences.[13]

ASCI's use of computational simulations for material modeling provides a clear example of how simulation interacts with theory and experiment. Not inconsequentially, ASCI's material modeling work also demonstrates the Initiative's impact on scientific progress in disparate fields well outside the fences of the nuclear weapons Laboratories.

Challenges of the Traditional Scientific Approach

Obtaining detailed scientific insight about some physical processes, such as how molten metals solidify, has always been a challenge, due partly to the limitations of theories and experiments. Theorists, of course, are limited by their basic understanding of the physical world and their imaginations. Material under extreme conditions—such as within high explosives, during the formation of stars, or in the operation of nuclear weapons—sometimes acts in highly unusual and often unpredicted ways. Consequently, development of accurate and useful theories which address material behavior under extreme conditions has been extremely difficult.

All three Laboratories conducted experiments to understand the dynamic behavior of materials (usually metals) that begin as solids, but transition to liquids, gases, or even plasmas as they are subjected to high speed, high power shocks. Among the sources used to generate shocks are explosives, high-velocity gas guns, and lasers. Various diagnostic instruments, including sophisticated cameras, high-powered x-rays (the resulting pictures are called radiographs), and laser-based tools, are used to capture how the material behaves as shockwaves travel though it.

Even though the diagnostic instruments used by the Laboratories were (and are) among the most sophisticated available, the usefulness of data they obtain is limited. Radiographs, for instance, are fundamentally 2D images. Computers can be used to transform them into 3D images, but this is possible only when multiple images are collected, in some experiments an extremely difficult task. The incredibly short duration of some of these physical events can also be an experimental limitation. During a hydro test, for example, an explosive charge is used to drive metal plates into small samples of a material, creating a shockwave which results in an event lasting only a tiny fraction of a second.

Nuclear weapons scientists often faced another challenge when conducting physical experiments. The high explosives and hazardous materials needed to generate an experiment's physical forces sometimes destroyed the very instruments that collect the data. These issues limit the scientist's ability to gather sufficient data in the experiments.

Use of Empirical Data

Recognizing these limitations, scientists have long combined experimental results with simulations, using the results to create new and better simulations of different physical systems. This approach can be described using Galileo's cannonball as it hits the ground. A sample of the iron from which the cannonball is made could be used to conduct experiments to produce data about its deformation properties under various regimes. That empirical data would then become inputs to the simulation and would be used to describe the model as different conditions are simulated. Because experiments can not capture every possible condition, the simulation would interpolate, or connect the dots, between data points.

Using empirical data in simulations is very efficient since it does not require the calculation of underlying physical model properties. However, it only works when there are sufficient experimental data available. Acquiring sufficient data to adequately describe the physics is feasible for an experiment involving a cannonball, a tower, and a relatively large amount of time. It is, however, far less feasible for nuclear weapons experiments, where the extraordinarily hostile environment and astonishingly short time scales make the acquisition of detailed experimental data extremely difficult.

Overcoming Physical Limitations

ASCI's materials simulations are used to calculate the performance of materials at molecular, atomic, and sub-atomic levels. This first-principles approach to the models allows the simulation to make predictions about how the materials would behave under a variety of physical conditions.

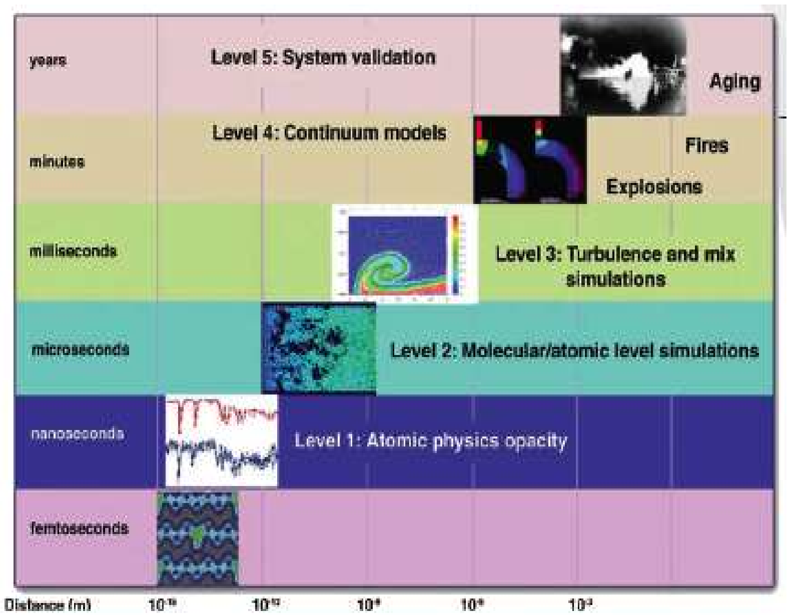

Figure 3-6. Physics occurs, and must be simulated, at many scales.

As noted earlier, one of the difficulties of simulation is that various physical effects occur on diverse scales of time and size. Corrosion or metal cracking of metals, for example, takes years to occur. Other effects, such as the mixing of materials during detonation, happen at the atomic level in picoseconds (trillionths of a second). Figure 3-6, from the 2001 ASCI Program Plan, provides a good overview of the types of physics that had to be simulated and the scales and timeframes that were required.[14]

Before ASCI the standard practice was to model multi-scale systems by creating seperate simulations for each scale (of time and distance) that was important in the physical process. This approach introduced inaccuracies. The Laboratories' approach, in response to ASCI's challenge, was to capture simultaneously as many of the different scales as possible. The catch was, where would they find enough computing power to handle the massive calculations needed to do this?

At a 2005 colloquium at LANL, Michael Ortiz of the California Institute of Technology outlined the issues involved in creating the needed multi-scale simulations. He made several important points:

- An effective multi-scale modeling paradigm will reduce or eliminate uncertainty and empiricism in the simulation of complex engineering systems.
- The ultimate goal of this paradigm is parameter-free (first-principles) predictive simulation.
- Material behavior occurs on multiple length scales and the underlying physics changes from scale to scale.

- The governing equations at each scale are determined by the physics; bridging the length scales is largely a mathematical and computational problem.[15]

Ortiz's last point is important. It observes that mathematical and computational limitations constrain our ability to appropriately model materials in a simulation. Before ASCI, the most powerful computing available permitted simulations of up to a few thousand atoms at most, and those for only very short periods of time. More powerful computers, made available in the Initiative's early years, enabled a vast leap to materials simulations capable of modeling a few million atoms, and shedding light on the atomistic behavior of materials. But this was still not enough power to enable capture of macroscopic behaviors of materials. Eventually, the staggering power of later ASCI computer platforms made it possible to simulate the interactions of millions and billions of atoms needed to create accurate predictions of material behavior.

Waiting for Power

Not surprisingly, materials modeling presented an additional complication for ASCI. The type of computing power it demanded was different from that needed for the nuclear weapons performance and safety calculations—largely due to inter-processor communication requirements. Because materials simulations only needed to capture data about how adjacent atoms interact, most data communications for parallel materials codes are limited to a few "nearby" processors. By contrast, performance and safety codes, which sometimes have to capture the movement of radiation across a system, demand much more sophisticated communication capabilities.

Although the first ASCI computing platforms were not powerful enough to realize the parameter-free, predictive simulation goal discussed by Ortiz, the Initiative always recognized the importance of materials modeling. A demonstration of this foresight was the temporary assignment of LLNL's Christian Mailhiot to ASCI at DP headquarters. Mailhiot's specialty was materials modeling research; he brought Weigand and other headquarters staff a valuable perspective on the potential power of the first-principle materials modeling needed to meet ASCI's mission.

In 2005, Mailhiot was asked about ASCI's support for materials modeling over the course of the Initiative. He said, "The support has been good and steady, especially from [LLNL's David] Nowak."[16] A physicist, Nowak had primarily worked on weapons design before moving to ASCI. Hence, he could appreciate the limitations of theory and experiment; he also had the forsight to advocate strongly for materials modeling despite the fact that many of the predictive capabilities would not be available for years.

A New Way to See the World

DP and its Laboratories began to see returns on their investments by the early 2000s, with the introduction of computing platforms that provided trillions of floating point operations per second (teraFLOP/s). Significant platform developments came in 2005, especially with the BlueGene/L installation at LLNL. This platform's architecture, consisting of thousands of lower-power processors joined together with several different types of networks, was particularly well suited for simulating how materials perform.

Even before BlueGene/L was fully delivered in fall 2005, its aptitude for materials modeling simulations was apparent. In the spring and summer of that year, Timothy Germann, Kai Kadau, and Peter Lomdahl of LANL used BlueGene/L to simulate the shock compression and release of porous copper. They used SPaSM, a molecular dynamics code, to simulate how a micron-sized cube of copper, consisting of 160 billion atoms, would react to being shocked.[17]

This research, and other work being done on the early versions of BlueGene/L, was revolutionary because it captured material properties on a much larger scale than was previously possible. One micron might be small—about one tenth the width of a human hair—but it was huge compared to any prior material simulations. This difference mattered; previous simulations could represent only a small part of a physical effect, such as the formation of a crack or a void. Germann, Kadau, and Lomdahl's research resulted in a paper that was submitted to SC|05 (Supercomputing 2005, an annual technical conference that covers supercomputing and communications issues). Their paper observes:

> The multi-billion atom simulations enabled by BlueGene/L now opens up to study a wide range of problems where an initial or (more often) a naturally emerging defect/microstructure length scale exceeds that of smaller-scale simulations. . . With such simulations [it is] now possible for more realistic interaction potentials, and turnaround times of a day or less, [therefore] we expect to see an exciting variety of previously intractable problems studied by MD [molecular dynamics] on BlueGene/L.[18]

One of the beauties of computational simulation of material behaviors is that simulations are not constrained by the limitations of real world data capture. With simulation, scientists do not need special equipment to capture data about how a material behaves in a very hostile environment, or over a very short time, such as a shockwave in a cube of copper. This information

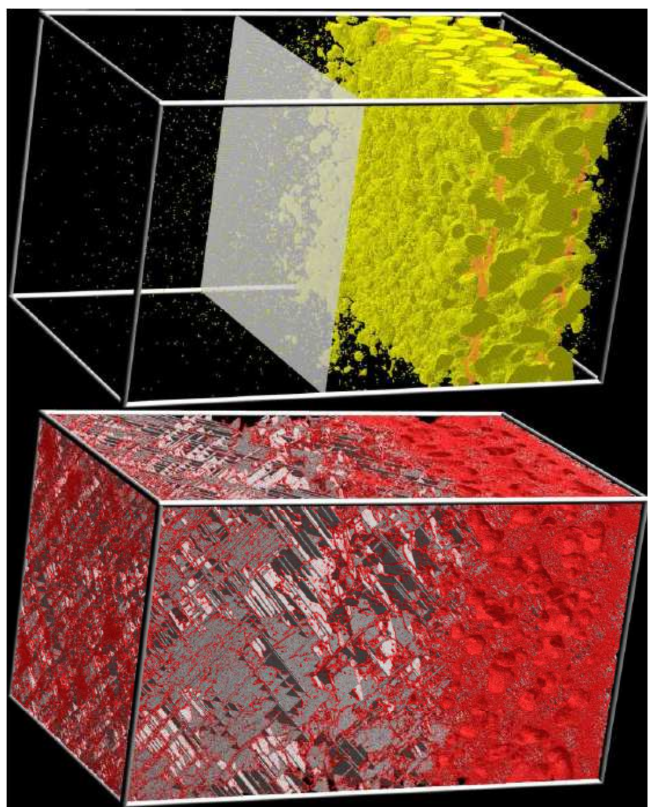

Figure 3-7. A 160-billion atom simulation of copper undergoing a shock.

is generated as part of the simulation; the scientist can simply access the required data. Of course, this very advantage points to the challenge of ensuring simulations are an accurate reflection of the physical world. This was another reason that ASCI insisted on verification and validation of material property simulations.

Unlocking the Door

ASCI's work on materials modeling has unlocked a tiny, armored door, ushering science into a new era. It is difficult, without resorting to science fiction, even to discuss the possibilities presented by the power of simulation when done with tera- and even peta-scale computing. Elaine Chandler, who managed the ASCI Materials and Physics Models program at LLNL, eloquently captured this future significance when quoted in 2004:

> Nearly a half century ago scientists dreamed of a time when they could obtain a material's properties from simply knowing the atomic numbers of the elements and quantum-mechanical principles. That dream eluded us because we lacked computers powerful enough to solve the complex calculations required. We are just now able to touch the edge of that dream, to reach the capabilities needed to make accurate predictions about material properties.[19]

Among the most exciting developments of these new molecular modeling capabilities is their viability for modeling biological systems. These systems operate at the cellular and molecular levels, and the new levels of computing power available have the potential to accurately simulate these interactions. It is a staggering possibility; biological systems are so complex that until recently very few believed simulations could ever be realized. And yet, in February 2006, the National Science Foundation published a report entitled "Simulation-Based Engineering Science: Revolutionizing Engineering Science through Simulations." On the use of simulation-based engineering in medicine, the report states:

> Most diseases (such as heart disease, cancer, stroke, and respiratory diseases) and their treatments (whether surgical, trans-catheter, or pharmacologic) involve complex physical responses and interactions between biological systems, from the molecular to organism scales. Simulation methods can therefore dramatically increase our understanding of these diseases and treatments, and furthermore, improve treatment.[20]

ASCI work on materials modeling done by the National Laboratories provides a great example of how simulation has become an integral part of how scientists and engineers now explore the world. Furthermore, the results of ASCI's work have dramatic promise to answer questions of broad interest to humanity. In many cases, theories and experiments have been inadequate to discover how physical systems operate. Materials modeling simulations provide scientists with tools to learn how nature works in scales of time, space, and energy that are cannot be observed directly in the physical world. It is fair to say that computational simulation can make the undetectable…visible.

Programming Models – *Creating the Applications*

The advent of parallel processing computers in the late 1980s and early 1990s offered new opportunities to create powerful simulation applications. It also created significant challenges in writing those programs. ASCI successfully used a number of different approaches, or programming models, to make the transition from serial to parallel high-performance computing platforms.

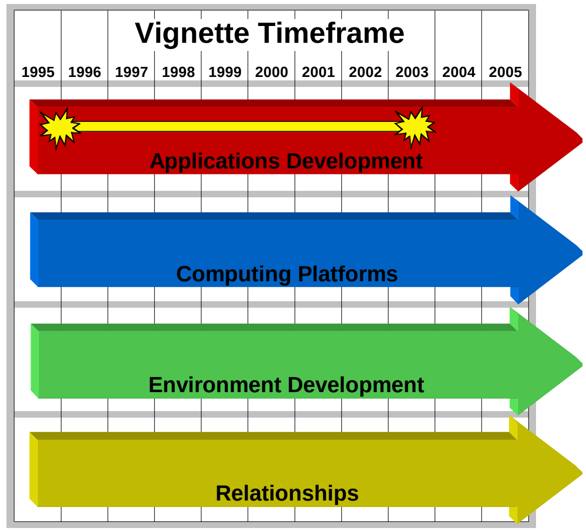

Savvy application developers, like those at the Laboratories, will always design their programs for the structure, or architecture, of the computer on which they will run. The demands, advantages, and drawbacks of the architecture must always be kept in mind. This mental picture of the system is known as the programming model, and it has a tremendous impact on how a programmer approaches building an application. ASCI's simulation applications were created during, and in many ways because of, a period of tremendous turmoil in the world of computer architectures.

At the start of ASCI, even if the market for high-performance computers had been stable, the weapons programs' needs for applications to better represent physics and dimensionality would have presented a significant challenge. Unfortunately, the prominent, serial-style architecture was not going to provide the computing power needed to run those applications. The fact that serial architecture failed to scale adequately threw the HPC industry into chaos.

Providentially, that very turmoil in the computer industry made the newer, more powerful simulations possible. The emergence of parallel processing computers using cheaper "commodity processors," provided ASCI a pathway to satisfy the computing power demands of the applications. The challenge for the application developers was in determining an appropriate programming model to use. In the early 1990s, a number of possible parallel computer programming models were emerging, and it was not at all clear which would be the most effective. Among the most important lessons ASCI learned was that changes in computer architectures happen continually and that often the best strategy is to choose programming models most adaptable to those changes.

Killer Micros

In 1989, LLNL launched a Laboratory Directed Research and Development (LDRD) project called the Massively Parallel Computing Initiative (MPCI) that studied the trend of using commodity microprocessors in parallel to support scientific research.

The subsequent reports of this project were published in 1991 and 1992 and became known as the "Attack of the Killer Micros." Livermore's Eugene Brooks led the research effort and edited the reports. He concluded the introduction to the 1991 version with this prediction:

> In sharp contrast to the consensus opinion of just a few years ago, very few computational scientists doubt that highly parallel machines will become the supercomputers of the future. It is clear that massively parallel machines will offer the Laboratory an entirely new computational capability in the 1992-1993 timeframe. Massively parallel machines, which provide an aggregate scalar performance of more than 80 times the capability of a conventional supercomputer, will be available during this time. These same machines will provide an aggregate vector performance, which is around 4 to 8 times faster than what will be obtainable from a conventional supercomputer. Conventional supercomputers will be *outclassed* [his emphasis] for scalar workloads and will lack luster for vector workloads almost more quickly than we can possibly adapt to the new technology. The Massively Parallel Computing Initiative is positioning the Laboratory to be able to take advantage of the new breed of supercomputer as they arrive. We will either be using Massively Parallel Computers for production computing by the end of 1993, or we will not be supercomputing.[21]

Starting with the world's first electronic computers, including the 1946 ENIAC and until late 1980s, almost all computer architectures consisted of a single Central Processing Unit (CPU) connected to a single bank of memory that executed instructions in a serial fashion. ENIAC was built at the Moore School of Electrical Engineering at the University of Pennsylvania and executed 333 multiplications per second (333 FLOP/s). Until that point, calculations such as these had to be done by hand using mechanical calculators that would take several minutes per calculation. The ability to perform hundreds of FLOP/s was truly revolutionary.

Over the next 40 years, advances continued to be made in computer performance. ENIAC's 18,000 vacuum tubes were replaced with solid-state transistors and eventually with microprocessors etched on silicon chips. The cables and switches used to program ENIAC were eventually replaced with paper tape, then punch cards, and later by electronic memory.

Until the advent of MPP computing, the greatest advances in performance were seen in the early 1980s with the introduction of the computers built by Seymour Cray. His company, Cray Research, developed specialized hardware and software that fully exploited an emerging technology known as *vector computing*. First introduced in the early 1960s by Westinghouse, the goal of vector computing was to dramatically increase the performance of mathematical calculations by using a large number of specialized math "co-processors," all under the control of a single CPU. The co-processors would then simultaneously perform the same floating-point operation, each on a different data value in a "vector" storage array. Hence in vector computing, the same algorithm can be simultaneously performed on a large data set. Cray's machines implemented a variety of clever hardware and software improvements, and eventually achieved normal operation speeds of 80 megaFLOP/s (million FLOP/s) with peak performance up to 240 megaFLOP/s, a many-fold improvement over all previous architectures.

The development of vector processing technology introduced a new programming model for application developers, and it took several years for them to adapt. Later, Cray Research offered systems that used multiple processors with a vector connection to the memory. From a programmer's point of view, there appeared to be a single processor connected to a single bank of memory—the classic serial computer. The problem was, as Brooks cited in his report, expected performance improvement from vector computing was slowing.

Parallel Computing Programming Challenges

Vector architecture performance was peaking just as SBSS demanded a vast and rapid increase in computational power. It was natural to ask if applications would be able to use MPP technology to meet that demand. No one was certain. Brooks and others, excited by The Killer Micro revolution in the early 1990s, recognized that an entirely new approach to creating simulation applications would be required.

Rather than dealing with one processor (or a few processors) interacting with a single bank of memory, in MPP computing applications would be spread across tens to possibly thousands of processors, each of which would have its own relatively small bank of memory. Data held by one processor would, if needed by another processor, need to be passed through a processor interconnection network. Because of this added complication, programmers had to consider several new challenges that impacted application performance. In a 2006 *HPCWire* article, Tom Sterling of Louisiana State University provided an excellent list of those performance problems:

- **Latency** - the number of cycles to a remote resource like local main memory (hundreds of cycles) or distributed nodes (thousands of cycles).
- **Overhead** - the amount of work (again measured in cycles) to manage the parallelism and (through an Amdahl argument) determine the granularity and therefore the amount of parallelism that can be effectively exploited.
- **Contention** - the delay time due to insufficient throughput such as chip pin or memory bandwidth.
- **Starvation -** the lack of useful work to be performed by the multiple [processor] cores because there is not enough user parallelism due to poor load balancing.[22]

With the introduction of parallel architectures, how a program was distributed on a computer became a critical factor in how well and how quickly the program would execute. Combined with the extraordinary challenge of ensuring that an application properly represents the physical world at the resolution needed for predictive simulation, many thought it would be impossible to write effective distributed programs, particularly ones that simulated nuclear weapons.

Understanding the Technology

The Laboratories were hardly idle in the world of parallel computing before ASCI. All three Laboratories experimented with parallel computers to determine if they could be adapted to nuclear weapons simulation. LANL had purchased the CM-2 and CM-5 computers from Thinking Machines. On August 23, 1990, Eldon Linnebur, Frank Bobrowicz, and Harold Trease of the LANL Applied Theoretical Physics (X) Division published a paper entitled, "Use of the Connection Machine in the Nuclear Weapons Program at Los Alamos," in which they concluded:

> The CM2 is the first massively parallel computer to be incorporated into the Nuclear Weapons code development effort. Its effectiveness and value in addressing an important nuclear weapons issue has been amply demonstrated. In addition, the CM2 offers much more potential for future improvements in terms of scalability of the architecture. In practical terms this means more processors of even greater power, more flexible connectivity and larger memories. We believe that the CM2 follow-on, the CMX [later known as CM5], offers the prospect of making three-dimensional calculations a more routine and more effective part of the X-Division's production environment.[23]

SNL similarly experimented with the Intel Paragon and with systems built by NCube. Ed Barsis, Director of SNL's Computer Research in the early 1990s, said in a 2005 interview, "At the time, SNL was experimenting with MP [Massively Parallel] computers and had shown that they could be used to solve real problems."[24] LLNL's Massively Parallel Computing Initiative also reached similar conclusions.

Programming Model Options

Having realized that parallel processing computing could be applicable to nuclear weapon simulation, Laboratory code developers faced a dizzying array of choices in how to write the applications. These choices would be driven by the various programming models that were available to developers writing programs for tens to hundreds to thousands of processors.

Though a simplification, parallel programming models could be categorized into four possible strategies describing how the parallel operations are handled. When appropriate, developers might create a hybrid model, using some or all of these strategies:

- The programmer is assisted by *libraries* of functions.
- The program is created within a *framework*.
- A *parallel programming language* is used.
- The *hardware and operating system* handle the parallel issues.

In the context of computer programming, a *library* is a collection of functions that are used repetitively. They allow a programmer to use simple instructions to direct a program to access the library, which in turn will execute the common tasks. In this parallel programming model, the programmers use libraries to simplify implementation of parallel execution within the code. For example, the mathematical functions, such as the square root, sine, cosine, and exponential, are generally gathered together in a library so the programmer does not have to write code into each application for these common utilities. Libraries can consist of simple or exceedingly complex and intricate tasks.

The best-known library used for parallel programming is MPI (Message Passing Interface). As the name implies, MPI is used to handle all the communications among the processors, and has a number of built-in functions ensuring that data get from place to place in the computer accurately and exactly when needed. MPI was born out of a number of efforts in the early 1990s, such as PVM (Parallel Virtual Machine) from the DOE's Oak Ridge National Laboratory. MPICH, which stands for MPI Chameleon, is a portable version of MPI developed at DOE's Argonne National Laboratory. MPICH became widely available and established MPI as a standard interface. Today there are a number of different implementations of the MPI interface, including versions supplied by the hardware vendors as well as open-source versions such as Open MPI. All three of the nuclear weapons Laboratories used MPI libraries effectively to create so-called *portable applications* that could be moved to different computer architectures with relative ease.

Programming parallel computers can also be accomplished using *frameworks*. Like a library, a framework contains code enabling the developer to perform important tasks without having to write the code for the same functionality repeatedly. However, where the libraries are subservient to, and invoked from, the application, the framework reverses this hierarchy. That is, the framework is a software "superstructure" within which the application is created, handling the high-level control. Frameworks allow a programmer to use relatively simple commands to invoke all the components of the application, while the framework itself implements the complicated actions needed to control the execution, communication, and data flow in the parallel application. One example of this that showed some success was POOMA (Parallel Object Oriented Methods and Applications), a framework developed by a team led by John Reynders at LANL. Another highly successful framework is known as Sierra and is used at SNL to link several applications together.

A "language" in the computing world consists of the human-understandable instructions that are converted to machine-understandable instructions and executed by the computer's processors. One way to address parallel computing issues is for the programmer to use a *parallel programming language*, that is, a language that has specialized logic built into it that directly handles the issues of interprocessor communication and data locality. An application written in such a language would by its nature execute on a parallel architecture. The most elementary instructions are in "binary code" consisting of 1s and 0s, describing the states (on and off) of a computer's electrical circuits can assume. Humans can write in binary code, but using higher-level languages allows humans to write in very simple terms what will become extremely complex binary instructions. Before a program written in a higher-level language can be executed on a computer, it must be compiled, a process which turns the lines of code, into a machine-understandable form. In the process of making that conversion, the compiler could be used to identify performance characteristics of a program, such as the need for particular sets of data, and then find efficient ways to execute it, such as "pre-fetching" data.

The most common high-level language used at the Laboratories when ASCI started was named FORTRAN (from formula translation). At that time, it was the scientific community's primary language for writing simulation applications. There was also growing interest then in the C language (and later in the C++ language). Unfortunately, none of these languages were designed to support parallel computing automatically, and programmers using these languages had to rely on other techniques, most commonly the use of libraries. In time, several parallel programming languages were created, including HPF (High-Performance FORTRAN) and UPC (Unified Parallel C), which used the compiler to deal with many of the parallel computing issues, but these were not adopted for the ASCI applications.

A fourth approach for supporting parallel operations connects all the processors to a single memory bank and leaves it to the hardware and the operating system to handle data movement between processors. This programming model recalled a similar feature of earlier computers in which there was a single processor connected to a single bank of memory (and from the point of view of any given processor this is exactly how the architecture appears). Computers of this type are known as Shared Memory Processor (SMP) systems. (This architecture is also called *symmetric multiprocessor*, using the same SMP acronym.) The approach has resulted in computers consisting of SMPs with as few as 4–8 processors and is often used to make up the individual multiprocessor "compute nodes" in the larger systems comprising thousands of nodes. SGI used this approach to apply the SMP architecture to the entire computing platform and produced systems that grew to more than 4,000 processors. This approach was used in several ASCI platform procurements.

The four programming models described above are not a complete set, nor are they mutually exclusive. There are many times that programmers would use a combination of different models as they sought to make both programming and executing the application as efficient as possible.

Ultimately, given this array of programming model choices, ASCI application developers learned over time that the best choices were the ones that allowed applications to be transferred, or "ported," to other architectures as smoothly as possible. In many cases the best practice was to use application-specific libraries, built on top of MPI, which handled all of the explicit communications. As long as the communication libraries were available on other computing platforms, moving applications was made much easier.

Moving Applications to Other Computers

Another hurdle facing application programmers at the Laboratories was the process of porting existing applications to the new parallel architectures. Many of these applications had been used for decades and were trusted to describe the results of underground nuclear tests. These programs had been written for serial computers, and most attempts to directly port these codes to parallel systems were unsatisfactory.

Despite the confidence and familiarity users had in the legacy applications, these codes had been written when generalized agreement with actual testing was required. The codes were not sufficiently detailed to follow exceedingly complex physical phenomena at the scales needed for use in a non-testing regime. Hence, the Laboratories decided in some cases to start over and to create entirely new applications for the new architecture. Often, a fundamentally new approach to the problem was required. Treating the problem as an opportunity, programmers developed new ways to treat the simulation of the physics that took advantage of the parallel operations. SNL demonstrated this particularly well with their PRONTO application, which was used on the world's first teraFLOP/s, massively parallel computer to simulate transient solid dynamics simulation for how nuclear weapon shipping containers would perform under accident conditions.[25]

Need for Flexibility

The Laboratories employed all four of the programming models described above. It was never a question of which model was best; rather, the important question was which model was best suited to a particular problem and the particular computer architecture. The critical factor in many cases was the nature of the physical processes to be simulated.

LLNL's Dale Nielsen wrote one of the articles for the 1992 *MPCI Attack of the Killer Micros Annual Report*. Nielsen discussed a heterogeneous view of computing and programming, not only in the types of computers available and how they were programmed, but also in how individual systems were employed. He wrote, "It has almost always been the case that no single parallel computer would support all of the individual physics packages of a single code because of inappropriate hardware architecture, software programming models provided by the manufacturer, or both."[26] Nielsen concluded that LLNL should adopt a programming approach that was as portable as possible. At Livermore this often involved using some form of MPI as the standard. In a 2005 presentation, LLNL's Mark Seager articulated how ASCI dealt with the issue of programming models. He said, "A key element of the solution was to know where to be rigid (programming model) and where to be flexible (platform architecture and vendor)."[27]

In 2002, DARPA, with the participation of ASCI, launched the High-Productivity Computer System (HPCS) project. One of its primary goals was to reduce the time-to-solution for the use of high-end computing systems, including the efforts needed to create new applications for the computers. The project funded research by IBM, Sun Microsystems, and Cray Inc. to explore combinations of hardware systems and software that will be needed to productively use systems at the petaFLOP/s scale (one petaFLOP/s is one quadrillion FLOP/s, or 1,000 teraFLOP/s). Work in this project still has not settled on an existing or a single programming model for the future computers.

It seemed that the issue of settling on a single, standard programming model was not in the cards. Nielsen's 1992 recommendation was still valid 13 years later. Even with the 2005 deployment of ASCI's newest computing platforms, Purple and BlueGene/L, the issue of parallel programming models had not been fully resolved. This was echoed in 2005 by Nick Donofrio, an Executive Vice President at IBM:

> What it boils down to is that there is really no right or wrong computing and programming model. It depends on how much effort you are willing to commit. The reality is that there may be some applications that cannot and should not be moved to the distributed processing programming model. But to take advantage of the technical advances, it is all about change in the HPC industry. Application programmers have to be willing to change everything. We need to motivate people not to abandon their installed base, but when appropriate, to migrate to another programming and computing model.[28]

The Initiative's success in developing applications was largely due to the shift to parallel programming models. It was the leadership of Vic Reis, Gil Weigand, and the Laboratory managers that helped to drive that change. It did not happen overnight, but with the enduring commitment from funding sources and the intellectual energy of many different code teams it did happen. As a result, not only is science-based stewardship better enabled, but the fields of parallel computing and simulation-science have been significantly advanced. Today, ASCI applications have performed simulations that enable scientists to arrive at deep and significant insight, well beyond anything that Brooks, Nielsen, and the MPCI researchers could have imagined in 1992.

---

[1] Drell, et al., *Science Based Stockpile Stewardship*, 10-12.
[2] Ibid, 4.
[3] Parker, "Strategic Computing Comes of Age," 14.
[4]U.S. Department of Energy Defense Programs,*Accelerated Strategic Computing Initiative Program Plan*, 1996, 11.
[5] James Peery, telephone interview by author, June 20, 2005.
[6] Parker, 15.
[7] Chrzanowski and MacGregor, eds., "The Eighties," in *Serving the Nation for Fifty Years: 1952-2002 Lawrence Livermore National Laboratory Fifty Years of Accomplishments*, 71.
[8] Kusnezov and Soudah, *ASC Program Plan FY05*, 3.
[9] Parker,17.
[10] Post and Votta, "Computational Science Demands A New Paradigm," 39.
[11] Post and Kendall, *Software Project Management and Quality Engineering Practices*, 24.
[12] Bruce Goodwin, interview by author, June 30, 2005, Building 132, Lawrence Livermore National Laboratories.
[13] U.S Department of Energy et al., ASCI Program Plan.(Department of Energy, 1 October 2002), 26.
[14] U.S Department of Energy et al., ASCI Program Plan 2001 (Department of Energy, 1 October 2001), 5.
[15] Michael Ortiz Multiscale Models of Materials: Linking Microstructure and Macroscopic Behavior (California; CalTech, 21 April 2005) 4.
[16] Christian Mailhiot, interview by author, 3 May 2005, telephone interview.
[17] Timothy C. Germann, Kai Kadau, and Peter S. Lomdahl 25 Tflop/s Multibillion-Atom Molecular Dynamics Simulations and Visualization/Analysis on BlueGene/L (Washington; Association of Computing Machinery, 12 November 2005), 1.
[18] Germann, Kadau, and Lomdahl , 8.
[19] Ortiz,18.
[20] Blue Ribbon Panel on Simulation-Based Engineering Science Revolutionizing Engineering Science through Simulation (National Science Foundation, 1 February 2006), 9.

[21] Brooks and Warrens, eds., *The 1991 MPCI Yearly Report: The Attack of the Killer Micros*, 5.
[22] Sterling, *Facing the Multi-Core Fear Factor*, quoted in Larzelere e-mail.
[23] Bobrowicz, *Numerical Experiments*, Appedix.
[24] Ed Barsis, Albuquerque, New Mexico interview by author, May 18, 2005.
[25] Attaway, et al., "Transient Solid Dynamics Simulations".
[26] Neilsen, *General Purpose Parallel Supercomputing*, quoted in Brooks, et al., *The 1992 MPCI Yearly Report: Harnessing the Killer Micros*, 11.
[27] Seager, *LLNL Platform Strategy*, 3.
[28] Nicholas Donofrio, telephone interview by author, August 23, 2005.

Chapter Four
Platforms – Power Plants for Simulations

Computing platforms provide the power that allows applications to simulate the physical world. In many ways, the availability of that power had been the limiting factor in how well those simulations could represent physical events in terms of the included physics, the resolution, and the dimensionality. As ASCI considered its mission to deliver massive new simulation capabilities to SBSS, the need for entirely new levels of computing power was clear. The challenge was to find the best way to deliver that power to the applications.

The following vignettes tell the story of how ASCI, DP, and the National Laboratories worked with the computer industry to build the needed platforms. The vignettes demonstrate how ASCI needed to change the basic approach to building the world's most powerful computers. They also tell the story of ASCI's impact on the HPC industry and how creating the platforms required changing the relationships between the Laboratories, DP, and the broader community of researchers and users.

Elements of a Computational Platform

At its most fundamental, a computer consists of relatively simple components that provide five basic functions:

- **Computational Processing Units (CPUs):** These are the components of the computer that execute the mathematics necessary to support applications. To use a computer, all applications are reduced to simple operations, such as additions, subtractions, multiplications, and divisions, which are executed by the CPUs.
- **Memory:** This part of the computer stores binary data (the "ones" and "zeros", indicating the "on" or "off" status in an electronic memory), known as *bits,* that the computer uses to represent data and the formulas that will be executed by the CPU. Providing access for the CPUs to the data contained in the memory is one of the most important factors controlling how much power a computer can deliver to an application. A very fast CPU will sit idle until relevant data can be delivered from the memory to the CPU for use in calculation.
- **Interconnection Network:** In older computers, there was only one CPU connected to one bank of memory. These were known as serial computers. As computer technologies advanced, several techniques were used to make the connection between the memory and the CPU faster. These included making the connecting wires shorter and using air or liquid cooling to decrease electrical resistance, increasing how fast electrons move in the computer. In the 1980s, the introduction of cheap, silicon-based microprocessors permitted a new approach to designing computers. Instead of a single CPU and bank of memory, several CPUs

could operate in parallel, each with its own bank of memory. An *interconnection network* was required to allow data to be passed from one CPU to another as necessary. How well and fast the interconnection network performed had a very large impact on the efficiency of the overall system.

- **Input/Output:** Another network connection is needed to allow a computer to communicate data outside itself. The Input/Output, or I/O, connects the computer with data stored externally on long-term media (e.g., hard disk drives). I/O is also used to connect computers to other networks, allowing data to be shared with other machines across a room, across town, across the country, or around the world. Finally, the I/O component allows human users to interact with the computer.
- **Operating System:** As a computing platform executes an application to conduct a simulation, another program runs in the background. This is called the operating system and is the software used to control the inner workings of the computer. If the operating system did not exist, the users would be required to deal with many mundane issues like deciding on exactly where to place data coming from the CPU in memory.

These five basic functions can be found in every computer. However, the functions can be provided in various ways, and the complete computer systems that deliver them can be complex, with thousands, if not millions, of parts. There are myriad options that must be chosen to create a computer. Together, these choices are known as the computer's *architecture*. The history of the ASCI platforms is largely the story of architectural choices made while a revolution was taking place in the industry due to the emergence of new technologies.

Measuring Computer Power

With the array of technologies that provided the functions necessary to deliver computing power, it was difficult to quantify the performance of different styles of computers. One simple way of doing that was to measure the specified performance of a system's components. For decades, CPU performance has often been measured in terms of peak FLOP/s, although there was general agreement that peak FLOP/s provides an incomplete picture of a computer's practical applied power. Measuring peak speed does not account for how well a computer can access memory, execute I/O, or how well parallel computers could pass data from one processor to another.

For this reason, computer scientists developed benchmarks to test the performance of computers as they executed actual programs. A wide variety of benchmarks was developed to test how well computers can perform a set of specified functions. The problem with any given benchmark is that while it measures the performance on some applications, it can not accurately measure performance on all types of applications. For example, some applications allow the CPU to use data stored in well-organized, easily accessible locations. Moving these data is very efficient and a computer running a benchmark that measures this type of data access could perform very well. On the other hand, some applications need to access data stored in very unpredictable, random ways.

A computer not designed for this type of memory access might perform poorly on a benchmark designed to measure that access.

Over the years, one benchmark has emerged as a consistent measure of the speed and power computers can bring to bear on scientific simulation problems. Known as Linpack, it measures the sustained rate, in FLOP/s, that a computer can deliver to solve a dense system of linear equations. The benchmark was introduced by Jack Dongarra of the University of Tennessee and has been used to track the performance of the world's most powerful top 500 computers. Dongarra and others have been tracking the performance of these computers since 1993 and publish the results in the semi-annual Top500 list. While there has been ongoing disagreement about the validity of using this benchmark to judge computer performance, the list has provided an excellent view of how computers and institutions have changed over the years. The Top500 list also provided excellent insight into the changes occurring in the 1990s to the industry that designed and built these powerful computer platforms.

Industry in Turmoil

In the early 1990s, the HPC industry was undergoing significant changes. Companies like IBM and Cray Research, which provided most of the computers used to support the analysis of underground nuclear test data, were slow to make the shift to parallel processing. Cray Research was particularly important to the National Laboratories because it had provided most of the HPC systems in the 1980s. Cray Research computers were based on vector processing, and the National Laboratories had become particularly adept at using vector computing to improve the performance of their applications.

As ASCI was beginning, companies that were early innovators of parallel computers were also having problems. The market for these machines had not developed as quickly as expected. While customers were limited by serial computing, they were also reluctant to make the very substantial investment, in both hardware and software, required to adapt their applications to parallel processing. In the early 1990s, companies like Thinking Machines, Kendall Square, and MasPar were going bankrupt or were being acquired by other companies. Even Cray Research, an acknowledged HPC leader trying to shift to parallel processing, would be purchased by Silicon Graphics in February 1996.

To supply the computers that the Laboratories needed to power simulation applications, the Initiative would have to work with the computer companies in a way that enhanced the commercial viability of those organizations. Reis and Weigand understood this very well from their experiences at DARPA. At the 1994 ASCI workshop, Reis made the point of telling the representatives of the nine computer companies in attendance, "It is important that industry tell DOE and the Laboratories what they think are their future directions."[1]

The ASCI Computing Power Goal

During the late 1980s and early 1990s, Laboratory scientists like Randy Christensen (LLNL), Frank Bobrowicz (LANL), and Ed Barsis (SNL) realized the potential power that parallel processing computers could unleash. In separate assessments, all of the Laboratories used early parallel computers for simulations of nuclear weapons performance and were pleased that they could adapt their applications to operate on this style of architecture. Understanding that it was becoming possible for computers to deliver massive amounts of computing power, the question became: how much power would be necessary? In 1994, LLNL's Randy Christensen and Charlie McMillan were asked to calculate the amount of computing power needed to achieve ASCI's simulation goals. Consider what they discovered:

JUSTIFICATION FOR THE NEED FOR 100 TFLOP COMPUTERS
TO SOLVE NUCLEAR WEAPON SIMULATION PROBLEMS

1. To simulate a nuclear weapon primary boost with sufficient resolution (dimensions are classified), today's 2D codes require 500 Cray YMP [the current, circa 1994, generation of supercomputers] hours.

2. Experiments have shown that primary boost is fundamentally 3D and to understand the effects of aging and/or changes in the weapons, they must be simulated in 3D. The move from 2D to 3D increases the compute time by a factor of 1000 to 500,000 Cray YMP hours.

3. Today's codes contain many empirical factors that are based on tests of existing weapons. As weapons age and changes are made, these empirical factors must be replaced by better physics. The calculation of the additional physics is expected to increase the compute time by a factor of ~100 to 50,000,000 Cray YMP hours.

4. A reasonable run time for a weapons analysis code is about 100 hours (4 days). This allows the analyst to remember what they computed and provides enough iteration to support problem solving. Therefore, 500,000 Cray YMP equivalents (50,000,000 YMP hours/100 hours) are needed to support 3D, better physics simulations of primary boost.

5. The peak performance of a Cray YMP is 333 [megaFLOP/s]. Thus the [FLOP/s] needed to support 3D, better physics simulations of primary boost is roughly 167 [teraFLOP/s] (500,000 Cray Equivalents * 333 [megaFLOP/s] per Cray).

NOTE: Please remember that the 167 [teraFLOP/s] figure is based on the problem and does not pre-suppose how the [FLOP/s] will be delivered. It is also based on our current understanding of the problem and the compute power needed do the simulations. As the codes are developed we anticipate that research will provide additional understanding that will probably increase (physics issues) and decrease (algorithms efficiencies) the compute power needed. Hopefully they will balance each other out.[2]

LANL and SNL later independently confirmed Christensen and McMillan's calculation. Hence, after consideration of the efficiencies that could be gained in how the applications were programmed, ASCI set a major goal for computing power at 100 teraFLOP/s.

This was really just an initial goal for the computing power that ASCI needed to deliver. Christensen and McMillan calculated that 100 teraFLOP/s of computing power would allow the Laboratories to cross a crucial threshold. Lower computing power necessarily means the simulations are run at low resolution. Below a certain resolution, it is impossible to determine with certainty if anomalous results are due to errors in the calculation (so-called round-off or discretization error) or due to previously unknown physical phenomena. Christensen and McMillan predicted that at 100 teraFLOP/s it would be possible to distinguish between those causes and the simulation results would not be impacted by errors in how the equations were solved. At that point, it would become clear where the simulations were having problems representing the actual physics of an event.

The following vignettes tell the story of how the National Laboratories and DP worked with the computer industry to build the platforms necessary to meet the goal set for ASCI. Over the decade in which these vignettes take place, the scientists and engineers working on the Initiative demonstrated tremendous innovation, dedication, and partnership to deliver these systems. The budget for the computing platforms was only about 15% of funding provided for ASCI over its first decade. Yet, that money was used in a way that made a significant difference in how the scientific insights would be provided to the stewardship of the nuclear weapons stockpile. These investments also helped change how the Laboratories worked together. Finally, ASCI's platform investments injected energy and funds at a key time to revive the industry that provided the world's most powerful computers.

ASCI Red – *Breaking the TeraFLOP/s Barrier*

In 1996 a computer named ASCI Red became the first to achieve performance of more than one teraFLOP/s, marking an important milestone toward creating the computing power necessary for ASCI's simulation applications. ASCI Red was important because the platform provided clear evidence that ASCI would have a significant impact on the scientific computing world. ASCI Red was also important because the process used to procure it demonstrated to the Laboratories that ASCI would operate ways very different from those of the past.

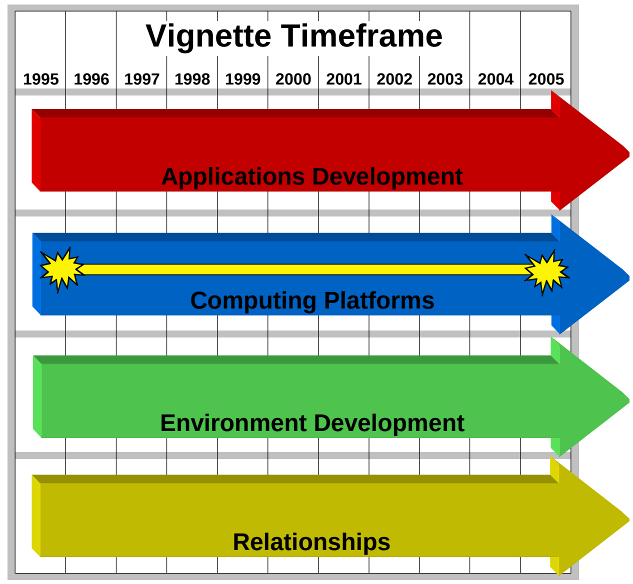

A crucial element of ASCI's success was that it simultaneously advanced a wide array of technologies that would be required to establish computational simulation as a fundamental partner, with theory and experiment, in modern science. This required providing teraFLOP/s computing platforms to the Laboratories as rapidly as possible. Shortly after the beginning of ASCI, a decision was made to buy the world's first teraFLOP/s computer. Named ASCI Red, this computer became an important symbol in several ways. For one thing, it was a harbinger of how ASCI would accelerate the delivery of simulation capabilities. Red would also, even in the decision to procure it, establish how ASCI would catalyze the relationships between the National Laboratories, DP Headquarters, and the computer industry.

Bishop's Lodge

The decision to pursue the world's first teraFLOP/s computing platform was made on January 5, 1995, at Bishop's Lodge, a hotel on the outskirts of Santa Fe, New Mexico. Outside, a light snow was falling, while inside, an early ASCI strategy meeting was in turmoil. The Directors of Computing for the three weapons Laboratories were in attendance; one was thrilled, the other two were in shock, and Gil Weigand realized he would have to call his boss. Weigand had just announced that ASCI's first computing platform would be built using MPP computer architecture and that it would be installed at SNL's Albuquerque site. Considering the extraordinarily rapid and simultaneous innovation across several fronts ASCI was demanding, it is scarcely surprising that the

Initiative's growth was accompanied by shock and disruption. Weigand, in a 2005 interview, recalled his impetus, "Vic [Reis] and I knew that we could not continue the 'status quo' in computing at the Laboratories."[3] Indeed, procurement of the first ASCI platform sent a message that ASCI would be managed in a new way.

It was clear that massive computing power would be needed to use simulations as an integral part of SBSS and the goal of attaining100 teraFLOP/s was set. Major technological leaps were required in several areas, including delivered performance, amount of memory, and I/O capabilities. While the platforms represented only one part of a capability to enable scientific insight through simulation, the high profile, the extraordinary size, and the enormous cost of each succeeding ASCI platform made the machines highly visible symbols of the Initiative, even before they were built. Naturally, they also became the subject of intense and sometimes emotional debate.

On the first day of the meeting at Bishop's Lodge, Weigand had asked to see each Laboratory's proposal for the first ASCI platform. He expected that LANL and LLNL likely would propose an approach based on the current generation of vector computing, as embodied by Cray Research, and on a modest level of parallel computing. Weigand knew that SNL's Ed Barsis, on the other hand, wanted ASCI to procure the world's first teraFLOP/s computer using MPP technology.

Weigand's decision was surrounded by controversy. This first ASCI platform would not be "owned" by any one Laboratory; rather, all three would use it in support of SBSS. It was named "ASCI" Red to reinforce that point, but it would be installed at SNL. Also, Weigand's announcement was made well before ASCI had received its initial funding. Weigand had issued a tall order: to buy a computer that had not been invented using money they did not have.

The announcement of ASCI Red aroused skepticism. It was to be the world's first teraFLOP/s computer, and it was to be delivered well before most people thought possible. The MPP architecture would use thousands of off-the-shelf processors with a high-speed, low latency network coupled to them. Although the concept of MPPs had been around for almost a decade and while SNL had met with some success employing it, it was not at all clear that the MPP architecture could be scaled up to the teraFLOP/s level, and that even if it could, it wasn't clear that it would not be overcome with hardware failures or that it would even be programmable in any practical way.

Impact on the Laboratories

Despite such concerns, Barsis was pleased. He felt SNL's success with earlier MPP computing systems proved their usability—and made Albuquerque the best home for a new teraFLOP/s system based on the MPP architecture. He was sure SNL would be able to make good use of the computer.[4] The computer center directors from LANL and LLNL, Hassen Dayem and Bill McCurdy respectively, were not so sanguine, having come to the meeting with every expectation that they would be getting the first ASCI computer procurements.

Early planning at LANL and LLNL had set the stage for ASCI and this likely influenced Dayem's and McCurdy's expectations. They also represented the Laboratories most adversely affected by the test ban. Weapons scientists at these two Laboratories had relied heavily on underground testing to assess nuclear aspects of weapons safety and reliability. SNL was responsible for the non-nuclear components which generally could be assessed without underground testing; hence, it was not as affected by the test ban as were the other Laboratories.

In 2005, Weigand said he knew he had dropped a bombshell at the Bishop's Lodge meeting: "That evening, I went back to my room and called Vic Reis at home. I told him what had happened and that he should expect a call from the Laboratory Directors [LANL and LLNL] either that evening or in the morning. Vic asked me, 'Was it the right thing to do?' and I told him 'Yes.' He said, 'Fine.'" [5]

The ASCI Red decision was considered bold, even shocking, but it turned out to be a sound strategic move, establishing the pace at which the initiative would move. In the early 1990s, a 7-times improvement in computing power in less than two years would be truly impressive. Pushing for a one teraFLOP/s computer so quickly signaled that ASCI took the "Accelerated" part of its mission seriously.

That quickness also demonstrated that ASCI would commit significant computer resources to support development of simulation at full scale. ASCI would simultaneously push the boundaries of all the technologies needed to build a predictive simulation capability. Programmers would have to learn how to create simulations running on thousands of processors by running simulations on computers with thousands of processors—and the environment developers would have to support both the MPP hardware and the applications programmers at once. The early introduction of the world's first teraFLOP/s computer would send notice to the rest of the world that DP and the Laboratories were quite serious about fulfilling ASCI's mission.

Massively Parallel Processing

ASCI's commitment to MPP architecture was daring, and skepticism about it not unfounded. Based on a simple yet powerful idea, MPP technology combines large numbers of microprocessors and memory systems in a fast, efficient network and focuses all these resources on a single problem. The challenge of this technology was that it demanded an entirely new approach to programming simulation applications. MPP architecture requires a simulation to be split into many "local" programs, each running on an individual processor or a small group of processors. As a local program proceeds, it eventually needs data from other local programs. The program in need then sends a message over the internal network to the other local programs, requesting the data. If the program has nothing else to do, it waits for the requested data to arrive, its processor idle. This situation has the potential for very inefficient use of the processors.

The time a processor waits for data depends on several factors. Has the data been generated? Is it available? How much time elapses in the sending and receiving of the data request, and in assembling the data into a return message? Also important is the location of the data and its "proximity" to the requesting processor. MPP architectures use a variety of "topologies," that is, interconnection pathways among the processors. Data requested by processor A from processor B will arrive more rapidly if there is a direct connection between A and B, but much more slowly if the data must pass through one or several processors "between" A and B in the topology. Programmers writing simulation applications would be required to keep track of how a simulation was divided into local programs, so that all local programs would progress through their problems at relatively the same rate. Programmers also would have to track the location of data to keep them topologically as close as possible to the processors that would need them.

At Bishop's Lodge in January of 1995, Barsis knew the challenges of programming MPPs were formidable, but he believed the approach promised to produce extraordinary new levels of computing power. He proposed to procure a system that would deliver more than one teraFLOP/s to the Linpack benchmark before the end of 1996. Barsis's confidence that a one teraFLOP/s computer would be delivered in less than two years left some of his colleagues at Bishop's Lodge highly skeptical. At the time, the two most powerful computers in the world were a Fujitsu-built system at the Japanese National Aerospace Laboratory, producing 170 gigaFLOP/s on the Linpack benchmark, and an Intel Paragon at SNL, which achieved 143 gigaFLOP/s[6] (a gigaFLOP/s is one billion FLOP/s). These machines reached less than one-sixth of the one teraFLOP/s goal. Many computer scientists did not believe the teraFLOP/s barrier could be broken before the turn of the century, but ASCI expected Red to be delivered three years before that.

Buying the Computing Platform

Once the decision was made to proceed with a teraFLOP/s computer, the challenge of procurement began. Before ASCI, computer procurements were controlled by individual Laboratories; DP Headquarters and other Laboratories offered only limited input. The ASCI Red decision pioneered new methods for later platform procurements. First, however, ASCI had to find money just to start the procurement because ASCI would not receive its first official funding until the fiscal year 1996 budget passed—in November 1995. Responding to the urgency of ASCI's mission, DP headquarters decided to use $15 million of Technology Transfer money for ASCI Red.

Technology Transfer was a relatively new budget line for DP. As the Cold War was winding down, some consideration was given to using the National Laboratories as research and development institutions that would provide technologies to U.S. companies in order to assist the competitiveness of those companies in world markets. Hundreds of millions of dollars were dedicated to this endeavor. For ASCI to use that money a "dual-use" criterion had to be satisfied. Dual-use required that Technology Transfer funds spent at the Laboratories would benefit both the nuclear weapons program *and* U.S. industry. ASCI convincingly demonstrated to DP financial managers that ASCI Red met the dual-use standard; the project could move forward.

Weigand and Headquarters insisted that the Red procurement be handled in the Tri-Lab spirit. The first fruit of this approach was a joint white paper demonstrating the need for teraFLOP/s computing power to support nuclear weapon safety and performance assessment. A Request for Proposals (RFP) was drafted and released on April 4, 1995, just three months after the Bishop's Lodge decision to proceed with Barsis's proposal.[7] Publishing an RFP so quickly was extraordinary and due to efforts by several SNL scientists, including Jim Tomkins, Art Hale, and Stephen Wheat. The RFP allowed bidders only 30 days to submit proposals, yet the industry responded. Amazingly, by the week of May 15, representatives from DP and the three Laboratories were able to evaluate proposals.[8] After a series of discussions with vendors on June 8 and 9, Intel was chosen to build ASCI Red.

The contract negotiations between SNL (on behalf of DP) and Intel were quite complicated. Essentially, a research and development partnership was created between the Laboratories and Intel. After all, a teraFLOP/s system was hardly listed in the Intel catalog. The system would require significant design effort, followed by a period of manufacturing and extensive testing. Though the challenges were substantial, and though all would have to be done on a tightly compressed schedule, the partnership proceeded.

On September 7, 1995, Secretary of Energy Hazel O'Leary announced the Intel contract:

> This agreement marks the advent of a new era in high-performance computing that will significantly benefit our national security. Computers of this scale will unlock the ability to confidently simulate nuclear weapons tests in the laboratory. This effort demonstrated a step forward for our scientific-based [sic] stockpile stewardship program. It emphasizes the Department's commitment to maintain the President's goal of preserving a safe, secure, and reliable deterrent without underground testing.[9]

Red would truly be massive. It would consist of 9,216 Pentium Pro processors, the most powerful Intel microprocessors found in desktop computers at the time. Two Pentium Pros, each with its own bank of memory, would be arranged, along with the interconnection network hardware, on a board the size of a small coffee table. These boards would be held in 85 cabinets, each the size of a large refrigerator. The total system would take up about 1,600 square feet of floor space, or about one-third the size of a basketball court. It would have 594 gigabytes of memory and access to two terabytes of hard disk storage.[10]

Skeptics

The announcement significantly raised the visibility of the Initiative and generated some criticism. Most of it centered on doubt that thousands of processors could be harnessed together to focus their power effectively on a single simulation application. Norris Parker Smith of *HPCWire* wrote a parody of this approach entitled, "Midsummer Night's Nightmare: teraFLOP/s Rex," imagining the computer abandoned in a dusty museum, never having fulfilled its promises and vanquished by SMPs, the shared-memory symmetric multi-processor systems.[11] A prominent university computer scientist echoed Smith's pessimism about ASCI Red in an e-mail, stating, "I also admire your moral courage to 'tell it like it is' while much of your reading community hasn't caught on. I still believe we will get to a teraFLOP/s, but it will be by 1999 and will be built entirely out of the little 'Essempees.'"[12]

Delivering the Computer

Despite the detractors, the contract was signed and Intel started engineering and manufacturing the system. The same Intel team that had produced the early Paragon systems took on the challenge of creating Red but the similarities between the systems largely ended there. SCI Red would use the just-introduced Pentium Pro processor and a newly designed interconnection network. The Intel team would also adopt Puma, SNL's

Figure 4-1. The ASCI Red system at Sandia.

lightweight kernel operating system, for the compute nodes (renaming their version Cougar), and they would use Unix for the nodes that acted as the interface for users and other computers. This approach to the operating system was critical because of the limited memory available on each node. The lightweight kernel approach eliminated much of the operating system activity that had the potential to disrupt application scalability. This also reduced complexity, thereby reducing the likelihood of software failures.[13] The ASCI Red system consisted of 9,216 processors, 640 disks, 1,540 power supply units, and 616 interconnection facilities. Failure of any one of these components with its associated hardware or software could disable the entire computer.[14]

Intel's design and manufacturing team was led by Justin Rattner, Director of their Scalable Server Laboratory. The team was quite experienced with the production of the previous versions of the Intel MPPs. Despite the complexity of the task, on October 17, 1996, little more than a year after Secretary O'Leary's announcement, the team ran the Linpack benchmark at 208 gigaFLOP/s on an 11-cabinet system, making it the world's most powerful computer.[15] That benchmark was quickly surpassed by a 327 gigaFLOP/s run on November 22, 1996.[16] On December 4, 1996, a full teraFLOP/s was achieved at the Intel facility.[17] The whole system was then packed up and delivered to SNL's facilities in Albuquerque. Red was online at SNL by June 12, 1997, when it ran the Linpack benchmark at 1.338 teraFLOP/s.

Weigand was quoted in a June 1997 *HPCWire* article, commending Red's makers: "This machine would not exist without Intel's dedication. Intel had the courage to pursue a goal many didn't think was possible, and they played a decisive role in shifting the industry away from highly customized vector processors to off-the-shelf building blocks."[18] Despite their success, Intel realized that by building MPP systems they were competing directly with their own microprocessor customers; they exited the HPC systems business within the next year.

With the ribbon cutting ceremony, ASCI Red entered into a long and distinguished service. The platform provided the vast computing power necessary for many simulations enabling scientists to better understand the nuclear weapons and their accompanying support systems. In November 1997, when Red had only been fully operational since June, a group of SNL scientists won the prestigious Gordon Bell award at the SC|97 conference. The Gordon Bell award annually recognizes the most important contributions in HPC applications. In just the few months that the Red system had been available, the SNL team ran a transient solid dynamics simulation application named PRONTO to develop insight about how nuclear weapon shipping containers would perform under accident conditions.[19]

In 1999. the system received an upgrade to its processors and memory boosting its power to 3.15 teraFLOP/s.[20] During a 2005 interview, Barsis commented on the ASCI Red process: "It was very successful in a number of areas. There were several technical innovations. First, we really tried to understand the reliability issues. We also designed in an upgrade path by over-sizing the network so we could switch processors. Finally, we wrote the contract in a way that if the machine did not meet the specifications, Intel would be obligated to [continue to] deliver hardware until it did." Barsis also described Red's impact on ASCI's mission: "Red has been used to solve problems that could not have been solved any other way. This included assessing safety issues and looking at thermal diffusion in nanodevices."[21]

Red Storm

ASCI Red was so successful that, when it became clear that the computer was about the reach the end of its useful life, ASC decided to purchase a new system built on similar architecture. In July 2004, SNL contracted with Cray to deliver Red Storm, a 41.5 teraFLOP/s system. Red Storm would also take advantage of off-the-shelf commercial technologies and would consist of 11,648 AMD Opteron processors, each with two banks of memory. The interconnection network consisted of a three-dimensional mesh using specially designed chips known as SeaStar.[22] Because the Opteron processor used the same instruction set as the Intel x86 microprocessors, and because the operating system, compilers, and MPP programming model were similar to those on Red, the National Laboratories would be able to simply move their application codes on to Red Storm—and have them run. It was a final gift from Red.

Figure 4-2. ASC Red Storm at Sandia.

Fulfilling Its Promises

Before ceasing operation in 2005, ASCI Red established many firsts for the Initiative. It was the first major Tri-Lab procurement—even before ASCI had received official funding. It was the Initiative's first large-scale computing platform. As a result, it exposed technology issues concerning the usability of this class of computer, including the need for larger, faster networks, and for large-scale visualization systems.

ASCI Red served for nearly a decade. During that time it delivered millions of processor-hours to a wide range of simulations, helping scientists to answer a great many important questions. Norris Parker Smith's dire prediction that Red would end up as a forgotten, useless hulk collecting dust in a computer museum was, apparently, completely incorrect. It turned out the bombshell decision Weigand made at Bishop's Lodge was, as he had assured Reis, the right thing to do. The approach ASCI followed to get the world's first teraFLOP/s computer built well before most computer scientists thought possible—with its enforced sense of urgency spawning an exceptional record of innovation—laid a strong foundation for ASCI as a program. It showed the way for much of the extraordinary decade that followed, a decade in which ASCI repeatedly expanded the very definition of high-performance computing and fundamentally changed the way scientific investigations are conducted.

ASCI Blue Pacific and Mountain – *Keeping the Industry Viable*

ASCI broke the teraFLOP/s barrier with ASCI Red at SNL, but that was not nearly powerful enough for the simulations required in stockpile stewardship. ASCI needed to dramatically increase computing power at the LANL and LLNL, as well. To begin this, the Initiative would work closely with computer vendors to build two additional multi-teraFLOP/s platforms. These systems would be the first ones built using a strategy of allowing the companies to deliver commercially attractive systems that were then combined into ASCI-scale resources.

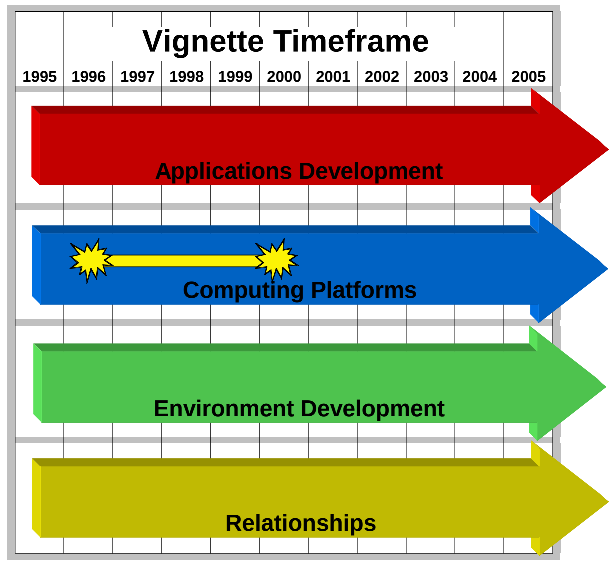

ASCI Red demonstrated that the Initiative would live up to the "Accelerated" part of its name, particularly in computer development. But Intel was getting out of the HPC business, and DP and the National Laboratories could not build the future systems by themselves; partnerships with other commercial computer manufacturers were necessary. Just as Red had demonstrated the effectiveness of pairing SNL's vision with Intel's design and production capabilities, the next two ASCI machines would not only deliver astounding computing power much sooner than had been thought possible, they would also establish mutually beneficial long-term relationships between industry and the National Laboratories.

The Curve

A graph appeared in a 1995 draft of the ASCI Program Plan. It became unofficially know as the "Curve," and continued to appear in plans and presentations throughout ASCI's existence. The "Curve" helped everyone in the Initiative visualize how far and how fast ASCI would have to move computing performance standards if they were to fulfill the fundamental mission for SBSS.

The graph's "Status Quo" characterizes computing power growth as predicted by Moore's Law. Gordon Moore of Intel had made the observation in 1965 that the number of transistors that could be placed, with minimum cost, in a semiconductor integrated circuit had doubled every 12 months and predicted that it would continue to do so. This observation was dubbed "Moore's Law" in 1970; in 1975, Moore altered his prediction

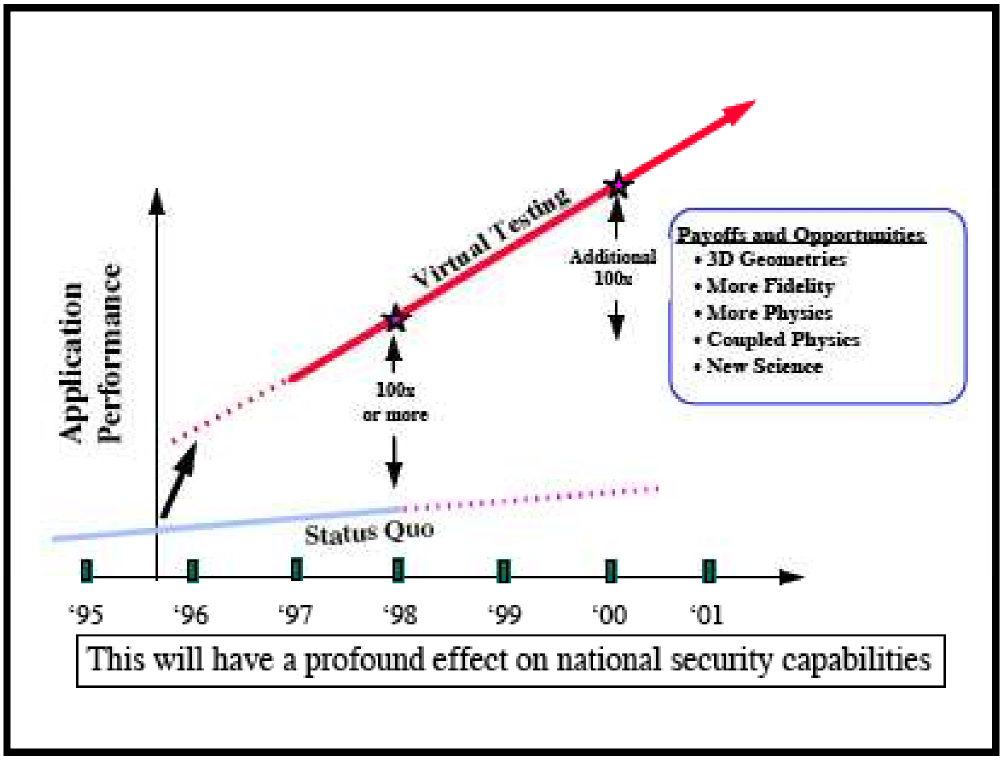

Figure 4-3. The ASCI "Curve" as it appeared in the 1996 *ASCI Program Plan*.

to be a doubling every 24 months. Observing that the increasing performance of the transistors would combine with the increasing density of transistors, an Intel colleague of Moore's modified the statement to the form most widely known: *for approximately constant cost, the performance of the transistors on an integrated circuit would double approximately every 18 months.* However, even though the computer industry was proving Moore's Law with steady geometric gains in performance, ASCI realized it would not be fast enough to meet their programmatic requirements and timelines.

Accelerating the Delivery of Computing Power

"Fast enough" for SBSS and ASCI meant satisfying the two constraints of creating the simulation tools 1) before the aging of the weapons altered their performance in unknown ways, and 2) before all the scientists with undergrounding test experience retired and would be unavailable to interpret the simulation results. It was a race against time. Nobody knew how rapidly the aging weapons might change, but the second constraint was rather predictable. It was expected that the last wave of scientists with testing experience would begin to retire around 2005, a decade after the onset of ASCI.

Accordingly, ASCI set the target at 10 years, giving themselves one decade to create a comprehensive simulation capacity. Within that decade, ASCI would have to supply computers capable of 100 teraFLOP/s. Moreover, ASCI would have to supply the necessary simulation codes themselves and the tools to gather and interpret the simulation results. Events would soon prove, as had been the case with the development of ASCI Red, that ASCI was serious about accelerated development.

TeraFLOP/s Summit

In early 1995, John Morrison was working for Hassen Dayem, director of the LANL computing center, on the initial planning for ASCI. Dayem was at Bishop's Lodge in January 1995 and was stunned by Weigand's decision. After that meeting, Morrison took on a more prominent role in the planning for the ASCI platforms. In 2005, by then Division Leader for Computing, Communications, and Networking at LANL, he reflected on the decision to put the first platform at SNL, "It was real interesting, after the fallout of the Bishop's Lodge meeting it was not only Hassen that was upset. Bill McCurdy at Livermore was also upset. Essentially what Gil [Weigand] told them in response was, okay, Los Alamos and Livermore, you guys work together on the next computer we buy, figuring that we could not do it."[23]

One has to suspect that Weigand knew that LANL and LLNL could, in fact, work together. He was counting on it. A meeting was convened at Gil Wiegand's request on March 9, 1995 (two months after the Bishop's Lodge decision was announced, but before the plans for the machine were finalized). Weigand felt the urgency of ASCI's mission and knew that if meaningful simulations were to happen, the three Laboratories had to get moving on teraFLOP/s platforms. Focusing on the high end of computing, he wanted the Laboratories' computing directors—Barsis, Dayem, and McCurdy—to be there, along with representatives having a focus on weapons.

The summit convened in the IBEX room at LLNL. It is a long conference room, where, as might be expected, there is a standard "white board" at one end. From the moment the participants were seated, a discussion of different architectural approaches was launched, and quickly grew to be a lively, free-wheeling exploration of a vast landscape of high-tech concepts: multiple processors, shared memory nodes, network interfaces, and many more. Eventually, someone suggested that for clarity's sake the two main architecture proposals under discussion be written on the white board so everyone would know what was being discussed.

SNL's Barsis picked up a red marker and wrote out the MPP approach, the basic architecture that would later become ASCI Red. When Barsis was finished, Livermore's McCurdy got up to outline LANL's and LLNL's thoughts on an alternative architecture. Picking up a blue marker, he sketched out a machine that would consist of a number of clustered SMPs. While the MPP approach had a single processor on a compute node, an SMP node could have many. On an SMP node, several processors would be attached to one memory subsystem, all sharing equal access. Each node would therefore have more memory accessible to each processor. The SMP nodes would then be interconnected with a network to allow clustering into more powerful systems.

In 2005, Morrison recalled how the LLNL and LANL contingents felt at the time about Barsis's MPP proposal, "Ironically, [we] agreed that although MPPs formed an effective architecture, we did not think that the marketplace was going to support them. So we made a conscious decision to use clusters of Shared Memory Processor (SMP) computers, figuring that they would be more commercially viable."[24]

Mitigating Risk

Clearly there were risks with both the MPP and the clustered SMP paths. It was not clear that one style of architecture could meet the separate simulation needs of the three Laboratories. It was also not clear which of the approaches would be viable at the teraFLOP/s scale. Furthermore, in 1995 it was not clear which networking technologies would be adopted by industry, and those that were available were very expensive.

The group of scientists sat in the IBEX room pondering these issues, staring at two possibilities sketched on the white board. Suddenly, the answer was obvious: ASCI had to execute *both* approaches. But there was one problem: money. The budget figures presented to Reis in February 1995 accounted for just one architecture. Weigand spoke up and said that he would take care of adjusting budgets to find the money. He did insist that all the platforms be viewed as “ASCI platforms” and not be referred to as the SNL, the LANL, and the LLNL machines.

ASCI was, he pointed out, a single program involving three Laboratories. Rather than talking about the Sandia architecture, or the Los Alamos/Livermore architecture, he suggested the architectures, and later the machines themselves, be given names unrelated to the Laboratories. Once again, the white board made the solution obvious. The MPP proposal, written in red ink, would become ASCI Red, and the clustered SMP proposal, described in blue ink, would become ASCI Blue.

Partnering with Industry

ASCI needed vast amounts of computing power and needed it fast. To get that power from an ailing HPC industry ASCI would have to create partnerships. But why were relationships with manufacturers suddenly deemed so important? After all, the government had been buying high-end computers for years. As things turned out, this fact was contributing to the problem.

Vic Reis had summed up the need for ASCI openness and cooperation during the September 1994 workshop. At one point during his talk he said, “The high end of the supercomputing industry was built to respond to defense needs and with support from the DOE weapons Laboratories, but with resources shrinking, action is necessary or the nation will be out of the supercomputing business.”[25]

Reis had spent time at DARPA. He knew that the performance demands placed on computer manufacturers by the defense community and the National Laboratories far exceeded the needs of commercial and industrial users. Until the mid-1990s, computer companies designed systems to suit their best-paying customer—the government. Unfortunately, those systems were not necessarily attractive to commercial users. When, as Reis described, the defense market shrank due to the end of the Cold War, the consequences of the business model focusing on the single customer (the government) became clear. Within a year of ASCI’s Santa Fe workshop, four of nine companies that attended would either declare bankruptcy or leave the supercomputer business.

This was a problem—and not just for the employees and shareholders of those companies. SBSS needed a healthy HPC industry to keep the nuclear stockpile safe. As the scientists and engineers of the Initiative turned to the question of the next computer platform, they knew it would have to be scaled to the performance levels needed by the

Laboratories but still be commercially viable technology. This would be the approach taken for building the ASCI Blue machines.

The ASCI Blue computer procurements would establish a new sense of partnership between the Laboratories and the HPC industry. Such relationships between the Laboratories and manufacturers would be critical to meeting platform performance targets. The Laboratories were contracting for systems that did not exist and had not even been designed. In fact, due to the development process, most of the ASCI platforms would take three to four years to be fully operational. Computer companies would have to make commitments to the performance and cost of the systems well beyond the technology horizons. It was with this challenge that the ASCI Blue procurement process started.

The Blue Procurement

The Blue procurement demonstrated, as had the earlier procurement of Red, that ASCI was serious about keeping a Tri-Lab focus. On August 14, 1995, an announcement was published that Defense Programs and the National Laboratories were soliciting ideas for delivery of a 3 teraFLOP/s system, and a Request for Proposals followed on February 20, 1996.

The ASCI Blue RFP was innovative in several ways. While it naturally specified mandatory requirements, it also laid out a series of more flexible goals or targets for the computing platform. Computer companies were to meet as many as possible but had the latitude to make compromises based on technology or cost restrictions. Proposals would be evaluated on the basis of how well the sum total of the goals was met.[26]

The procurement was also innovative because it envisioned not just one computer but at least two: an Initial Delivery (ID) system that would prove the concept and be used for rapid application development, and, two years later, the final Sustained Stewardship TeraFLOP/s (SST) platform. Nominally a 3 teraFLOP/s machine, the SST must deliver 1.0 teraFLOP/s on a hydrodynamics benchmark code because that would be more representative than Linpack of applications used by LLNL and LANL. The RFP also envisioned a Technology Refresh (TR) for the ID system to keep it on the leading edge during the long SST development timescale.

The RFP required that proposals be submitted on March 26, 1996. However, in late February, Silicon Graphics Incorporated (SGI) acquired the financially troubled Cray Research. As both companies were potential bidders on the project, an extension was granted so that they could craft a new joint proposal.

Evaluating Proposals

The proposal evaluators from DP and the National Laboratories met at LANL in May of 1996. They quickly focused on two interesting proposals, though for very different reasons.

One was from IBM, who proposed a 3 teraFLOP/s system based on the SP-2, their existing SMP technology. The SP-2 would be scaled up (i.e., employ a faster interconnect) and scaled out (by placing more processors in each of the SMPs). IBM's approach was to base the ID system on single processors and networks of existing SP-2 systems for rapid delivery. The TR system would be deployed a year later with four-processor shared-memory nodes and the same interconnect as the ID system. Eventually, the number of processors per node would be increased to eight, and a new interconnection network would be deployed for the SST platform.

The other interesting proposal was from SGI. They proposed to cluster Origin2000 computers (each a 32-processor SMP) interconnected with HIPPI (High-Performance Parallel Interface) for the ID system. Their SST machine was based on scaling up the SMP to 4,096 processors with a cache-coherent, distributed, shared-memory architecture. This approach offered much larger shared-memory nodes and used special networking and operating system technologies that would provide users with a "single system image." Every processor on the computer would be able to directly access any memory location without having to generate an explicit message.

Both proposals were technically aggressive and provided solid plans to reach new levels of computing power, albeit with different architectural approaches. While IBM took a somewhat more conservative route, their plan was attractive because it was backed by IBM's decades of solid computer engineering. The SGI plan was technically very appealing because it offered the opportunity to have thousands of processors sharing the same memory space. But was their approach too technically aggressive? And would the new SGI/Cray Research marriage introduce coordination difficulties?

Two Blue Machines

There was only money in the Initiative's budget for one ASCI Blue system. Based on the recommendations of the Laboratories, Reis and Weigand decided to proceed with both the IBM and SGI systems; they would revise the budget to cover the costs.

There were several compelling reasons behind Reis and Weigand's decision. First was security. In 1995, networking technologies were not secure enough to allow nuclear weapons scientists at one Laboratory to productively use the computing platforms at another. ASCI had to provide secure, readily accessed computing platforms to programmers at both LLNL and LANL. This would accelerate their ability to write tera-scale applications. Furthermore, two systems would provide additional risk mitigation. ASCI, for the nation's sake, truly could not afford to fail.

Once it was decided to procure two ASCI Blue computing platforms, it was quickly resolved that LANL would get the SGI system and that LLNL would get the IBM.

Figure 4-4. ASCI Blue Mountain at Los Alamos.

This was the beginning of a long and productive partnership between LLNL and IBM. Of course, if there were to be two Blue systems installed, they needed distinct names. The IBM system at LLNL became Blue Pacific, and the SGI system at LANL became Blue Mountain.

Figure 4-5. ASCI Blue Pacific at Livermore.

Delivering the Platforms

Both systems were fully delivered in early 1999, though not quite as originally proposed. During the intervening years, difficulties were encountered with the technology development required for both SST systems. In the spirit of partnership, compromises were made on issues like the processor type, the number of processors per shared memory node, and the speed and capacity of the interconnection networks. The final Blue Mountain, for example, consisted of 48 Origin 2000 computers, each with 128 shared-memory processors, for a total of 6144 processors. The individual machines were interconnected using the HIPPI networking technology.[27] Blue Pacific, as eventually deployed, consisted of 5,856 processors with 4 processors per shared memory node and interconnected with IBM's high-speed TB3 networking switch.[28] The flexibility in the RPF was rewarded with successful platforms, as both systems peaked at more than 3 teraFLOP/s on the Linpack benchmark.

On the Curve

Along with Red, the ASCI Blue machines helped the Laboratories and ASCI leap onto the "Curve." One sign of this, to which ASCI's scientists and engineers could look with some justified pride, was the semi-annual Top500 supercomputer list. In November 1995, as the ASCI procurements were just starting, the list showed SNL's Intel Paragon in second place at 142 gigaFLOP/s of computing power. LANL's Thinking Machine was in tenth place running at 59.7 gigaFLOP/s, and LLNL's Cray Research T3D was well down on the list but still delivering a respectable 25.3 gigaFLOP/s[29].

By November 1999, the Laboratories had leaped up to first, second, and third place on the lists. ASCI Red at SNL had just undergone an upgrade and was delivering 2.3 teraFLOP/s, Blue Pacific at LLNL was delivering 2.1 teraFLOP/s, and Blue Mountain at LANL was delivering 1.6 teraFLOP/s.[30] In just four years, ASCI, in cooperation with industry, accomplished a 90-fold increase in computing power.

ASCI put the Laboratories on the "Curve," but there was more to the mission than delivering phenomenal computing power. The Red and Blue procurements also nourished critically important relationships. Each of the National Laboratories worked in a whole new way with each other and with DP. The procurements made the "One Program – Three Laboratory" strategy a reality. Doing business in this new way may have been disconcerting for some, but results clearly indicated that ASCI could accept nothing less than a full Tri-Lab effort in order to deliver a simulation capability that would enable scientific insight. The design, delivery, and deployment of the ASCI Blue platforms also established important partnerships between the Laboratories and computer companies that would pay dividends over the decade to come. Without those partnerships, it is doubtful ASCI could have so quickly ascended the "Curve."

Linux Clusters – *Providing Cost Effective TeraFLOP/s*

Using simulation for scientific investigation places varying demands on a computer system. Occasionally scientists will need most or all of a platform's power applied to a single problem (*capability* computing). At other times, they run a large number of jobs, each using small amounts of power, but in aggregate requiring huge amounts of power (*capacity* computing). While the Red and Blue platforms delivered capability computing, ASCI developed capacity computing through clever and wise technology investments, notably by boosting cost-efficient, Linux-based "clusters" to terascale performance.

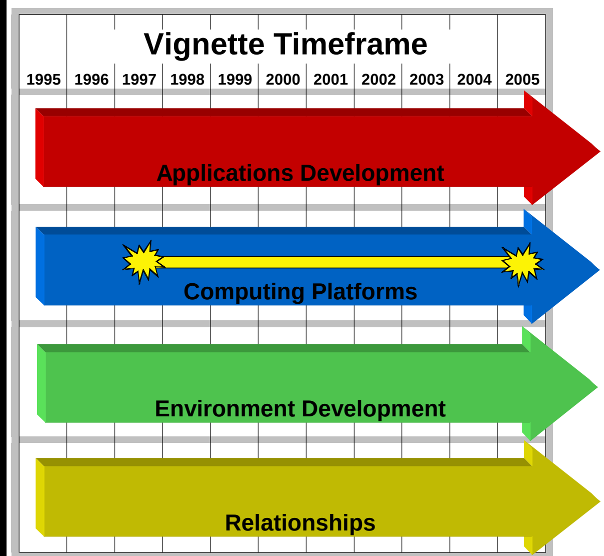

The ASCI Red and ASCI Blue platforms gave the Initiative a solid start toward fulfilling its mission. Novel management, Tri-Lab cooperation, and innovative partnerships with manufacturers all demonstrated the commitment of ASCI's individual members to the ultimate goal. Without question, however, advancing processing power at such a rate was proving to be expensive. By the late 1990s, ASCI leadership felt a need to create more cost-effective computing systems, and to develop some competition for ASCI's existing vendor-integrated platforms. As the decisions to install both MPP and SMP platforms showed, architectures that can deliver terascale power do not all have to be the same, and ASCI was about to venture into another dramatic new area of high-performance computing: terascale Linux clusters. A *cluster* is a group of linked computers working together so closely that in many ways they act as a single machine. Clusters have been around for many years, and the invention of the cluster has been ascribed to many different organizations. That clusters could be linked to form supercomputers, and that this could be cost-effective, are ideas of the mid- to late 1990s, and were enabled by the confluence of several trends, most notably the dramatically dropping price of low-end commodity computers, the development of high-speed high-capacity networking capability, and the rise of Linux, an open-source, public-domain, portable operating system.

ASCI's built-in flexibility allowed it and the Laboratories to become pioneers in developing, deploying, and employing cost-effective terascale Linux cluster computers. Much of the Linux cluster development was built on research done during the early 1990s by the Laboratories, NASA, and others to explore networking many commodity desktop computers together to create one high-performance system. Initially, there were significant technology shortfalls, or gaps, that kept these systems from performing at terascale levels, but those gaps would not go unfilled.

In fact, Linux clusters have been so successful that by 2005 these systems represented 72% of the systems on the Top500 list and accounted for nearly 49% of the cumulative 2.29 petaFLOP/s of the entire list. They are accelerating the use of computational simulation as a peer to theory and experiment, not just for nuclear weapons, but for a wide range of diverse research programs.

If ASCI Red, Blue Pacific, and Blue Mountain were successful, why did ASCI need to worry about additional capability? Further, if Red and the Blue machines justified the huge monetary investment expended in creating them, why did ASCI need to seek more cost-effective platforms? The answers to these questions turn out to be quite complex, closely linked, and involve several factors.

Capability vs. Capacity

It takes a lot of computing resources to deliver simulations sufficiently detailed to permit basic scientific discovery. When a problem from continuum physics, such as the physical operation of a nuclear weapon, is posed on a computer, it must be discretized, or sampled, and converted into huge systems of equations involving many discrete variables. This process naturally involves truncation and estimation of quantities, which produces uncertainties in the discrete system and its solution. To be reliable, the simulation must be run at a sufficiently fine scale to force these natural errors to be so small that they are insignificant. Christensen and McMillan determined that a minimum of 100 teraFLOP/s of sustained computing power would be needed to *begin* to simulate the operation of a nuclear weapon with sufficient accuracy to overcome the discretization uncertainties. This estimate represented the required computing *capability*. In this setting, a system's *capability* refers to how much computing power can be brought to bear by that system to work on a single problem. There are some simulations that cannot even be attempted unless there is adequate computing capability. ASCI had to focus early on rapid delivery of high capability so that scientists and engineers with underground nuclear test experience could validate the simulations.

There is, however, another factor affecting how much computing is required to enable scientists to achieve true scientific insight: the number of simulations that must be run. A measure of a computer's ability to run large numbers of smaller simulations is called *capacity*. For using simulation to facilitate discovery, capacity is just as important as capability.

For instance, thousands of small simulations, or jobs, may be run in the course of writing and testing simulation applications. Developers typically build an application by writing simulation code during the day and then extensively testing the code overnight in a process known as regression testing. One application may be run on many different problems to ensure that mistakes, or bugs, were not introduced with newly added lines of code. Bugs in a code can be extraordinarily difficult to find and remove; for this reason developers typically do not make further changes to a code if the previous set of changes has not passed the regression testing.

Another task that requires running many small jobs, and therefore high capacity, is sensitivity analysis. This is a process that determines how sensitive the results of a simulation are to variations in the input data, and seeks to discover whether small changes in the input simulation parameters will make large, possibly unexpected, changes in the output. A common, albeit rather "brute force" approach is simply to run the simulation many, many times, making small changes to the input. This approach requires significant capacity. Other approaches use intricate mathematical analysis superimposed atop the simulation, which also requires significant capacity.

Balance

For ASCI to be successful, its computational systems had to provide the proper balance of computing capability and capacity. The challenge for DP and the Laboratories was that this had to be done with limited budgets. The 2003 JASON study, *ASCI Requirements*, reinforces this point in its conclusion:

> To summarize, our major concern is in providing as soon as practical much needed capacity to the ASCI program lest the resulting very large oversubscription in *Capacity* becomes unmanageable. In addition, a road map must be developed to deliver to the program machines of requisite *Capability*.[31]

Another factor driving ASCI toward lower cost architecture was the need to establish credibility of the idea that scientific discovery could be facilitated by computational simulation. Although ASCI's work on stockpile safety and performance was essential to national security and often classified, it did not exist in a vacuum. The Initiative's researchers belonged to a broader scientific community. As ASCI's efforts to accelerate computational capabilities began to bear fruit, the issues of how to demonstrate the validity of the resulting science and how to share the capability surfaced at the Laboratories.

Institutional Computing Demands

The three Laboratories focus primarily on nuclear weapons, but they also support a number of other scientific research programs funded by a variety of government agencies. For example, programs in the high-energy physics, biology, climatology, geosciences, computer science, mathematics, and other areas are funded by the Department of Energy's Office of Science. Another funding partner for the Laboratories is the Department of Defense, supporting programs in conventional munitions, as well as war fighting capabilities. The Department of Homeland Security funds projects ranging from counterterrorism to nuclear non-proliferation to information analysis, among others. Each Laboratory also has a small amount of internal funding, under the Laboratory Directed Research and Development (LDRD) program, set aside for general research. As ASCI started to deliver on the new computational capabilities it became apparent that many

Laboratory researchers investigating these other programs also needed ASCI-level computational capabilities to do their science.

Involving non-nuclear-weapons scientists was important for ASCI. Nuclear weapons science is by necessity classified; yet, a way was needed to convince scientists in other domains that simulation could provide a means to scientific insight, supplementing the theory and experiment that had served for centuries. Dona Crawford, Associate Director for Computation at Livermore, captured it perfectly in a 2005 interview, "If only weapons scientists had had access to the big machines, the rest of the Laboratory would never have bought us saying 'simulation was a peer to theory and experiment.'"[32]

The Laboratories needed, and wanted, to provide ASCI-level computing to non-ASCI researchers, but the problem, once again, was funding. Funding for other researchers was spread out among a number of other DOE programs as well as non-DOE sources. The Laboratories had to find a more cost effective means to deliver terascale computing platforms. Fortunately, a solution was already underway. The same MPP architecture of the late 1980s and early 1990s that had inspired ASCI Red provided the answer.

Beowulf Clusters

Several companies took the MPP approach to system design and employed as many commercial-off-the-shelf (COTS) parts as possible to take advantage of economies of scale. The resulting systems largely consisted of parts from desktop PCs. Using large quantities of parts designed originally for commercial computers made building system cheaper. By the early 1990s, researchers were asking if it was possible to build high-performance platforms by networking together the PCs themselves to work on a single large problem. In 1994, NASA researchers Donald Becker and Thomas Sterling at CalTech started the Beowulf Project to examine this possibility.

Becker and Sterling used clusters of PCs running the open-source Linux operating system. Their first system consisted of sixteen, DX4 processors that communicated via an Ethernet network.[33] Just like the ASCI Red and ASCI Blue systems, the Beowulf processors would use message passing to exchange the data among the processors.

Figure 4-6. Typical Beowulf cluster configuration.

Theoretically, the Beowulf approach to cluster PCs could scale to the multi-teraFLOP/s levels required for both ASCI's capability and capacity needs, but in the

1990s that was unconfirmed and key technologies were missing. Fortunately for ASCI, a program called PathForward was already in place and primed to address such gaps.

PathForward

PathForward was launched in ASCI's early stages to deal with several computing platform technology gaps by funding development in those areas. PathForward initially focused on the very large ASCI systems, such as ASCI Red and the two ASCI Blue machines. PathForward later focused on the SMP cluster systems that would be known as ASCI White, ASCI Q, and eventually, ASC Purple. Fortuitously, many of its solutions could also be applied to Linux clusters.

In late 1997, the National Laboratories and Defense Programs formed a team which included Paul Smith, Norm Morse, and Gary Kent of DP; Karl-Heinz Winker of LANL; Art Hale of SNL; and Bob Deri and Dick Watson of LLNL.[34] The team evaluated whether hardware and software technology trends would lead to the performance needed to meet ASCI's platform component requirements in the necessary time. They identified which critical technologies would fall short; those then became the focus of PathForward funding.

A number of technologies met the team's criteria. There were hardware gaps such as the interconnection networking technologies required to supply processors with data from both near and distant system memories. This was considered particularly vital since a parallel computer's interconnection was often a bottleneck, potentially paralyzing overall performance. Software development lagged in areas such as distributed operating systems and the environments that would allow programmers to create applications spread across thousands of processors.[35] To address these obstacles, LLNL conducted, on behalf of ASCI, a comprehensive PathForward procurement in 1997.

One of the first PathForward contracts pushed for development of interconnection technologies. Among the first results were advancements made to the Quadrics interconnection network, a technology used for the 20 teraFLOP/s ASCI Q capability platform installed in the early 2000s at LANL. Over the following years, PathForward contracted to advance areas in which performance was too sluggish for ASCI's needs, including the development of the Lustre file system and the porting of software tools (e.g., the TotalView debugger) to the Linux operating system. Though PathForward technologies initially were to be applied to ASCI's large, single-job-concentrated, capability-focused platforms, many of its contracts were also applicable to ASCI's large, multiple-job capacity platforms. Finally, PathForward paid further dividends when previously-contracted technology advances were also used to scale Linux cluster systems to unprecedented levels of performance.

Multi-programmatic Cluster Resource

While PathForward was important for delivering some important enabling technologies, somebody still needed to determine if the cluster approach would scale to teraFLOP/s. In a 2005 interview, Mark Seager, LLNL's lead for the Platforms Strategy Team, described how LLNL dealt with that issue:

> In late 1999, we did a gap analysis to examine the current state of Linux clusters. We decided that we needed people in-house that could build

> these things. We started taking people from the UNIX group and added them to the Linux group. By the winter of 2001 we thought we could field the first cluster.

LLNL's first clusters resulted from the Parallel Capacity Resource (PCR) procurement, resulting in two systems, a 128 dual-processor node system and an 88 dual-node system. These successful systems led to a much more ambitious undertaking—the Multiprogrammatic Capability Resource (MCR). Seager's 2005 comments noted that:

> We built the MCR cluster… its success convinced the users that Linux clusters could be useful.[36]

Figure 4-7. Multiprogrammatic Capability Resource (MCR) at Livermore.

Users were convinced. LLNL's MCR was among the world's first systems using Beowulf PC-cluster architecture to deliver multi-teraFLOP/s power. Interestingly, the funding for MCR did not come from ASCI. Livermore management, appreciating the need to support the Laboratory's non-ASCI research, created LLNL's Multiprogrammatic and Institutional Computing (M&IC) program by collecting funds from those various non-ASCI research programs that would benefit, and invest them into developing what became the MCR. An RFP was issued in May 2002. Linux NetworX, the successful bidder, delivered a rack-mounted cluster system consisting of 1,152 nodes (or PCs) with two Intel Xeon processors per node. The MCR took advantage of the Quadrics interconnection network, one of the technologies earlier nurtured by PathForward. It also used another early PathForward investment, the Lustre file system, to keep track of data. Fully installed in March 2003, the computer was able to deliver a peak 11.2 teraFLOP/s of computing power, with a sustained 7.63 teraFLOP/s employed on the Linpack benchmark. The MCR, in some ways a big pile of desktop PCs, was then the fifth most powerful computer in the world.[37]

Clusters for Weapons Simulations

The MCR was a success, but it was dedicated to non-weapons research. LLNL was impressed enough, however, to use ASCI funds to contract with IBM for the ALC (ASCI Linux Cluster) system. It consisted of 960 nodes with two Xeon processors each and using the Quadrics interconnection network. It delivered a peak compute power of more than 9 teraFLOP/s.[38]

By 2005, LLNL's most powerful Linux cluster was focused, like the MCR, on non-nuclear weapons research. Thunder, built by California Digital in 2004, consisted of 1,024 nodes, each with four Intel Itanium2 processors. A PathForward investment, the Quadrics interconnection, was once again put to use to provide node to node networking. Ultimately, Thunder provided peak computing power of 22.9 teraFLOP/s.

Mirroring LLNL's success, LANL and SNL also recognized how powerful and cost effective the Linux cluster approach could be. In 2003, LANL contracted with Linux NetworX for Lightning. This system was funded by ASCI and consisted of 1,280 nodes with two AMD Opteron processors per node. Myricom provided the interconnection network. At peak, the computing platform delivered over 11 teraFLOP/s of computing power. A LANL press release quoted John Morrison, head of the computing division at the Laboratory, on the promise of clusters: "The Lightning system that Linux NetworX offers is a cost-competitive way to meet our growing need to run large, important calculations and get results in a few days. A system of this magnitude will provide a valuable proving ground for large-scale, practical cluster computing, building on the exciting development of open-source tools by the larger high-performance computing community."[39]

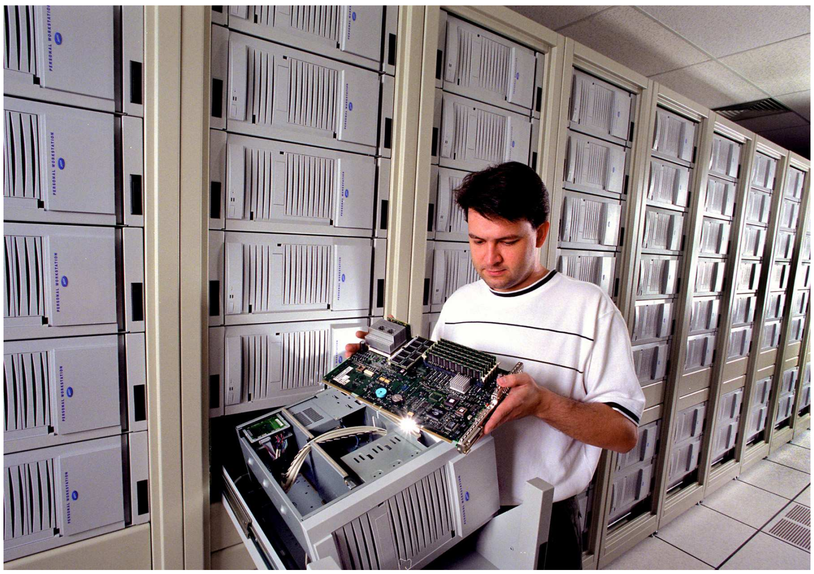

Figure 4-8. CPlant at Sandia.

During the late 1990s, SNL experimented with clustering distributed commodity computers in a project known as CPlant. CPlant used a variant of Linux to build on ASCI Red's architecture by clustering Compaq workstations equipped with Alpha processors that were developed by the Digital Equipment Company (DEC). A special adaptation to its operating system, called "portals," reduced the overhead required to send data from one processor to another.

Then, in 2005, SNL purchased Thunderbird from Dell. The system boasted 4,096 nodes, each with two Xeon64 processors. A Cisco InfiniBand interconnect provided networking. Thunderbird's peak computing power of more than 64 teraFLOP/s made it number 5 on the November 2005 Top500 list.[40]

Compromises

Scalable and cost effective as it is, the Linux cluster style of computing is not necessarily the best architecture for the simulation of nuclear weapons. Serious compromises have to be made in the performance of various components, including the memory-to-processor connection and the interconnection network. Such compromises directly impact the computers' ability to apply their massive computational power to simulation problems. On the other hand, the beauty of this architecture is that it leverages the desktop commodity PC market. Therefore, many of the fixed costs of developing and manufacturing components are spread across millions of units. Another thrifty advantage of the Linux cluster systems is that the Laboratories are able to specify exactly what components will make up the system.

The final costs of the Laboratories' Linux cluster systems are not publicly available. However, based on press releases from the Laboratories and the manufacturers it is safe to guess that each system cost somewhere in the low tens of millions. This appears to be a bargain when compared to the ASCI capability computing platforms (e.g., Red and Blue) whose costs ran into the low hundreds of millions. In fairness, however, it must be noted that most of the computer science necessary to make the clusters work was technology developed in the creation of the large capability machines—leveraging this know-how made it possible to develop the "bargains" of the clusters.

Payoffs

More was gained from the Linux cluster experience than just some of the world's most powerful computers. In the 2005 interview, Seager explained, "The lesson here is that, by building similar systems over and over, you get better at it. Perhaps 'one-offs' are not a good investment to meet capacity computing demand. There is just too much overhead cost associated with them."[41]

Ultimately, the Laboratories' investment in Linux cluster technology through ASCI and other programs served ASCI's mission, served the research needs of non-weapons scientists, and in fact served to move the computer industry toward more innovative growth. PathForward investments accelerated deployment and helped tear down technological barriers. Since LLNL's MCR broke into the top 10 on the Top500 list in 2002, computers with Linux cluster architecture have regularly appeared with the world's most powerful computers and are often built at a fraction of the price of the large specialty machines. Bo Ewald has served as executive vice president and chief operating officer of SGI, and as president and chief operating officer of Cray Research, Inc. In a 2005 interview, while chairman and CEO of Linux NetworX, he commented on ASCI's impact on his industry, "ASCI did help the world move towards a commodity based style of high-end computing which has been good for small companies like Linux NetworX. ASCI certainly helped customers understand that they can solve big problems with clusters of small computers."[42]

BlueGene/L – *Collaborative Innovation at its Best*

As ASCI matured, it became increasingly clear that traditional approaches to building high-performance computers would require inordinate amounts of space and power. While this was sustainable up to the 100 teraFLOP/s goal, going beyond that would require innovative approaches. BlueGene/L is one example of a different way to deliver massive amounts of computing power. It resulted from a partnership between LLNL and IBM that identified the technology opportunity and nurtured it into reality.

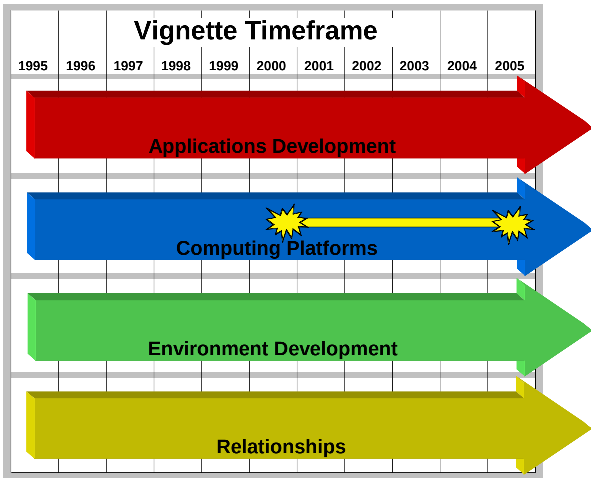

By the year 2000, ASCI had successfully procured five large capability computing platforms. These were ASCI Red, ASCI Blue Pacific, ASCI Blue Mountain, ASCI White, and the just-finalized contract to procure the ASCI Q platform from Compaq Computing, to be installed at LANL. As these procurements progressed, the machines were getting physically larger and were consuming greater amounts of electrical power. The Laboratories agreed that, although this was sustainable in the short term, ASCI needed to explore new ways to deliver computing power, especially as it considered procurement of the platform that would finally deliver the long-sought 100 teraFLOP/s.

For this reason, the Initiative started the Advanced Architectures program. This effort led to the IBM design for BlueGene/L, which was later selected to be manufactured at full scale. By the end of ASCI's first decade in 2005, BlueGene/L would provide a peak of 380 teraFLOP/s. (Editor's note: BlueGene/L was further expanded in 2007 to achieve 596 teraFLOP/s.)

Space and Power

Supercomputers, like ASCI Red or the ASCI Blue machines, are housed in rooms that nobody would suggest are comfortable. These rooms are very large and brightly lit. They also tend to be rather cold, drafty, and noisy. The computers themselves tend to be rather boring, generally consisting of hundreds of large refrigerator-sized racks of equipment with lots of cables running in and out. (Long gone are the days when computers offered the visual distraction of large spinning reels of magnetic tape.) One or two people might be seen moving around the room conducting maintenance on the system, which mostly consists of unplugging one part, then plugging in another. Systems are monitored from a control room that usually has windows looking out over the computer floor.

As the ASCI computers became more and more powerful, providing the physical space they required, along with the cooling, wiring, and maintenance accessibility, became an increasing challenge. Also, these computers consumed vast amounts of electricity. ASCI Red, the world's first teraFLOP/s system, delivered peak computing power of about 2 teraFLOP/s. The machine itself took up 148 square meters and consumed what was seen in 1996 as a staggering 1.6 megawatts of electricity.[43] That, it turns out, was nothing.

By 2000, new ASCI systems dwarfed Red in size and power consumption. In that year, Compaq Computers (which was later acquired by Hewlett Packard) was contracted to build and install the ASCI Q system at LANL. Q, which provided a total of 20 teraFLOP/s of peak computing power, was huge. At 1,266 square meters it was almost 10 times as large as ASCI Red and consumed 2.5 megawatts of electricity.[44]

DP and the Laboratories were concerned by the ever-growing physical size and energy appetite of the ASCI platforms. But even delivering an impressive 20 teraFLOP/s, Q was only one-fifth of the way to achieving the 100 teraFLOP/s necessary to begin simulation at the required scale for SBSS. Much more computing power was still required for ASCI's simulation needs. A crucial question was whether platforms could be built that would deliver even more computer cycles in smaller systems while consuming less energy. The design and eventual procurement of the BlueGene/L system proved they could.

Setting a Direction

Paul Messina came to ASCI from the CalTech in 1998. When Gil Weigand was promoted to DP Deputy Assistant Secretary for Research and Development, Messina took over as DP lead for ASCI. (In 2000, Weigand left the Department of Energy for new challenges in the private sector.) One of Messina's goals was to take a more aggressive role in defining the architectures for future computing platforms. Jose Munoz was at DP with Messina and was in charge of the Simulation and Computer Science program element. In 2005, he reflected on Messina's priorities, "They [ASCI] had been buying systems that were directed by the vendors for delivery in the near term. Messina wanted a medium-term system developed, which is where the idea for BlueGene/L came from."[45] To those ends, Messina created a new sub-program element, Advanced Architectures, which was led by Munoz. The element's scope was described in ASCI's 2001 Program Plan:

> In examining the prospects for future high-end systems, we have recognized that accelerating research into advanced computing architectures is important and could pay off by providing substantial leverage to future ASCI platforms. . . . The goal is to explore architectural alternatives that are not constrained by today's market forces.[46]

Advanced Architectures

The Advanced Architecture sub-program element was first funded in fiscal year 2001. But where were those funds to be invested? In 2000, David Nowak and Mark

Seager, part of the ASCI leadership team at LLNL, learned of a research program at the IBM Thomas J. Watson Research Center which seemed perfect for goals of the Program Plan.

According to Seager, he and Nowak were visiting the research center to look at technologies related to computing platform interconnection networks. In that context, one of the IBM researchers started to describe a system they were designing as part of the CHAMPS project, a computer specifically designed for Columbia University to make only Quantum Chromo Dynamic (QCD) calculations. Seager's reaction was that, while the approach to the interconnection was interesting, ASCI would actually be interested in the entire computer.[47]

The CHAMPS architecture was based on the idea that it was not necessary for any particular component of the system (e.g., processors, memory, interconnection network) to push performance boundaries. Rather, the aggregate performance of huge numbers of components would deliver computing power to the scientific simulations. This would allow backing off of the performance requirements for individual components; designers would have the opportunity to make components that were much smaller and used much less energy, but would entail using tens or even hundreds of thousands of processors to reach the 100 teraFLOP/s goal.

LLNL advocated researching the BlueGene/L architecture, but the other Laboratories and DP had legitimate concerns about its viability and usability. Just the idea of programming tens of thousands of processors seemed problematic. Nobody had programmed more than a few thousand, and it was feared that the communications, timing, and load-balancing problems would not scale, or would be extremely difficult to manage, or even insoluble. Due to such concerns, an external review panel was organized, which met in Berkeley, California, on February 28, and March 1, 2001. The panel included chair Michael Levine, Robert Borchers, Norman Christ, Candace Culhane, William Dally, and Robert F. Lucas. All were prominent computer scientists, drawn from a variety of government agencies, universities, and non-DP National Laboratories.

They examined the overall approach to the architecture and its potential impact on the use of the simulation applications to enable scientific discovery. In their report, the panel "strongly support[ed] proceeding with the R&D phase of this project." They did offer this cautionary advice: "The program should be structured to optimize the time to working applications rather than to minimize the time to having a full-scale system on the floor."[48]

Research and Development Contract

Based on this review, LLNL and IBM entered into a research and development contract on November 9, 2001, as part of the Advanced Architectures sub-program. At the contract's announcement, LLNL's Nowak commented on its significance:

> This represents a new thrust, very different from the approach taken by the main line of ASCI machines. Up until now, ASCI supercomputers have been designed to address the entire spectrum of numerical simulations required of the stockpile stewardship effort. This new BlueGene/L innovation can address an important subset of those

computational problems, those that can be easily divided to run on many tens of thousands of processors.[49]

LLNL's Lynn Kissel became program development leader. He later characterized the research and development contract approach as quite good since it was not tied to a full-scale platform delivery contract. This allowed the Watson Research Center computer scientists to explore the possibilities of the BlueGene/L architectural approach.[50]

"Novel Concepts"

Meanwhile, on April 22, 2002, the Initiative issued the Request for Proposals (RFP) for ASC Purple, a 100 teraFLOP/s computing platform to be installed at LLNL. ASCI had long ago set the end of 2004 as the date they would require a 100 teraFLOP/s system. The RFP incorporated the innovations used for the other ASCI computing platform procurements. It laid out a set of minimum requirements while giving the computer companies the flexibility to expand and enhance the performance of the system. The RFP also provided a chance for the companies to propose options for other novel computer architecture concepts. IBM took that opportunity to propose, as part of their bid to build Purple, an option to also build a large-scale BlueGene/L system.

On November 19, 2002, after long and complicated negotiations, Secretary of Energy Spencer Abraham announced a $290 million contract between LLNL and IBM for not only the 100 teraFLOP/s ASC Purple system, but also for a 360 teraFLOP/s BlueGene/L system.[51] This announcement, made at SC|02 conference in Baltimore, should have stunned the high-performance computing community. Instead, it was overshadowed by talk of the newest "most powerful computer in the world," the Japanese Earth Simulator.

Japanese Earth Simulator

The Earth Simulator was built by the Japanese computer manufacturer NEC. It provided a peak computing power of 40 teraFLOP/s. The machine was conceived in the late 1990s and was built in part because NEC received large quantities of funds for research and development from the Japanese government. As its name suggests, the computer was used to simulate phenomena on the global scale, such as earthquakes, weather, and climate. Earth Simulator employed SMP architecture, similar to that used by the Laboratories prior to the beginning of the ASCI

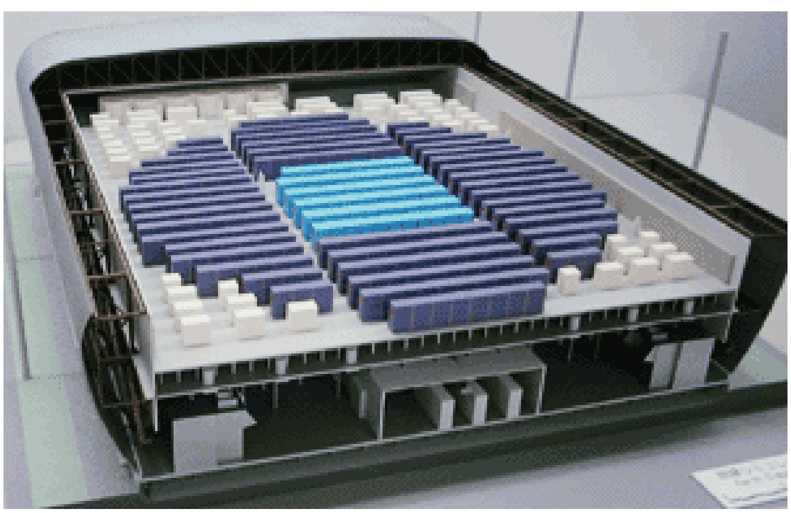

Figure 4-9. A model of the Japanese Earth Simulator and the building that houses it.

program. The Earth Simulator, however, used the SMP on a much larger scale than had been done elsewhere. This architecture provides a very fast connection between processor and memory and provides access to data without the need to communicate via messages among processors. Also, the interconnection network, known as a "crossbar," is extremely efficient as it provides a direct connection from any one processor to another. All these features make computer cycles more efficient when working on simulation problems as less time is spent waiting for data to show up. However, this efficiency comes at a price—in terms of money, in the large amounts of electricity consumed by the computer's high-performance components, and in the space required to house the system. In fact, one entire floor of Earth Simulator's building was dedicated just to cabling for the crossbar. Ultimately, Earth Simulator covered 3,200 square meters of floor space, consumed 10 megawatts of electricity, and reportedly cost $350 million.[52]

An Alternative Approach

The initial architecture design goals for BlueGene/L were that it would consume less than one megawatt of electricity and would take up only 230 square meters of floor space. Yet, it would eventually have 65,536 computing nodes. The computer uses an innovative arrangement of two computing processors on the same node. Both can be used to do computing, or one processor can be used to do computing while the other is used for communication tasks. Hence the system was designed to have 131,072 processors (now commonly referred to as "cores") that would deliver a staggering 360 teraFLOP/s of computing power.[53] (Because the configuration can be set to have only one processor computing while the other communicates, delivering half the possible computing cycles, the BlueGene/L performance is sometimes quoted with two numbers, e.g. it is sometimes described as a 360/180 teraFLOP/s machine. We will use only the single, larger number, representing the full computing configuration.)

The heart of the BlueGene/L system is an Application Specific Integrated Circuit (ASIC) and is a complete "system-on-a-chip" that includes two IBM PowerPC 440 processors, five interconnection network interfaces, and 8 megabytes of embedded fast access memory.[54] The system-on-a-chip is joined with nine memory chips to create a compute node providing 5.6 gigaFLOP/s of peak computing power and one half gigabyte of memory. Two of these nodes are then mounted onto a compute-card that is about 2 inches tall by 8 inches wide, about the size of a standard desktop computer memory card.

The computer is built to scale by assembling larger and larger collections of compute cards. Sixteen compute cards are put together in a node-card, which becomes one shelf in a rack of 64 node-cards. BlueGene/L also uses link-cards that connect the computer with devices for data input and output. The final 360 peak teraFLOP/s configuration consists of 64 cabinets.

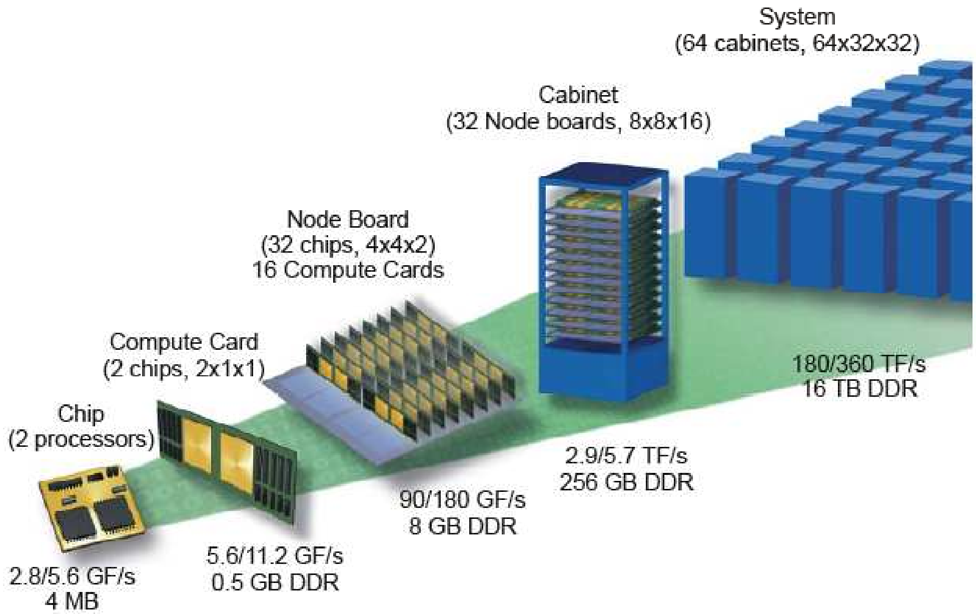

Figure 4-10. BlueGene/L architecture.

BlueGene/L incorporates application-specific network innovations. The processors are connected into networks using several different configurations, or topologies. First, the *3D torus network* provides excellent performance when a processor communicates with another processor that is "nearby" in the torus. Second, its *binary tree network* is used to quickly pass data to or from a large number of processors. Finally, a *barrier network* allows quick synchronization of the 131,072 processors.[55] In addition to the application-specific networks, BlueGene/L uses two other networks: a standard gigabit Ethernet used for data input and output, and one to provide system control and failure diagnostics.

One of the biggest potential problems in building a computer of this complexity is how to deal with the inevitable hardware and software failures. With so many parts, it is rather amazing that BlueGene/L runs at all. ASCI computers are among the most complicated electro-mechanical machines ever conceived. Early skeptics considered the millions of parts in the ASCI systems and predicted they would only run for a few minutes.

IBM and the LLNL BlueGene/L team took these challenges seriously and designed a number of features to boost reliability. This started with the system-on-a-chip approach to the computing node. By using a processor design drawing relatively little electrical power, they were able to increase reliability. Also, the single system-on-a-chip was able to replace up to 50 of the chips that would appear in a standard desktop computer. Since heat wreaks havoc with electrical circuits, IBM also paid close attention to cooling. BlueGene/L is air cooled but uses an unusual arrangement of airflow that provides especially efficient heat exchange. Providing this airflow path is why the computer appears to incline in one direction by about 15 degrees, unlike the standard upright refrigerator style cabinets of most high-end computers.

BlueGene/L's design team also took special care with its operating system. As with ASCI Red and Red Storm, a very simple operating system, known as a lightweight kernel, is used on the compute nodes. One benefit of a stripped down operating system is that it will not inadvertently issue conflicting commands that would cause a failure.[56]

Figure 4-11. BlueGene/L, showing the characteristic slanted cabinets designed to provide airflow for cooling.

Building the Platform

For three years IBM and LLNL scientists worked together daily to identify and resolve design issues so that BlueGene/L would operate as promised. This close collaboration exemplified ASCI's inclusive approach and deepened the productive partnership between IBM and LLNL that had begun with ASCI Blue Pacific.

IBM and the Laboratories held annual workshops that facilitated early appreciation of the complexities of system design. This exposure allowed the Laboratories' application programmers to get a head start on modifying their codes for the unique BlueGene/L architecture. Two such workshops were held in Reno and Lake Tahoe, in fall 2003 and in fall 2004. To assist in the transition to the new system, emulators were created, which would handle a program just as the final BlueGene/L architecture would.

On October 27, 2005, Ambassador Linton Brooks, head of the NNSA, officially dedicated both the ASC Purple and the BlueGene/L systems. In 2004, using only one-quarter of the final system, BlueGene/L had run the Linpack benchmark at 70.72 teraFLOP/s, catapulting it past NASA's Columbia system (51.8 teraFLOP/s) and the Japanese Earth Simulator (35.8 teraFLOP/s) to the head of the Top500 list. By the time of the dedication a year later, the full BlueGene/L system had achieved 280.6 teraFLOP/s on Linpack. At the dedication, Brooks highlighted the continuing importance of partnerships with computer manufacturers: "BlueGene/L points the way to the future and the computing power we will need to improve our ability to predict the behavior of the stockpile as it continues to age. These extraordinary efforts were made possible by a partnership with American industry that has reestablished American computing preeminence."[57]

Enabling Scientific Insight

Considerable effort was required to design new—or convert existing—applications to the novel BlueGene/L architecture. But the results were impressive from the moment the system was switched on, and BlueGene/L proved its value by enabling users to achieve scientific insight almost immediately. On November 17, 2005, at the annual SC|05 conference, Frederick Streitz of LLNL presented a paper entitled, "100+ TFlop Solidification Simulations on BlueGene/L." Like his ASCI predecessors who had used ASCI Red at SNL, Streitz and his co-authors won the prestigious Gordon Bell award. In the paper they described a very complex simulation of the solidification of molten tantalum and of the formation of the solid metal's grain structure.

Streitz and his team of seven other LLNL and IBM researchers simulated the actions of over 500 million atoms. The team included LLNL's James Glosli, Mehul Patel, Bor Chan, Robert Yates and Bronis de Supinski, and IBM's James Sexton and John Gunnels. Prior to Strietz' work, such simulations had been limited, by a lack of computing power, in the number of atoms that could be simulated. This had resulted in a simulated view of only one small part of a grain of solidified tantalum. Streitz's team, however, captured details of the solidification process that had never been observed before, *even in physical experiments*.

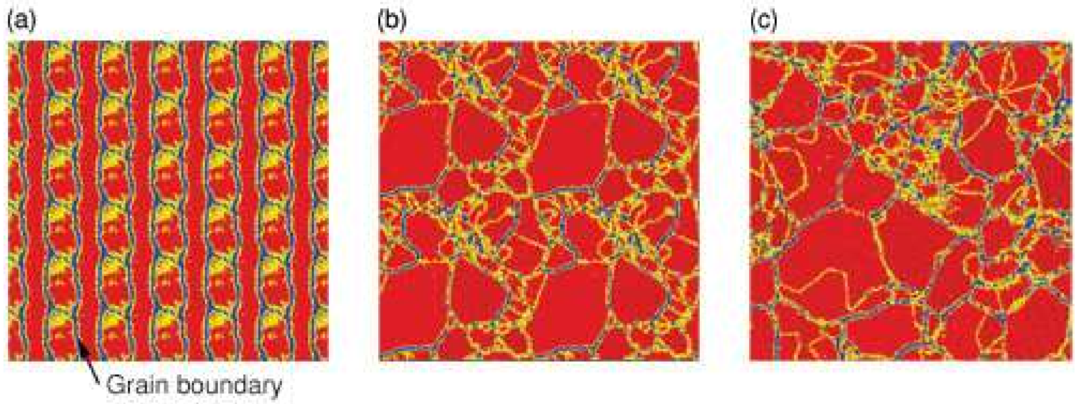

Figure 4-12. Atomic-level simulations of molten tantalum solidification and grain formation, using (a) 64,000 atoms, (b) 2 million atoms, (c) 16 million atoms.

Remarkably, even as Streitz and his team were running the tantalum solidification simulation, a variety of other simulations were running on BlueGene/L. A code called Miranda was being used to study the behavior of Rayleigh-Taylor instabilities, dislocation dynamics were being simulated with a code named ParaDiS, scientists running a code named SPaSM were learning how materials react to shock conditions, and researchers were using a code called FLASH, developed at the University of Chicago, for simulations to study nuclear flashes on the surface of neutron stars. Shortly after it was officially dedicated—even as parts of it were being installed—the BlueGene/L system was running at least eight different simulation applications.

There had always been questions about how useful the BlueGene/L architecture could be, but the proof was at hand. BlueGene/L could provide computer cycles to scientists and engineers with extraordinary efficiency. The tantalum solidification simulation used ran continuously, at more than 100 teraFLOP/s, for seven hours.[58]

At the BlueGene/L dedication, LLNL Director Michael Anastasio stated the importance of the computer: "The early success of the recent code runs on BlueGene/L represents important scientific achievements and a big step toward achieving the capabilities we need to succeed in our stockpile stewardship mission. BlueGene/L allows us to address computationally taxing stockpile science issues. And these code runs provide a glimpse at the exciting and important stockpile science data to come."[59]

During its early days, the Initiative set the 100 teraFLOP/s target, but nobody anticipated a system quite like BlueGene/L. Each evolutionary platform (ASCI Red, Blue Pacific, Blue Mountain) had enabled new scientific insights and stretched the boundaries of computer science. But as these systems grew, requiring ever more space and power, ASCI managers and researchers seized upon a new approach. Nicholas Donofrio of IBM said at the BlueGene/L dedication:

> The partnership between the National Nuclear Security Administration, Lawrence Livermore National Laboratory and IBM demonstrates the type of innovation that is possible when advanced science and computing power are applied to some of the most difficult challenges facing society. Blue Gene/L and ASC Purple are prime examples of collaborative innovation at its best. Together, we are pushing the boundaries of insight and invention to advance national security interests in ways never before possible.[60]

ASC Purple – *Fulfilling the Promise of Power*

In 1994 ASCI determined that at least 100 teraFLOP/s of were needed to fulfill the mission. At that time it was not clear to anyone if that goal could be met. A decade later, the 100 teraFLOP/s ASC Purple computing platform was dedicated at Lawrence Livermore. Not only did that computer deliver the needed power, but it was also done in a way so that important nuclear weapon applications and environments were immediately available to the Lab's scientists and engineers.

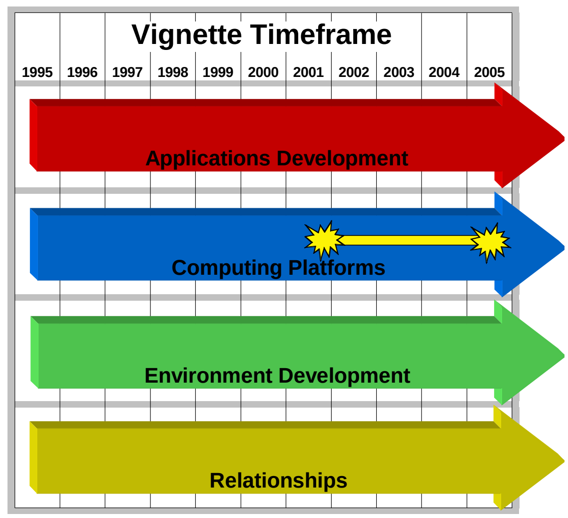

On October 27, 2005, Ambassador Linton Brooks, Administrator of the National Nuclear Security Administration, formally dedicated ASC Purple (the "I" in ASCI having been dropped in 2004) and BlueGene/L. At the time, ASC Purple represented the culmination of a lengthy and deliberate evolution of the ASCI's platforms, while BlueGene/L represented the birth of a new approach to HPC. Purple indeed delivered the 100 teraFLOP/s it was designed to achieve—and did so only a decade after Christensen and McMillan calculated the need for it. ASC Purple represented a 4,000-fold increase in the amount of computing power available to the Laboratory's nuclear weapons scientists and engineers.

The importance of ASC Purple is its direct lineage from the earlier ASCI Blue Pacific and White computing platforms; simulation applications developed for those previous systems migrated easily onto Purple. These included codes that simulated the operation of nuclear weapons under a variety of conditions and could address questions about the safety and performance of the weapons. The clustered SMP architecture approach of the ASCI Blue Pacific and Blue Mountain, White, and Q was particularly well suited to those sorts of applications.

Importance of 100 teraFLOP/s and Beyond

Numerical errors can cause quite a bit of confusion to scientists who interpret the simulations. Fundamentally, the results of computational simulations are just numbers. Granted, those numbers are often turned into very interesting and even beautiful visual images, but these pictures are just another way of presenting huge quantities of data. The data are generated from the equations used to represent things like the changes occurring to a physical system due to outside forces such as severe shocks.

One of the difficulties of using simulations for scientific discovery is that uncertainties in the computational results depend upon how, and how precisely, those numbers are calculated. In part, the uncertainties arise from the fact that the actual physical properties are not fully understood or not represented correctly in the equations. Until ASCI, uncertainties in the results of simulations had always been caused by a combination of this incomplete representation of the physics and a host of purely numerical phenomena (i.e., discretization error, round-off error). Without sufficiently powerful computers, it would never be clear if uncertainty in results came from deficiencies of the numerical solutions of the equations or from the incompletely represented physics. A major goal of the ASCI sequence of SMP platforms was to resolve the calculations to the point where the uncertainties in the results were caused solely by the inaccuracies in the representation of the physics, with no significant contribution from the numerics. It was thought that achieving 100 teraFLOP/s would resolve this question. ASC Purple was the first machine of this lineage with sufficient computing power to accomplish that goal.

Mitigating Risks

On the way to 100 teraFLOP/s, ASCI and the computer companies had envisioned and deployed systems that went far beyond the state of the art at the time of each proposal. Each milestone on the way to ASC Purple had been risky for both the companies and the Laboratories.

All of the platform procurements required innovative statements of work, creative risk management, intense negotiations, and close collaboration with the manufacturers. The procurements also incorporated a variety of novel features to mitigate risk. Design specifications, for example, were defined in a way that gave the manufacturers flexibility and yet ensured that the resulting systems would provide the necessary performance.

ASCI managed this by specifying a certain balance in the architecture. HPC platforms must achieve the best possible performance balance between its subsystems: the processors, the memory system, the interconnection network, and the input/output systems. If one component system lags behind the others, it can create a performance bottleneck where one system is swamped while the others sit idle. Part of achieving “balance” is often a matter of trading system performance for system cost.

The ASCI Blue RFP, released on February 12, 1996, was the first to specify particular performance ratios. It specified that for every FLOP/s of processor capability the architecture would include 0.5 to 1 byte of memory, 10 to 100 bytes of disk storage, 1 to 3 bytes-per-second of memory access, and 0.0125 to 0.125 bytes-per-second of interconnection network performance.[61] Five years later, the Purple RFP specified amazingly similar ratios. It stated that for every FLOP/s from the processor the platform would provide not less than 0.5 bytes of memory, 20 bytes of disk space, 1 bytes-per-second of memory access, and 0.1 bytes-per-second of aggregate interconnection network speed.[62] Specifying the performance values as ranges (or as lower limits) allowed manufacturers the design freedom to tailor performance among the system components while controlling overall system cost.

Another ASCI risk-mitigation tactic was to require that manufacturers provide early delivery of smaller systems. For example, the Purple RFP originally called for an Early Delivery Technology Vehicle (EDTV) by the fourth quarter of FY 2002 and for a system demonstrating the Purple technologies in the fourth quarter of FY 2004. The full system was due in the first quarter of 2005 (October 2004). This schedule assumed the Purple contract would be issued by July 2002. Because of the procurement's innovative complexities and the aggressive technology goals for Purple, the selection process and contract negotiations took longer than expected. Further complicating the process was the exciting development that IBM proposed to include BlueGene/L as part of the contract.

Negotiating the Purple Contract

The Livermore negotiation team was led by Mark Seager, who had also been principal technical negotiator for ASCI Blue Pacific and ASCI White and had overseen their eventual delivery, installation, and use. Negotiators were faced with two main issues. The first was the risk involved with both the Purple and BlueGene/L systems. In a 2005 interview, Seager described one way they handled the risk:

> "We wrote [into the contract] a clause whereby every so often we would review the progress and, if necessary, sit down together as partners and renegotiate the deliverables, schedule and/or payments. This was known as the 'partnership language.' This allowed us to deal with very formidable problems as they developed. In addition, the 'unwind language' in the contract allowed either party could bail out of the contract without incurring heavy penalties if these partnership negotiations failed."[63]

The other challenge was to keep the contract within the Initiative's budget. One point of leverage Seager had in this arena was LLNL's experience with the MCR Linux cluster system, which had been purchased with non-ASCI funds. MCR showed that the cluster style of architecture was viable at the multi-teraFLOP/s scale. While IBM had originally

Figure 4-13. ASC Purple at Livermore.

proposed an SMP system for their Purple Early Delivery Technology Vehicle platform, later negotiations led to them building what became the ASCI Linux Cluster. Seager and the Livermore team were able to use these systems to lower the overall cost of the contract. They also teamed with Lawrence Berkeley National Laboratory to convince IBM to produce a lower cost, 8-way shared memory node. Finally, Livermore chose not to use IBM's job-queuing system LoadLeveler, further lowering the price. Every economy counted, especially since budget cuts in FY 2003 made it unclear whether there would be enough funds for both the Purple and BlueGene/L systems. As it turned out, Seager's hard negotiations paid off; the savings realized during the negotiations made it possible to procure both systems.[64]

Delivering the System

During the summer of 2005 the 300-ton Purple computer was loaded onto trucks at IBM's Poughkeepsie, New York, manufacturing facility, where it had been fully assembled and tested. Once the trucks arrived at LLNL's newly constructed Terascale Simulation Facility the complicated process of installing and wiring began immediately.[65] Purple had to be installed with military precision, since a failure of any of its several hundred thousand parts could dramatically impact the system's overall reliability.

The 197 refrigerator-sized cabinets housed nearly 12,300 IBM Power5 processors and 50 terabytes of memory divided into 8-processor shared memory nodes controlled by IBM's AIX operating system. Secondary storage was provided by more than 8,000 hard disk drives that together could store more than 2 petabytes of data. The computer had more than 200 miles of cabling, much of it used in a high-speed interconnection network that was known as the Federated switch.[66]

A Path to Scientific Insights – Delivered!

Purple was the realization of ASCI's early "ultimate goal," a 100 teraFLOP/s computer with enough computing power to address problems at so fine a resolution that anomalies could definitively be attributed to phenomena of the physics rather than to numerical error. Purple's introduction increased confidence in simulation results; until Purple, scientists tended to assume that uncertainties in simulation results were caused by the computer's inability to fully solve mathematical equations. Since then, scientists may be just as apt to suspect a previously unknown phenomenon of physics.

In a 2005 article, Seager wrote, "Purple delivers the entry level computing power required for the full weapon system simulation capability the ASC program needs to fulfill its vital Stockpile Stewardship mission."[67] It is interesting that he characterized Purple, the culmination of ASCI's efforts in computational evolution, as "entry level," because in fact Purple represented not an end, but a beginning. Not surprisingly, even as Purple was validating its construction by revealing never-before observed details of real problems, forward-thinking scientists were already envisioning the giant leaps of knowledge that might be attainable with still more powerful tools, and they set about devising systems that could compute at the petaFLOP/s scale.

Truly, Purple and BlueGene/L thus represented a launch point far more than they embody an objective. Ambassador Brooks said it best at the October 27, 2005, dedication of the ASC Purple and BlueGene/L systems:

> The unprecedented computing power of these two supercomputers is more critical than ever to meet the time-urgent issues related to maintaining our nation's aging nuclear stockpile without testing. Purple represents the culmination of a successful decade-long effort to create a powerful new class of supercomputers. BlueGene/L points the way to the future and the computing power we will need to improve our ability to predict the behavior of the stockpile as it continues to age. These extraordinary efforts were made possible by a partnership with American industry that has reestablished American computing preeminence.[68]

---

[1] Summary of ASCI Workshop, 1.
[2] Christensen and McMillan, *Justification for the Need for 100 Tflop Computers*, 1.
[3] Gil Weigand, telephone interview by author, December 19, 2005.
[4] Ed Barsis, interview by author, Albuquerque, New Mexico, May 18, 2005.
[5] Weigand, interview.
[6] Top500.Org, November 1994 Top500 Supercomputer List, www.top500.org/list/1994/11/100
[7] Art Hale, e-mail to Alex R. Larzelere, April 4, 1995.
[8] Alex R. Larzelere, e-mail to Art Hale, April 20, 1995.
[9] "DOE Orders Teraflops Supercomputer for Nuclear Test Simulations," *HPCwire*, September 8, 1995.
[10] Mattson and Henry, "An Overview of the Intel TFLOPS Supercomputer," 2.
[11] Norris Parker Smyth, "Midsummer's Night's Nightmare: Teraflops REX," *HPCwire*, July 28, 1995.
[12] Elliot Swan, e-mail to Norris Parker Smyth, August 11, 1995.
[13] Mattson, 10.
[14] Ibid., 2.
[15] Art Hale, e-mail to Gil Weigand and Alex R. Larzelere, October 17, 1996.
[16] Gary Kent, e-mail to Art Hale and Alex R. Larzelere, November 26, 1996.
[17] Alex R. Larzelere, e-mail to ASCI Execs., December 16, 1996.
[18] Jan Rowell, "Full ASCI Red Goes On-line at Sandia," HPCwire, June 20, 1997.
[19] Attaway, et al., "Transient Solid Dynamics Simulations".
[20] U.S. Dept. of Energy Defense Programs, *Advanced Simulation and Computing Programming Plan* 2003, 30.
[21] Barsis, interview.
[22] U.S. Dept. of Energy Sandia National Laboratories, "Red Storm to be Assembled in New Mexico".
[23] John Morrison, telephone interview by author, June 23, 2005.
[24] Ibid
[25] *Summary of ASCI Workshop*, 1
[26] U.S. Dept. of Energy, "Statement of Work," in the *ASCI Blue RFP*, 11.
[27] U.S. Dept. of Energy, *ASCI Blue Mountain 3-TOps System: On the Road to 00 TeraOps*.
[28] U.S. Dept. of Energy Lawrence Livermore National Lab, *System Attributes: ASCI Blue-Pacific System*.
[29] Dongarra, Simon, and Meuer, "Top 500 List for November 1995".
[30] Dongarra, Simon, and Meuer, "Top 500 List for November 1999".
[31] Schwitters etal, *JASON Study – Requirements for ASCI,* 75.
[32] Dona Crawford telephone interview by author, May 6, 2005.
[33] Merky, *Beowulf History*.
[34] Deri et al, *ASCI PathForward Project*, 3.

[35] U.S. Dept. of Energy, *ASCI PathForward Description*, 5.
[36] Mark Seager, telephone interview by author, June 21, 2005.
[37] Dongarra, Simon, Meuer, and Top500.org, *Lightning Details*, Online: http://www.top500.org/system/6547
[38] Ibid.
[39] Danneskiold, "Los Alamos to get Lightning computer system from Linux Network".
[40] Hans Meuer and Top500.org, *Top 500 list for 11-2005*.
[41] Seager interview.
[42] Bo Ewald, telephone interview by author, July 28, 2005.
[43] Wu-Chun Feng, "Making a case for Efficient Supercomputing," 8.
[44] Parker, "Riding the Waves of Supercomputing," 9.
[45] Jose Munoz, telephone interview by the author, July 23, 2005.
[46] U.S. Dept. of Energy, *ASCI Program Plan 2001* December 1, 2000, 28.
[47] Mark Seager, telephone interview by the author, June 21, 2005.
[48] Levine et al., BlueLight Project Review, 2.
[49] IBM, *IBM and Department of Energy's NNSA Partner to Expand IBM's Blue Gene Research Project*.
[50] Lynn Kissel, telephone interview by the author, May 12, 2005.
[51] U.S. Dept. of Energy Lawrence Livermore National Laboratory, *IBM to Build World's Fastest Supercomputers*.
[52] Parker, 9.
[53] IBM, Lawrence Livermore National Laboratory, and the BlueGene/L Team, *An Overview of the BlueGene/L Supercomputer*, 2.
[54] Ann Parker, "Into the Wild Blue Yonder with BlueGene/L," 23.
[55] Parker, "Into the Wild Blue," 24.
[56] Ibid.
[57] Don Johnston, *World's fastest supercomputers dedicated at Livermore Lab*.
[58] Frederick H. Streitz, James N. Glosli, Mehul V. Patel, Bor Chan, Robert K. Yates, Bronis R. de Supinski, James Sexton, and John A. Gunnels, *100+ TFlop Solidification Simulations on BlueGene/L*.
[59] Johnston.
[60] Ibid.
[61] "Option Blue – RFP," (LLNL and LANL) 2 February 1996), 8.
[62] "Statement of Work," Advanced Simulations and Computing (ASCI) Purple (Livermore, California: Lawrence Livermore National Laboratories, 15 February 2002), 21.
[63] Mark Seager, interview by author, telephone, 21 June 2005.
[64] Ibid.
[65] Parker, 4.
[66] "IBM Demos ASC Purple Milestone Supercomputer," HPC Wire, 22 July 2005 http://news.taborcommunications.com/msgget.jsp?mid=436979 (16 February 2006).
[67] Ibid.
[68] Don Johnston, "World's Fastest Supercomputers Dedicated at Livermore Lab".

Chapter Five
Environments for Simulation Capabilities

Users and developers of the ASCI simulation capabilities operate in a very complex environment. They have had to deal with many different computer architectures, some with unique ways of organizing computations and data. A number of different technologies are used to write the applications that run on parallel computers, and many more are used to create the models to be simulated. et another set of tools is used to process, interpret, and visualize results. To add to the complexity, usually there is more than one technology that can provide a given function. Finally, because ASCI users often work in a classified environment, they must interact with a number of different security systems.

One of the reasons that the ASCI environment is so complex is that it needs to support, simultaneously, many different types of users. For instance, the scientists and engineers are ultimately trying to develop a scientific understanding of how nuclear weapons would operate under many different conditions. But those scientists and engineers depend on another group of users, who write new applications. This latter group consists of the developers, and they have to run their applications on the large platforms to ensure that changes to the programs will work properly on those platforms. In turn, developers depend on others to create system software, debuggers, and so forth.

An issue for ASCI has been to provide a robust, highly flexible *user environment* that enables all types of users to deal with the tremendous complexity of the applications and computing platforms. This challenge was made even more difficult as the Initiative progressed, as its applications and platforms were evolving rapidly. It was vitally important that the ASCI environment keep pace with the changes in the applications and platforms. Without a robust computing environment for the ASCI users, it is doubtful the advances for enabling scientific insight through simulations would have been possible. It is useful, then, to describe more fully exactly what is meant by "user environment," or more pertinently, to describe the elements of an ASCI environment.

ASCI User Environment

Although the ASCI environment supports a variety of users, not all users depend on all elements of the environment. For this reason, it might be useful to describe the ASCI environment as being organized into the following general elements:

Application Development – This element of the environment provides the tools and resources needed to assist users in developing and maintaining the very large and complex applications needed to conduct simulations. In addition to writing lines of code for an application, this includes functions such as debugging the code, understanding and tuning code performance, and implementing software quality practices.

Model Building – For every application, a mathematical model must be built to represent the physical systems to be simulated. The model consists of the geometry of the physical parts, along with information about the materials comprising the parts and the physical environment that will affect those parts. Very often, a large part of model building includes the creation of the finite element meshes that are used to enable a mathematical representation of the physical world.

Data Handling – The ASCI applications and computer platforms were designed to produce huge amounts of very detailed data about physical world simulations. The ASCI environment's data handling capabilities were required to scale well beyond anything that had previously been developed. The necessary capabilities included the file systems that could organize data and format it for future access, archives for long-term data storage, and tools for transmitting data from memory to archive and back.

Post-Processing and Visualization – As an application conducts a simulation, the data representing the evolution of the system are saved. Typically the simulation acts by "time-stepping," that is, by computing new values for the state of each variable, at a sequence of discrete points in time. As there are often many variables at each of many thousands (or millions) of spatial locations, and many thousands of time steps, this creates extremely large data sets to be saved for later examination. Depending on how the computer and the application were designed, it may be a huge job to concatenate the results data into a state where they can be analyzed. This element of the environment must provide tools to make assembling and organizing the data as easy as possible. Because of the enormity of the data sets (far beyond what anyone could analyze in a lifetime), tools and systems are needed to create visual representations of the data to help the user to understand the results.

Infrastructure – The user environment infrastructure provides basic functions that for many different purposes. These include such systems as high-performance networks and security technologies. Users at one Laboratory need to access and run applications on platforms at other Laboratories, a critical need in light of ASCI's Tri-Lab focus. And, since that much of ASCI's subject matter is classified, the work can not be done without appropriate security. As in so many other aspects, the infrastructure needs of the ASCI applications and platforms went well beyond any technologies existing at the start of the Initiative.

Using the ASCI Computational System

In some ways, the work done by teams building an environment is invisible, until something goes wrong. Since the beginning, the scale on which ASCI operated—the size of the platforms, the accelerated pace, the classified nature of the research—the challenge to provide an environment consisting of several evolving technologies has been tremendous. ASCI had to anticipate the needs of the user, as it could not afford to have problems crop up that would bring progress to a halt while "fixes" were devised.

Before describing specific technology problems confronting the ASCI environment teams, it is useful to examine how the platforms were used, in order to appreciate what the environment must support. The following is one example of a path that a user might take through the ASCI environment to use simulation to enable scientific insight.

A scientist trying to simulate a physical experiment must first study all the relevant information about the elements in the physical system. That information could include data from past experiments, theoretical calculations, and past simulation results. The geometry of the object undergoing simulation must be recorded in detail. In addition, the physical properties of the object, the materials of which it is made and the chemical or even atomic properties, must be understood and described concisely. Without this information, the simulation cannot portray the correct material properties or behavior. Finally, the scientist must consider the media in which the object of interest resides (or through which it travels) and the nature of any other objects or media with which it interacts.

Building a Model

With an understanding of the objects and their physical characteristics, the modeler next creates a mathematical model describing the interactions of all the elements in the system. Accomplishing this involves developing equations that describe the physical laws governing the objects and the physical environment in which they act. Next, a finite element mesh covering the simulated object(s) and the surrounding media is created. This step can be exceedingly complex, and can involve creating many specialized meshes for different parts of the simulated region that must be spliced together or overlapped, requiring many specialized equations that apply only at mesh boundaries, junctions, or overlaps. These meshes are generally so complex that highly sophisticated computational tools must be built just to create the meshes.

The meshes can then be populated with data about the physical properties of the various materials that make up the items. In some cases, the material data can be represented by formulas; in others, data are obtained from detailed tables of information derived from experiment (that is, from empirical data).

The model setup includes specifying a time-step for the simulation, that is, the amount of simulated time that elapses between times at which the simulation calculates the current locations and conditions of the materials in the model. It is crucial that to the time-step be specified sufficiently small to capture all of the physical events, some of which occur very quickly.

Developing the Application

The next step in the process is to select an application code to run the simulation. The initial study process will have included a review of existing simulation applications to determine if one were valid for the expected conditions. If no valid application exists (as has generally been the case throughout ASCI's history), and if no existing code can be suitably modified, then an application must be created (and the scientist would need a research and development team to create one). In this process, programmers create a structure for the application and build modules that execute the mathematical models of the relevant physics and materials of the problem. Those modules are tested on computing platforms to determine if they can be successfully executed. If not (a very common occurrence in the early development of a simulation) then debugging tools must be used to find the exact cause of the failure. Debugging on a massively parallel computer is particularly difficult because of the tremendous number possible of program elements, also known as tasks or threads, which are simultaneously executed.

Once the application is executing properly, developers must "tune" it to make it operate as efficiently as possible. Tuning includes studying application performance, and looking for bottlenecks that can be eliminated (for example, places in the simulation where thousands of processors sit idle while one processor computes a data value needed by the others to proceed). Finally, the application would undergo verification (to ensure it operates as developers intended), and validation (to ensure it is a correct reflection of the physical world). This is done in a variety of ways, often by comparing simulated results with data from actual experiments (where such data exist) or by comparing simulation with theoretical predictions, or by comparing with other applications designed to compute the same (or similar) physics.

Running Simulations

Once the application is developed, verified, and validated, the scientists are ready for production simulation. To operate the application they must submit a "job" to the computing platform. The user must specify a number of computational parameters—the number of processors to be used on the parallel computers, for example, and where the resulting data should be stored. Any "initial conditions" for the simulation, such as temperatures, pressures, velocities, elevations, etc., are input as part of the description of the job.

An important parameter that must be selected is how often the simulation should pause and store what are called "restart dumps." This is a complete description of the current state of all the variables, with sufficient detail about the recent history of the ongoing simulation to be used to restart an application if for some reason it does not run to completion. A restart dump enables the application to continue after a stoppage (a system crash or failure of critical hardware, or even an undiscovered logic error) without having to start the simulation again from the beginning. This is particularly important to the ASCI simulations, many of which run continuously for days, weeks, or even months. When a job stops unexpectedly, the computer reloads the most recent restart dump and proceeds from that time-step. These dumps also provide data about a simulation in progress and can be accessed to ensure the simulation is running as expected. The user could also specify that additional data dumps be made to preserve the intermediate state of the system during the simulation of a particularly important intermediate event.

Moving the Resulting Data

Once the simulation is complete, the user must gather the data for analysis and interpretation. Depending on how the simulation is written, this data can exist in several different locations. For temporary data storage, some computing platforms use hard disks connected directly to the processing units. On other systems, processing nodes do not have attached disks and, in all cases, long-term storage is located off the computing platform. Users must therefore make sure the results are transferred somewhere that allows interpretation. Of course, before it can be interpreted, the data needs a little work.

Some simulations are written so that the data are sent to separate files for each time-step and for each processor. In this situation, even data from a single time-step might include hundreds (even thousands) of files. Frequently, before a team can view the results, they must concatenate data files, a process that assembles the data into a small number of large files. While manipulating the data, users might also conduct operations on it, looking for certain aspects of the results that are pertinent to understanding what occurred in the simulation. After all these post-processing operations, the data giving results for the simulation are ready for interpretation.

Understanding the Results

An application computes, at each time step, data for each of many thousands of elements in the model. Were the data printed on paper in numeric form, users would have to sift through thousands of pages to understand even the simplest of simulations. For this reason, ever since the earliest use of computational simulation, results have been converted into some visual display. This could be as simple as a two-dimensional graph of, for example, temperature as a function of time, but as simulations have become more complex, ever more sophisticated methods to examine the resulting data have been developed.

One method is to plot the location of each element on the surface of an object undergoing simulation for a single time-step. From this information, a picture could be assembled showing the location of the object (and its shape at the specified time). The picture can be printed, displayed on a computer monitor, or even projected on a theater display. If this were done for several time-steps, the pictures generate a movie of object, showing its motion and evolving deformation. Innovative display technologies have been developed that allow users to see this sort of movie in stereoscopic vision, creating a "virtual reality" display of the simulation results.

When the simulation results have been displayed, analyzed, and interpreted, the data can be moved into long-term storage. Archiving is provided in the ASCI environment by the High-Performance Storage System (HPSS), which simplifies the process by providing a single point of entry to the archives and the data catalogs. HPSS takes care of moving the data onto storage devices, usually magnetic tapes, and keeps track of the whereabouts of the data. HPSS could also move (migrate) the data to different storage systems, depending on the access requirements.

The overall process described above summarizes the process of conducting computational simulation in the pursuit of scientific discovery. The process has been employed countless times, in many different fields and by many teams of scientists, engineers, and developers. While not invented by ASCI, per se, the Initiative was instrumental in providing the impetus, the astonishing platforms, the applications, and the environments that have led to simulation attaining widespread currency throughout the wider scientific community.

Simulation at a Distance

A remarkable feature of the ASCI user environment is that users do not necessarily have to be physically close to the computing platforms and storage systems they use. High-speed encrypted networks allow users at one National Laboratory to securely access the systems located at another. Scientists simulating a nuclear weapon (or anything else) could develop the application at LANL, then create the model with data from SNL, and use a platform at LLNL to run the simulation. Finally, using ASCI tools, the user could then create visualizations of the simulation results and display them in an immersive 3D theater located at LANL. Of critical importance is that all of this can be accomplished while maintaining the appropriate levels of security, even for highly classified applications, input data, and results.

The following vignettes provide a sampling of technological issues that confronted the builders of the ASCI user environment, ranging from debugging tools to file system formats, and from visualization systems to high-speed data networks. Organizing the work on these diverse technologies provided the Initiative with some of its greatest challenges.

In some ways, the user environment is not unlike the umpire of a baseball game—the better it does its job, the less it is noticed. It was far easier for ASCI managers to measure and evaluate applications or platforms the user environment, because the applications had specific, easily measured goals. “The world’s fastest computer” is readily measured, and is far more enticing, than “one heck of a debugger.”

Much of the work on the ASCI environment technologies took place more or less “in the shadows,” never resulting in a ribbon-cutting ceremony. Despite this and many technological and programmatic difficulties, enormous achievements were made on the user environment. Moreover, this area was critical to ASCI’s ultimate success. Without a high-quality environment in place to support users, the applications would have never have been successfully developed and the computing platforms would never have reached their full potential.

Parallel Programming Tools – *Enabling a New Approach to High-Performance Applications*

> A big part of the success, but also part of the challenge of ASCI, was the fact that often new technologies were under development at the same time they were needed by other elements of the Initiative. Perhaps this was best seen in the tools used by the ASCI applications developers. These tools not only required many technical innovations, but also forced ASCI to develop team relationships so that the needed technologies would be available in time to support the creation of the advanced applications.

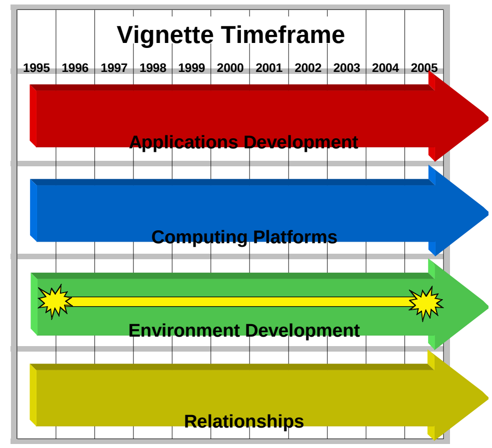

ASCI's planners recognized early that no single institution would be able to provide all of the technologies needed to achieve the objectives of the Initiative. The July 1994 ASCI Strategic Planning Document discussed who was expected to participate in technology developments:

> Individual ASCI projects will be executed by 'virtual' teams. Members of an individual team will come from all three DP laboratories. Where appropriate, membership will include other DOE laboratories, other government agencies, and industry. [1]

This approach worked especially well for technologies with relevance beyond just the simulation of the nuclear weapons. Fortunately for ASCI, the tools to create programs for massively parallel machines fall into this category. Developing the technologies of vital interest to the nuclear weapons community and to commercial and basic science researchers allowed ASCI to collaborate closely with industry and academia; this partnership ensured that new technical capabilities were available when they were needed.

The Parallel Programming Toolbox

All programmers use a variety of tools to create software applications, including the following tools, which proved essential for developers in the ASCI environment:

- **Debuggers** – used to identify, characterize, and correct errors in a program.

- **Compilers –** used to turn human-understandable programs into a computer-understandable form. Compiling programs also offered an opportunity to analyze the code and optimize it to execute more efficiently.
- **Performance analysis tools –** used to capture and display data about how a program executes on the parallel computers. These tools are used to help identify opportunities to make program execution more efficient.
- **Libraries –** used to provide an efficient way to call or use functions that are needed many times in a program.
- **Parallel algorithms –** the mathematical means to represent various functions, in this case designed to operate well on a parallel computer.

Such tools are vital to program efficiently and effectively on parallel computers, regardless of the application. LLNL's Mary Zosel helped drive development of these tools. In a 2005 interview, she discussed the problems of building the tools in ASCI's fast-paced environment:

> One of the greatest challenges we faced was the fact that our programming environment was a moving target. We had been working [before ASCI] to optimize the vector [style computer] programming environment, which we dropped. hen we moved onto the MPP environment of ASCI Red, and then the SMP environment of the ASCI Blue machines, White, and now ASC Purple. It seems that we were always running to catch up with the architecture choices.[2]

Fortunately, the technical challenge in developing advanced parallel programming tools generated broad interest in industry, academia, and the other National Laboratories. Commercial companies were attracted to potential profits from of research results, while academics saw the tools as "enabling technologies" that would facilitate the conduct of scientific enquiry. The 1996 ASCI Program Plan reflected these diverse interests in the strategy proposed to create these tools:

> The ASCI application development [programming] environment must provide the ability to rapidly develop complex applications that are efficient, scalable, portable, and maintainable. ASCI will work with academia and commercial computer companies to create the advanced development tools, methodologies, and standards needed to make this happen. [3]

But how was ASCI to build effective teams from academia, the other National Laboratories, and commercial companies? Although Laboratory engagements with academia or industry were common at the time, they had always been formed by individual DP Laboratories acting independently. When ASCI began there nobody had much experience executing a Tri-Lab program, and ASCI's Parallel Programming Tools element provided a perfect opportunity to refine the process of doing that.

Building on the SuperLab Experience

Fortunately, there was an example of early Tri-Lab computer science collaboration—the early 1990s SuperLab, led by Steve Berggren of DP Headquarters. SuperLab was an effort to coordinate high-performance computing research among the three weapons Laboratories using then-emerging Internet technologies.

Certain SuperLab technologies were demonstrated for Vic Reis on October 3, 1994. Afterward, LLNL's Hank Shay reported that, "Reis and the entire group were favorably impressed. Fundamentally, SuperLab offers the possibility of linking the resources of the three labs and to synergize the programs DOE is trying to develop (science-based stockpile stewardship, etc.)."[4] SuperLab's success was due in part to its novel management approach—it was run by a committee with representatives from Headquarters and from each of the three Laboratories. Building on the SuperLab model, ASCI was governed, from the onset, by an Executive Committee in the Tri-Lab fashion.

This management approach was employed repeatedly during the course of the Initiative, and a number of Tri-Lab committees were formed to target specific technology areas. Development of parallel programming tools, for example, fell under the Problem Solving Environment (PSE) program element, which later became part of Simulation and Computer Science (S&CS).

To execute a Tri-Lab project within PSE, ad hoc committees were formed to articulate the problem and to develop a funding plan to create solutions. The ad hoc committee recommendations were reviewed by the PSE committee and ultimately by the ASCI Executive committee. The Executive committee maintained a global view of ASCI, and worked to ensure balance between the various technology development activities. Once funding was identified by the Executive and PSE committees, the committee on parallel programming tools prepared a summary of the proposed projects that would become part of the Implementation Plans for a given fiscal year.

Options for Working Together

Having an effective management structure gave the Laboratories flexibility for engaging outside organizations in the development of tools technology. Several different approaches were taken, including:

- Doing the work at one, two, or all of the three DP weapons Laboratories
- Collaborating with other National Laboratories
- Working with universities, and often funding independent work at universities
- Contracting with commercial companies

Selecting from among these approaches, the PSE management and the specific technology committees needed to weigh costs, risks, and potential benefits carefully. Work at the both the DP and non-DP National Laboratories was expensive and there were issues with accessing available resources. Universities were less expensive but often had limited interest in pursuing ASCI-scale technologies and had constrained access to required tera-scale hardware systems. Contracting with commercial companies was certainly possible but only when research goals aligned with a company's business plans.

The Whitepaper – A Strategic Planning Tool

The process of identifying research goals and execution plans often began with a Tri-Lab white paper. These helped build a consensus about the technology need, delineate specific performance requirements, and elucidate ideas about how the need could be satisfied. White papers were used many times to identify needs for the ASCI environments program element.

An excellent example is the white paper, "ASCI Debugging Requirements," published March 28, 1998, and written by Jeff Brown of LANL, Mary Zosel, Rich Zwakenberg, and Mark Seager of LLNL, and Alan Williams of SNL. The introduction of the white paper stated the technical need:

> The Department of Energy Accelerated Strategic Computing Initiative (ASCI) program requires an advanced quality parallel debugging capacity with features and scalability that will keep pace with rapid improvements in compute platforms. Current debugging technology is not sufficient to meet ASCI programmatic needs for application code development.[5]

The paper covers specific issues with the types of applications that were under development, the expected rapid introduction of new computing platforms, and the challenges that programmers would face with platforms that were expected to scale up to 20,000 processors. Performance requirements for the debugger were also specified.

"ASCI Debugging Requirements" suggested that the best approach to creating the debugging technology ASCI needed was to build on an existing tool that could be scaled to the required performance levels. The paper also pointed out the value of dealing with a commercial company that would be able to provide ongoing support. The authors analyzed existing tools and found the product TotalView, the property of the Dolphin company, to be the best technical candidate to scale up to meet the requirements. The Tri-Lab committee's white paper concluded:

> The ideas described above represent our current thinking on debugger features of interest. New ideas will arise as time goes on and priorities may change, especially as we gain user experience with debugging ASCI-scale applications on Ultra-scale large systems. Core issues such as language support and scalability will remain top priority. An ASCI/Dolphin partnership that leverages the existing TotalView code base to add advanced debugging capabilities will greatly accelerate the deployment of much needed debugging technology required by the ASCI program.[6]

Based on this white paper, the three Laboratories and DP agreed to pursue a relationship with Dolphin. This led quickly to several questions: What form of agreement would ASCI enter into with Dolphin? Which part of ASCI would fund the activity? Which Laboratory would take the lead in executing the agreement?

TotalView

ASCI managers decided to pursue a standard commercial contract with Dolphin. (An alternative might have been to form a Cooperative Research and Development Agreement, or CRADA.) The contract was to be executed by the LLNL procurement office. The managers decided to fund this activity through PathForward. They also decreed that the contracts for computing platforms would be written specifying that the software systems would interface with the TotalView debugger.

Early in 1999, Dolphin split into two parts, Dolphin and Etnus. TotalView became the responsibility of Etnus, and LLNL entered into a research and development contract with Etnus to provide new features in the TotalView tools. A description of the project said:

> The scope of this effort has three thrust areas. The first is to extend the debugger within the TotalView framework to support application development that scales to tens of thousands of processors. Second, the debugging framework must provide greater functionality (e.g., aid the user in debugging complex C++, OpenMP, and Fortran 90 codes). Third, the debugging framework must be made more versatile, allowing the user to use performance analysis tools while using TotalView.[7]

All of the DP Laboratories were involved in guiding Etnus and reviewing its progress. Furthermore, each DP Laboratory would prototype new versions of TotalView, a process in which users test and apply evolving software to expose problems and find solutions to improve its functionality. The Laboratories used emerging versions of TotalView to develop ASCI applications. This provided ASCI scientists with experience running complex parallel applications on the most powerful computing platforms in the world. Without TotalView, or a tool with similar capabilities, it is unlikely that ASCI applications would have been ready in time to take advantage of the computer power as it was being delivered. In turn, valuable feedback from the ASCI scientists enabled Etnus to continually improve their software. The collaboration between the Laboratories and Etnus was extremely productive.

TotalView became a de facto HPC commercial standard. By 2005, many companies, among them IBM, SGI, Hewlett-Packard, Cray Inc., and Sun Microsystems, were shipping TotalView with their systems. This was a powerful demonstration of the wider value of ASCI-funded research to the industrial sector. The ASCI contract helped Etnus do well, which further ensured they would be there to provide support for TotalView. Etnus continued to add new features and functions to the tool that benefit a host of customers, including the National Laboratories.

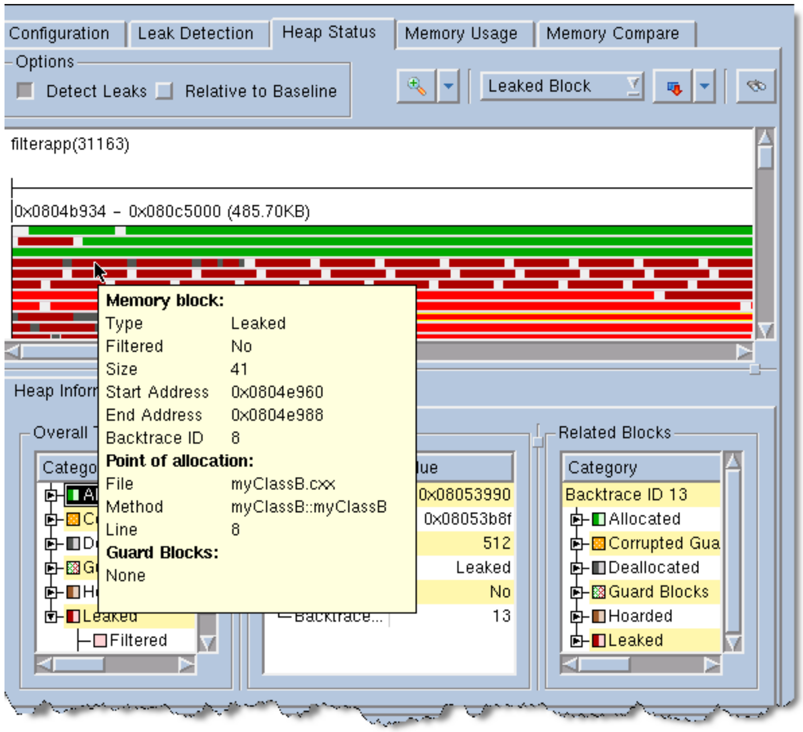

Figure 5-1. TotalView screen shot.

Other Collaborations

TotalView is but one example of how technology developments were made in the area of programming tools. Other examples include the continued improvements to the MPI libraries, which handle moving messages among the processors while programs run in parallel. The MPI work exploited the full range of partnering options available to the Laboratories. In addition to in-house development, universities and non-DP National Laboratories were engaged and commercial contracts were placed. The open-source software community was not ignored. For example, interfaces to MPI in other tools, such as TotalView, were developed. Their efforts resulted in numerous products used to build ASCI applications, including OpenMP, LA-MPI, Open MPI, Vampir, PAPI, Tau, and Valgrind. Speed and efficacy were critical. These programming tools had to be created concurrent with development of the applications needed to meet the ASCI mission.

Naturally, not all projects were as productive as were TotalView or Vampir. In some cases, technical goals proved to be too ambitious; in others, the needed technology did not lend itself to collaborative, inter-organization ventures. Productivity is, however, relative. While some projects did not produce tool technologies that were directly relevant to ASCI applications development, they always provided important technology or human-relations lessons that could be applied to other projects.

Demonstrated Success

By 2006, ASCI applications were running productively on the 100 teraFLOP/s ASC Purple platform. This was possible only because the programming tools that could support this scale of computer existed. Creating those tools was an exceptional accomplishment but was mostly overshadowed by the celebrity of the massive platforms.

Once Purple was running applications successfully, Mary Zosel reflected on the achievement: "One of the biggest accomplishments has to be the fact that people have successfully been able to develop codes for the ASCI computing platforms. There were predictions at the outset of the program that it would be impossible to program and debug code for the large numbers of parallel processors that were needed to achieve teraFLOP/s capabilities."[8]

Scalable Visualization – *Insight from Data*

Computational simulations used for scientific discovery generate huge quantities of data. Interpreting simulation results is facilitated by turning the data into visual images. As with the other ASCI capabilities, at the beginning of the Initiative the traditional methods for visualization were not adequate for the scope and complexity of the applications and platforms. This meant that ASCI had to develop new approaches and technologies to meet its mission needs.

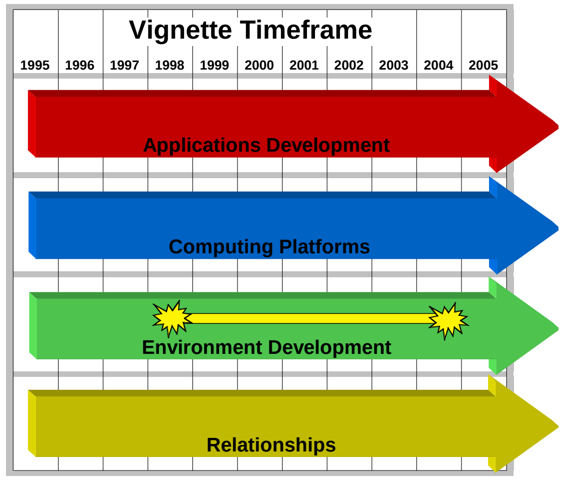

Simulations running on powerful computers create mountains of data representing the state of materials as each of the millions or billions of finite elements specify the state of physical objects at each time step. The problem for scientists is that they must dig through all of this data to find the critical information that is relevant to the question they are trying to address. Discovering the few bits of data that can unlock scientific insight is akin to finding a grain of rice on the beach.

Scientists sometimes have rough guesses about where and when the crucial events occur in the simulation, but those are often only rough guesses. Perhaps it was the initiation of a particular physical process, such as a turbulent flow; or it might be the change in state of a particular material. When, after careful searching, the information they want is found, scientists study it closely so they can understand the details of the simulated physical event.

Finding and correctly interpreting those bits of important data was a focus for the Initiative from the first program plan, published in 1996, which said:

> The [simulation] capabilities developed in the ASCI program will be used by weapon designers to help make crucial judgments concerning the safety, reliability, and performance of the weapons in the U.S. enduring stockpile. Making good judgments will depend on their ability to interpret and understand the data available to them. Given the massive amounts of data involved, they will depend on graphically oriented data comprehension applications. These highly flexible applications must

> (1) allow the designers to directly examine all aspects of the simulation results, (2) provide powerful analytical capabilities with customizable and extensible human-oriented graphical interfaces, and (3) handle the massive amounts of data that will be generated by the ASCI code-platform combination. Such tools do not exist today.[9]

Data and Visualization Corridors

In 1998, Paul Smith of DP and John van Rosendale of the National Science Foundation organized a series of three workshops on Data and Visualization Corridors for Large Scale Computation. In the final report on results of the workshops, they wrote:

> Across the government, mission agencies are charged with understanding scientific and engineering problems of unprecedented complexity. The DOE Accelerated Strategic Computing Initiative, for example, will soon be faced with the problem of understanding the enormous datasets created by [teraFLOP/s] simulations, while NASA already has a severe problem of coping with the flood of data captured by earth observation satellites. Unfortunately, scientific visualization algorithms, and high performance display hardware and software on which they depend, have not kept pace with the sheer size of emerging datasets, which threatens to overwhelm our ability to conduct research. Our capacity to manipulate and explore large data sets is growing only slowly, while human cognitive and visual perceptions are an absolutely fixed resource. Thus, there is a pressing need for new methods of handling truly massive datasets, of exploring and visualizing them, and of communicating them over geographic distances.[10]

The problem for users of computational simulations is that results data, in their simplest form, are just *bits*, that is, the binary “1s” and “0s” generated by computing platforms. In this form, it is impossible for anyone to understand simulation output. Even after the bits are converted into base-10 numbers more meaningful to humans, the output of an ASCI-scale simulation is vast; it is impossible for anyone to ferret out meaning from the raw data.

Before the start of the Initiative, simulation output was often small enough that scientists could interpret it in its numeric form. As late as the early 1990s, offices at the National Laboratories often overflowed with reams of fan-folded printouts, and computer monitors often glowed with numbers or graphs of simple one-, two-, and occasionally three-dimensional line plots. These approaches were not feasible for ASCI-scale simulations. The Initiative had to transform the approach users took to interpreting simulation results. This had to begin by changing the way scientists and engineers viewed simulations results, but the question was, “How?”

The human brain likes images. It is very well adapted to viewing them, interpreting them, and synthesizing understanding from them. A movie is an excellent example—what is a movie but a series of still pictures? A movie watched on television is merely a changing tableau of individual dots called *pixels*, or picture elements. The human brain receives those rapidly changing still pictures or shifting arrays of pixels and interpolates between individual pictures or dots to perceive a sense of motion, of imagery changing over time. A human's ability to understand tremendous amounts of data through pictures makes visualization a powerful way to interpret and communicate the information generated by computational simulations.

Handling and Visualizing Simulation Results

Turning data sets into visual images at the ASCI scale involved creating new technologies. There are three major components required to build a data visualization system:

- Data manipulation tools
- Rendering engines
- Display devices

Data manipulation is usually performed after a simulation has been completed and is referred to as "post-processing," although it is sometimes done concurrently with the physics calculation. Several operations are employed to extract and merge the data into a single file. Then scientists and engineers process the data to expose relevant information. One way of doing this is to extract "iso-surfaces" or "volumes" in the data. An iso-surface represents points in the data where values, such as the simulated material's temperature or the pressures experienced, are equal. An advantage of the surface/volume extraction process is that resulting data files are normally much smaller and thus easier to handle than the original simulation data. Once surface extraction is completed, the next step in the visualization process is rendering.

Rendering reads surface or volume data and creates multi-colored or gray-scale images, with highlights and shadows, creating the appearance of a solid (or sometimes the appearance of a translucent solid or volume). There are two forms of rendering. The first, "off-line" rendering, is done separately from image display. Hollywood uses off-line rendering to create animated motion pictures. These movies require a tremendous amount of computing power to render individual images (or "frames" of the movie), but that power can be spread out over many individual computers and over many weeks or months. The huge computational effort and lengthy time required are both worthwhile, since once a film is rendered and the resulting movie is completed it can be viewed as often as desired.

The other form of rendering is interactive and is done as the user explores the data. In this method the images are displayed immediately, as they are being rendered. This type of rendering is used to interactively explore scientific simulation data (this is also how the images must be generated to play computer video games or for "virtual reality" displays). As users choose to view the data from different perspectives, surfaces must be transformed into images extremely quickly, smoothly representing the viewer's passage through the data.

How the images are viewed depends on display systems. Display can be viewed on simple desktop computer monitors or on large theater screens. An extraordinarily powerful technique is to "immerse" the users *in* their data. This can be done in several ways. The simplest is to display a pair of stereoscopic views (one for each eye of the user) that create the effect of 3D images. Another method, known as a CAVE (Cave Automatic Virtual Environment), projects 3D images of the data on the walls, floor, and ceiling of a small room. Users in the room feel as though they are walking through the simulation results. Tools for this type of display were in their infancy when ASCI commenced (the first primitive CAVE was demonstrated at the SIGGRAPH conference in 1992).

The Scalability Issue

In their 1998 report, Smith and van Rosendale reported that ASCI computer scientists recognized in the 1990s that then-current data handling technologies did not provide the performance needed by ASCI. Visualization, at that time, was performed with technologies that would not scale well enough to handle the huge datasets expected from ASCI's applications and computing platforms, or from other data gathering systems such as satellites.

Soon after ASCI began, data visualization was handled primarily by systems made by Silicon Graphics, Inc. These SGI systems combined specialized processors, collectively known as "graphic pipes," with standard computing platforms. The computers handle data manipulation and surface extraction, after which the graphic pipes rendered the images and drove desktop or theater displays. The system worked well, easily handling the simulation results being generated at that time. All three DP Laboratories continued for many years to deploy large "multi-pipe" SGI systems, fondly called "Reality Monsters." Because their systems worked so well, in 1998 SGI held a virtual monopoly on high-end visualization tools. Unfortunately for users, the SGI systems were very expensive, ran only proprietary operating systems, and only a few graphic pipes could be used in parallel, resulting in bottlenecks when huge datasets were processed.

Having developed an understanding of the technology problems ASCI faced, the Executive Committee refined its approach in 1997, adding several new program elements. One of those elements was the Numerical Environment for Weapons Simulation (NEWS), which was later incorporated into the Visual Interactive Environment for Weapons Simulation (VIEWS). Working with academia, industry, and other National Laboratories, NEWS was to develop technologies needed for scalable simulation systems. "Scalable" meant that even as the quantity of data from the simulation grew incredibly large, the performance capability of every part of the visualization system would be able to grow to accommodate the data. Several basic requirements for such systems became apparent.

It was clear the system would have to handle datasets as large as several terabytes. Displays would have to be very large physically and have very high resolution to allow users close examination of details of the results. Users must be able to quickly navigate through a dataset. The use of motion in examining data was understood to be a powerful means for humans to detect and find interesting visual features. Finally, as ASCI expected users would not always be in the same location as the computer, the visualization system must support remote users. With a clear idea of what new capabilities were needed, ASCI then had to create them.

Visualization Software

Over the course of the Initiative, ASCI ran several projects to develop technologies to enable scientists to understand the simulation results. In 2001, ASCI published a Technology Prospectus that described the requirements, and the plan to meet them, for the Simulation and Computer Science (S&CS) program element (originally known as the Problem Solving Environment (PSE) element). The section on Data Management and Visualization included the following:

> The extremely aggressive ASCI visualization requirements, as with the aggressive ASCI computer requirements, have led to the development of a visualization architecture that depends on scalable hardware and software in order to attain the required performance. We are focused on research and development, on applying and combining commodity cluster technology, and tiling [using multiple projectors] of commodity display technology to reach our goals.[11]

New technology developed in the commercial marketplace, along with inventions spearheaded by the National Laboratories have significantly changed the world of data handling and visualization over the last decade. ASCI projects were fundamental in enabling these changes.

Three of the projects focused on software for the manipulation and visualization of large datasets. ASCI contracted with a company called CEI (Computational Engineering International) to develop Ensight Gold. This software enables users to manipulate ASCI-scale datasets, render visual images, and then interact with those images in a variety of ways. Ensight was particularly useful for simulations run remotely by LANL staff on the ASCI White computer located at LLNL. Ensight converted the data into many small files in Livermore and then rapidly moved them over secure networks to Los Alamos, where scientists used Ensight Gold and SGI graphic systems to render and display images from the files.

Two other ASCI software tools projects, VisIt and ParaView, also provided important data manipulation capabilities. Both used open-source software as building blocks, such as the Visualization Tool Kit (VTK, from a company named Kitware) to visualize and interact with data. The use of open-source software in these tools was important because it fostered greater engagement with the external community and provided a more robust tool-development effort. Both VisIt and ParaView provide important data manipulation capabilities for ASCI users, and in recognition of its importance, VisIt was awarded an R&D 100 award in 2005.

Leveraging PC Games

As important as these software projects were, the commercial market for PC video games sparked what is arguably the most significant advance in ASCI's visualization efforts. In a PC game, data are generated by the interaction of the players with the game, and graphic artists, working with the game designers, devise rules by which that data are converted to images. Although not generated by physics simulation applications running on high-performance computers, the PC game data undergo the same process described earlier to create visual images. The speed of image rendering and display is crucial to the success of a PC game and is accomplished in large part through the use of special-purpose high-speed processors known as "graphics cards."

ASCI recognized that the power and speed of graphics cards made them very attractive for processing the visualization data for the Initiative. One difficulty in taking advantage of commercial graphics cards was that, as powerful as they were, the cards and the computers in which they were installed could not handle the size of the ASCI datasets. To address this, the National Laboratories collaborated with Stanford University and the University of Virginia to develop a novel approach: use clusters of PCs, each employing a powerful graphics card, to render images that would be shown on very large displays. The partnership led to the Chromium software (another R&D 100 award winner) in which the rendering work is divided among hundreds of computers and graphics cards. After rendering, Chromium re-combines the resulting images for use in large scale display devices. This technology allowed the Laboratories to take advantage of lower cost, commercially available graphic cards, while providing the ability to handle large volumes of data and increased resolution.

To run parallel visualization software such as Ensight, VisIt, ParaView, and Chromium, and to make effective use of high-powered commodity graphics cards, it was necessary to construct a new kind of Linux-based, scalable visualization cluster. These clusters, similar to those being built by ASCI to run *capacity* weapons calculations, were constructed from PC microprocessors and emerging high-speed interconnects. The VIEWS program recognized early that this kind of approach could likely match or significantly outperform more expensive SGI Reality Monsters and, furthermore, were more easily expandable by adding scalable units.

As a modern follow-on to the small prototype PC clusters built at Stanford and other universities with help from VIEWS, all three DP Laboratories have deployed very large, production Linux-based visualization clusters with hundreds of nodes, some of which were large enough to have been included on past Top500 lists.

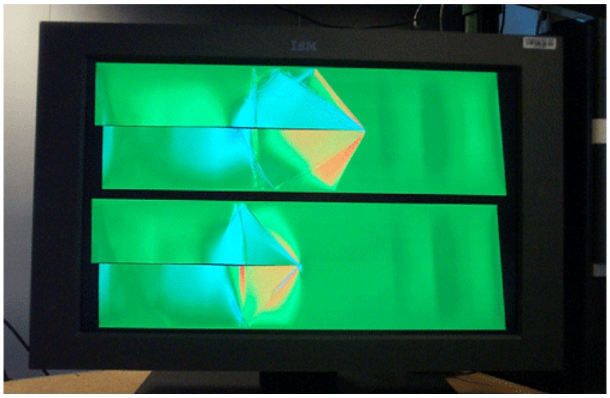

Figure 5-2. Bertha display.

Displays

The VIEWS program element made significant investments in very-high-resolution displays. Bertha, a 22-inch 92 million pixel LCD display built as the result of a small contract with IBM, had such fine resolution that it was impossible for a person to tell the

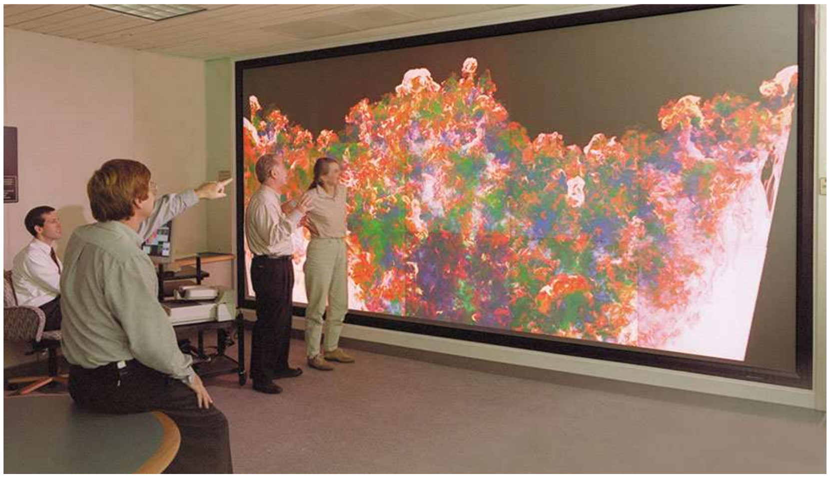

Figure 5-3. LLNL power wall.

difference between the display and the printed page. Bertha subsequently became a commercial success and was marketed by a number of display vendors.

On the other hand, a "power wall" was not only high resolution, but also quite large. The term was originally coined at the Laboratory for Computational Science and Engineering at the University of Minnesota, one of the earliest academic collaborators of the VIEWS program. Power walls used multiple video projectors, carefully aligned so that they effectively merged individually displayed images into one huge image. In 2000, Daniel R. Schikore, Richard A. Fischer, Randall Frank, Ross Gaunt, John Hobson, and Brad Whitlock of LLNL published a paper in IEEE (Institute of Electrical and Electronics Engineers) Computer Graphics and Applications that described in detail how these displays could be built and driven by appropriate parallel software. LLNL uses a number of large, high-resolution power wall displays in both its classified and unclassified facilities.

Figure 5-4. SNL VIEWS corridor.

SNL built a truly massive power wall at the Albuquerque site. This display measured 12 meters wide and 3 meters tall, and used 48 carefully aligned projectors to form an image consisting of over 60 million pixels.[12]

LANL built a CAVE display system that was also huge, with one large center wall about 4 meters tall. Smaller displays were installed on the ceiling, floor and on the sides of the large center screen. All was arranged so that a user would feel completely immersed in the simulation data that was projected with a resolution of over 43 million pixels. In the CAVE, users can feel as if they are standing right next to a huge molecule, for example, with the surrounding space extending to infinity in three dimensions.[13]

"Insight, Not Numbers"

The work done to enable post-processing and visualization of ASCI simulation data was essential to building comprehensive simulation capabilities. These tools were critical to building applications because generally the only way to verify and validate the codes is to run a simulation and then to examine and analyze the results. These visualization systems make it possible to examine the massive data sets efficiently, and the high resolution makes

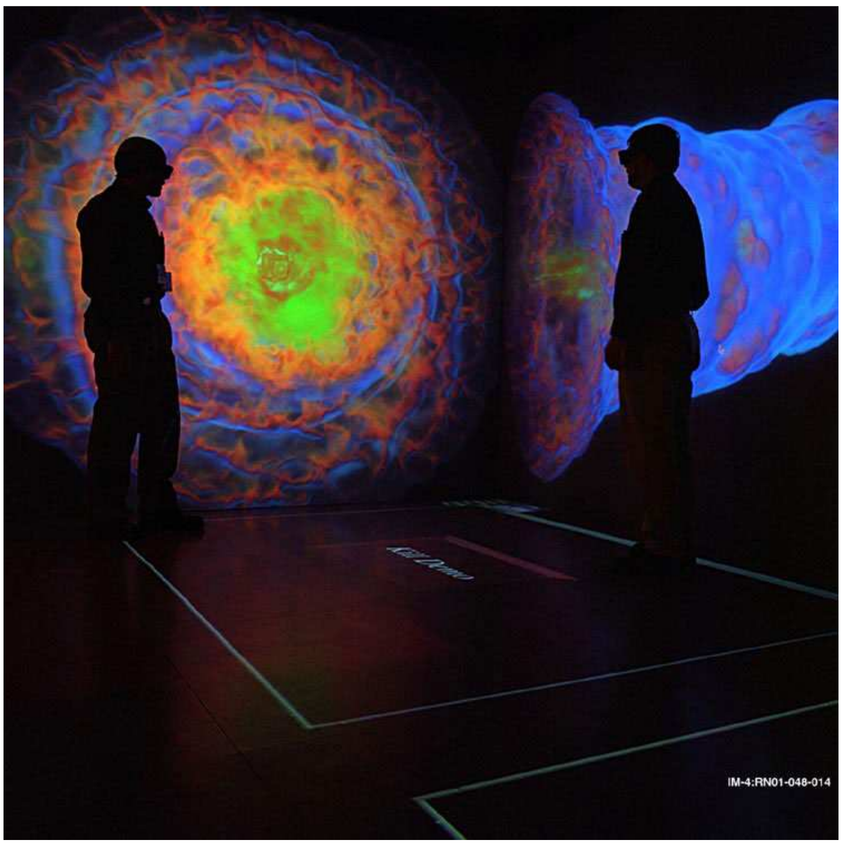

Figure 5-5. LANL CAVE.

it possible to examine the data effectively. The VIEWS program was driven in large part by the well-known quote of Bell Laboratories engineer R. W. Hamming: "The purpose of computing is insight, not numbers." ASCI took this environment challenge seriously.

In the foreword to their 1998 workshop report, Smith and van Rosendale quoted Weigand on the subject of handling large datasets: "ASCI is important to the Nation; visualization and data handling are critical to ASCI and everyone doing large-scale simulation."[14] Happily for ASCI and the United States, those capabilities were successfully delivered.

---

[1] Steve Berggren, Accelerated Strategic Computing Initiative: Strategic Planning Document (Department of Energy, 8 July 1994), 16.
[2] Mary Zosel, interview by author, 17 May 2005, Lawrence Livermore National Laboratory, California.
[3] Accelerated Strategic Computing Initiative: Program Plan (Lawrence Livermore National Laboratories, 1 September 1996), 17.
[4] Hank Shay, e-mail communication, October 4 1994.
[5] Jeff Brown, et al., ASCI Debugging Requirements (Los Alamos, Lawrence Livermore and Sandia National Laboratories, 28 March 1998), 1.
[6] Ibid., 8.
[7] Matt Wolfe, Laurie McGavran, and Neil Pundit, Scalable Programs Debugger (Etnus): Project Overview (Lawrence Livermore National Laboratory), Online at: http://www.llnl.gov/asc/pathforward_trilab/scalable_parallel_debugger.html
[8] Zosel.
[9] U.S. Department of Energy Defense Programs, ASCI Programming Plan 1996 (Department of Energy, September 1996), 17.
[10] Paul H. Smith and John van Rosendale, Data and Visualization Corridors: Report on the 1998 DVC Workshop Series (California Institute of Technology, September 1998), iv.
[11] Paul C. Messina, et. al., ASCI Technology Prospectus (Department of Energy, July 2001), 39.
[12] ASC Highlights 2003 (Sandia National Laboratories, March 2004), 52.
[13] Laura Monroe, Steve Stringer, and Jeff Brum, "Extreme Resolution Visualization Enabled New Discoveries." Scientific Computing (February 2005), 26-31.
[14] Smith and van Rosendale, vii.

Chapter Six
Partnering to Deliver Insight

From its inception, the Initiative was defined by teamwork. Its goals were well beyond the abilities of any one institution. ASCI harnessed the intellectual power of the National Laboratories, academia, and industry to enable advances that brought about a sea change in computational simulation. The leadership provided initially by Vic Reis and Gil Weigand and later sustained by Paul Messina, Bill Reed, Dimitri Kusnezov, and the ASCI Executive Committee was essential in establishing and maintaining productive partnerships. This may well be one of ASCI's greatest achievements.

Nuclear Weapons National Laboratories

Teamwork with and among Laboratories has a history going back as far as the establishment of the National Laboratory system. During World War II, the Manhattan Project had sought and harnessed world's best scientific minds, much as ASCI would do over four decades later. The need for the "best and brightest" continued after the shooting war gave way to the Cold War. In order to keep the nuclear weapons program on the cutting edge of research, the federal government engaged universities to manage and operate many of the nation's research National Laboratories. This might seem odd since universities are known for openly sharing research results while the government, with serious security concerns, decided to maintain control of the information necessary—and the infrastructure used—to design, manufacture, test, and deploy nuclear weapons.

Balancing the competing concerns of security and open accessing to university scientists resulted in what became known as Government Owned, Contractor Operated (GOCO) Laboratories. In the case of LANL and LLNL, the Maintenance and Operations (M&O) contractor was the University of California (UC). SNL was, for many years, operated by the Western Electric division of Bell Telephone, which later became American Telephone and Telegraph (AT&T). In 1993, the Martin Marietta Corporation (which merged in 1995 with the Lockheed Corporation to form Lockheed Martin) took over the M&O contract for Sandia.

The three DP Laboratories were tasked to review the scientific quality of each other's work. LLNL was established in 1952 to speed the design of the thermonuclear weapons. As time went on, both LANL and LLNL would conduct peer reviews of each other's work in highly secure environments. This system, as it required them to research similar questions, often resulted in competition for funding between the Laboratories. SNL focused on the engineering aspects of the weapon systems and, therefore, was not in direct competition, but participated in scientific reviews with LANL and LLNL.

DP was not the only customer for the research done at the weapons Laboratories. Work at the Laboratories was diverse, with activities sponsored by a number of programs, including DP, the Office of Science, and other elements at the Department of Energy. Each Laboratory also sought and conducted research sponsored by the Department of Defense and, later, the Department of Homeland Security. The different Laboratories worked for and in the national interest, but that interest was best served through a healthy competitive spirit between the Laboratories.

The competitive aspect of their relationships meant the Laboratories did not always collaborate well. But ASCI leaders recognized early on that the task at hand was too large for any one institution to handle, and that collaboration among the Laboratories, closer than ever before, was essential. ASCI managers intended to avoid the common "not-invented-here" syndrome, in which scientists and engineers at one Laboratory tend to eschew hardware, software, or tools devised outside their own Laboratory. ASCI wanted good ideas; to get them, ASCI wanted the people with the good ideas to talk to each other, both inside and outside the DP Laboratories.

Defense Programs Headquarters

DP was, and remains, a key component in inter-Laboratory relations. After World War II, the United States wanted to ensure that nuclear power and weapon technologies were in the hands of the civilian government and not the military. For that reason, the Atomic Energy Commission (AEC) was created to oversee all work at the weapons Laboratories and production plants. In 1975, after the energy crisis of the early 1970s, the AEC was replaced by ERDA, the Energy Research and Development Agency, which was focused on alternative energy research. In 1977, ERDA was consolidated with other energy-related agencies to create DOE.

DP became the DOE element providing oversight for the nuclear weapons program, and in 2000, DP became part of NNSA. Three primary elements comprise the DP organization: Headquarters (HQ), the Operations Offices, and the Field Offices. The HQ offices are located in Washington, D.C, and in Germantown, Maryland. The Operations Offices are based in major cities near the Laboratories. For the nuclear weapons Laboratories, the Operations Offices were originally located in Albuquerque, New Mexico and Oakland, California (although those offices were later merged into one, in Oakland). Finally, the Field Offices are co-located with the Laboratories.

At the time ASCI began, DP was charged with providing program direction and oversight for the Laboratories and the plants. This generally did not include providing technical direction, since the expectation was that each Laboratory was itself best qualified to set its own technical agenda. The Laboratories performed peer review amongst each other, to monitor technical performance of the Laboratories as a group. HQ set high-level policies and budgets for Laboratory activities, while the Operations Offices and Field Offices were involved in day-to-day management of their respective Laboratories.

ASCI would operate in a very different way. The demands of the urgently needed technologies rendered the traditional way of doing things insufficient. From the earliest planning, Initiative management was based on the idea that ASCI could succeed only if the Laboratories were to work together, and if HQ took a much more active role in setting technical direction and facilitating interactions between the three Laboratories, a philosophy encapsulated in the "One Program – Three Laboratories" Strategy.

Broader Research Community

ASCI also needed to establish and maintain relationships with research universities, relationships that served several purposes. Most obvious of these was that university researchers were the developers of many of the unclassified technologies useful to develop parallel applications and visualization. Additionally, universities produced people trained in computational science who would accept employment at the Laboratories, working on classified elements of the Initiative. Perhaps most importantly, the universities became a very public testing ground for some of the concepts necessary to advance simulation as a peer to theory and experiment. This solved an important problem for ASCI.

For hundreds of years, science has been a great, open, often contentious, conversation. Research is published not for pride, but for critique, to add to a collective body of knowledge. The Initiative, while pushing the frontier of high-powered computational simulation, essentially supported nuclear weapons; hence, much of the work to understand the science behind the weapons could not be published in open literature. Security issues precluded many of the accomplishments of the Laboratories from becoming part of the conversation of science or from benefiting from wide critique. ASCI management realized, however, that similar unclassified work, done with or at universities, *could* be openly published. The *approach* ASCI was taking to develop simulation capabilities could be published and critiqued via university research, even if work on specific ASCI projects remained secret.

Beyond the universities, other research entities would also prove useful to the ASCI enterprise. In particular, ASCI would establish partnerships with other, non-DP National Laboratories. Examples of profitable relationships included those ASCI had with Argonne National Laboratory in the areas of visualization and MPI, and also work on high-performance storage systems, conducted with both ORNL and the National Energy Supercomputing Center (NERSC) at Lawrence Berkeley National Laboratory. Initiative scientists also had important discussions with their counterparts in other government computational simulation programs, including programs at the National Security Agency, the National Science Foundation, and the Department of Defense’s High-Performance Computing Modernization Program.

Computer Industry

The relationships between the Laboratories and companies in the U.S. computer industry were also critical to the success of ASCI, as commercial computer companies played a vital role in providing much of the technology ASCI needed. While the Laboratories had always been first, or among the first, to acquire next-generation computers as soon as the vendors built them, the approach taken by ASCI was fundamentally different. At the time, the Laboratories were asking for computers that did not exist, and had not been designed. The Initiative required companies to prepare proposals for systems that did not appear in any price list, and in some cases, were merely notions. Computer companies would not only deliver the systems, they would be integral partners in the design of the systems. Contracting for the computing platforms had to be done in a whole new way.

Part of this approach was using contracts that were set up as research and development efforts. Contracts for these R&D efforts could not be structured as simple hardware procurement contracts; they had to be much more flexible. A true partnership was needed to deal with the unexpected results of R&D. The computer companies had to have the confidence to expose problems as they became apparent, and the Laboratories needed to be open to compromises. Often, the compromises turned out better than the original specifications.

The following vignettes demonstrate how ASCI innovatively cultivated these partnerships. These partnerships took place in a complex environment and required abundant cooperation and coordination between the people at the Laboratories and Defense Programs. The effort was worth it: the results of the collaborations often exceeded the original expectations.

Academic Alliances – *Harnessing the Power of Universities*

Universities have a unique relationship with the nuclear weapons Laboratories, providing many people and much fundamental research used for stockpile stewardship. From the start, ASCI recognized that universities would play an essential role in developing simulations with resolution that enables scientific discovery. The Initiative established a number of university relationships that have had a tremendous impact on the use of large-scale computational simulations.

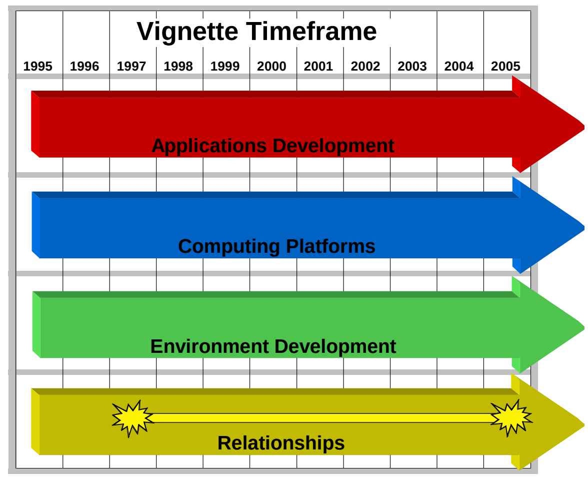

Relationships between the National Laboratories and universities are as old as the U.S. nuclear weapons program. During World War II, the University of California (UC) was selected to operate “Site Y,” the Manhattan Project’s secret laboratory at Los Alamos, and most of the scientists working there came from academic institutions. This was not surprising since, at the time, universities were the only source of theoretical and experimental scientists in nuclear physics. After the war many returned to those institutions, but maintained connections to the weapons program, and the universities continued to be an important intellectual resource for the weapons program. For more than half a century, UC continued to operate LANL and, beginning a few years later, LLNL as well. This business relationship allowed scientists to hold joint appointments with UC and a Laboratory and fostered collaboration between Laboratory and UC scientists.

University Contributions

ASCI leadership recognized the necessity for academia to contribute to the development of unprecedented simulation capabilities. Announcing the Initiative’s emphasis on cooperation with academia in 1996, Gil Weigand identified ASCI’s motivation to pursue these relationships, and presented areas where he believed that universities could help:

- ASCI simulation and computing problems are so hard that Laboratories cannot solve them alone.
- Develop a broad consensus that simulation is an appropriate means of ensuring confidence in the safety, performance and reliability of the stockpile.
- Help train the next generation of stockpile stewards.[1]

An important reason why ASCI needed universities is that many tools and techniques developed for classified weapons-related problems could be applied to unclassified problems by academics. Taking advantage of the open nature of university research, these tools and techniques could be validated in an open, peer-reviewed process, allowing wide implementations of advancements in simulation. Universities could demonstrate openly that simulation could function as a peer to theory and experiment. The 1996 ASCI Program Plan recognized the importance of universities:

> The shift to high-performance computing and science as the basis for confidence in the stockpile poses complex theoretical and practical problems in computer science and the physical sciences that are worthy of study by the best and most creative minds of the Nation. Engaging the efforts of individuals and groups at universities, other government agencies, and industry through Strategic Alliances and collaborations will be critical to the success of ASCI.[2]

Academic Alliance Levels

The implementation of this strategy was elegant and effective. ASCI established three levels of *Academic Alliances*. Level 1, the highest, created a small number of Alliance Centers at research universities. These Centers received enough funding, over a long enough period "to develop critical-mass efforts dedicated to long-term ASCI issues, such as high-confidence simulations."[3] The Level 1 Alliances were selected and managed jointly by all three Laboratories and DP.

Where the Level 1 Alliances addressed broad questions of developing simulation capabilities, the Level 2 Alliances focused on particular technology issues. Level 2 Alliances were funded to support individuals or small teams of researchers to address issues of interest to the Laboratories. Finally, Level 3 Alliances were narrowly focused on a single Laboratory's near-term activities and usually only involved a few researchers, normally an individual professor and some students.

ASCI also made time available on their terascale computing platforms to run unclassified university simulations. Access to this unprecedented level of computing capability for university research was a major asset for ASCI and also a huge inducement to the universities.

ASCI moved quickly to establish the Alliance Centers. Planning for the Centers began as soon as ASCI was funded, in FY 1996, and was led by Thuc Hoang and Merrell Patrick of HQ, with assistance from Dick Watson of LLNL, Dona Crawford of SNL, and Ann Hayes of LANL. In 1997, Bob Voigt of the College of William and Mary joined the group as part of the HQ team. Watson, in a 2005 interview, looked back on that early planning. He recalled, "I started having meetings with the other Laboratories and HQ to design the alliance program in 1996. Vic and Gil had envisioned the 3 levels at that time. I thought that was an excellent idea. First we set up the Level 1s, then the 2s. The level 3s were always a decentralized function. From the start, all three Laboratories were really committed to the Alliances."[4]

Proposals Process

One of the first decisions was that one Laboratory should administer contracts with the universities. This was intended to ensure that the contract administration would be identical for all of the Alliance Centers. After discussions among HQ and the Laboratories, LLNL assumed this task. However, all three Laboratories were responsible for technical guidance and oversight of the Centers, and this was accomplished through the Tri-Lab Alliance Strategy Team (AST). The AST was responsible for facilitating the competition to select the Level 1 Alliance Centers and the later Level 2 projects, as well as the implementation of the contracts. The team consisted of the people mentioned above plus Derrol Hammer of the LLNL Procurement Department. The competition process was widely credited as one of the most open and fair processes ever used to establish university research projects. In 2005, Watson commented, "The interesting thing was that we got a lot of feedback—even from those people that had been down-selected at the white paper level—that the process was very fair."

Rivalry for the Level 1 Alliance Centers was keen, which was not surprising the Centers were guaranteed ASCI funding of about $4 to $5 million per year for five years and possibly longer. The competition was announced in a November 11, 1996, Department of Energy press release in which Ernie Moniz of the White House Office of Science and Technology Policy said, "The DOE ASCI program is strategically leveraging the administration's scientific and technology instruments at U.S. universities to substantially advance our ability to numerically simulate scientific problems of national significance and unprecedented scale."[5]

The competition for the Level 1 Centers started officially with a bidders' meeting at a hotel at the Dallas-Fort Worth airport on December 5, and 6, 1996. The agenda called for ASCI to present program details and give the universities the opportunity to ask questions. The conference was divided into three main topics. First, Gil Weigand presented an overview of ASCI. Next, the competition process was described; and finally, there was a discussion of the technical areas of interest presented by representatives from the three Laboratories. More than 40 universities were represented at the meeting; naturally, very lively discussions ensued.

Hammer presented the competition process, which was, he explained, designed to be as burden-free for the universities as possible and yet provide the Laboratories with enough information to make selections. The process consisted of three major steps. The first was the preparation of “pre-proposal” white papers. These brief papers (approximately 10 pages) were to describe the technical scope proposed for a university center, how it would be organized, and its expected cost. HQ and the Laboratories, along with outside academic, government, and industry representatives, would review the pre-proposals and suggest to the universities whether a full proposal was advised. Armed with that feedback, the universities would decide if they wanted to submit full proposals. At this second stage, proposals were much more extensive and allowed the universities to go into detail about how they would build a Center to address ASCI’s goals. Universities with the most promising submissions would then be visited by HQ and Laboratory evaluators, who would judge their suitability as ASCI Alliance Centers.

Selections

In the months following the December 1996 pre-proposal meeting, teams of scientists and engineers at the universities created about 80 pre-proposals. After the review of the white papers, approximately 40 final proposals were prepared and subjected to thorough peer review.

On July 31, 1997, Secretary of Energy Federico Peña announced the award of five Level 1 Centers to universities, saying in a press release, “President Clinton has challenged us to find a way to keep our nuclear stockpile safe, reliable and secure without nuclear testing. We’re going to meet his challenge through computer simulations that verify the safety, reliability and performance of our nuclear weapons stockpile. I believe these Alliances will produce a flood of new technologies and ideas that will improve the quality of our lives and boost our economy.”[6] In the same press release, Ernie Moniz of the White House described the expected impact of the Alliance Centers: “The ASCI academic alliance program will help accelerate the preeminence of American universities in large-scale simulation, a methodology of rapidly increasing importance and enormous promise for leading-edge science. Moreover, we anticipate a ripple effect, as the DOE ASCI corporate and university partnerships enable the American science and technology community to understand complex physical systems.”[7]

All five Level 1 Centers focused on developing and demonstrating large-scale, integrated, multiscience simulation capabilities. The goal was to create simulations in unclassified subjects but of a similar complexity to those being done at the Laboratories. The following Centers were selected:

- The Center for Integrated Turbulence Simulations in Propulsion Systems at Stanford University
- The Computational Facility for Simulating the Dynamic Response of Materials at the California Institute of Technology
- The Center for Astrophysical Thermonuclear Flashes at the University of Chicago
- The Center for Simulation of Accidental Fires and Explosions at the University of Utah

- The Center for Simulation of Advanced Rockets at the University of Illinois at Urbana/Champaign[8]

Competition for the Level 2 Alliance projects started later, towards the end of 1997, with a Request for Proposals. The Level 2 projects selection process would result, for a number of universities, in "a single award … expected to average $200K-$400K per year for three years of focused, single-issue investigations."[9] The Tri-Lab group who prepared the RFP identified the following technology areas of interest to ASCI:

- Data manipulation, visualization, and their integration to enable "end-to-end" solutions for managing, assimilating, and delivering terascale scientific data to desktops of designers, analysts, and code developers.
- Scalable parallel computational algorithms for teraFLOP/s systems (1000s of processors). Topics in computational mathematics, software, and algorithms in computational physics/engineering (for example, radiation diffusion/transport and mechanics) of interest to ASCI.
- Scalable parallel software tools for effective use of massively parallel terascale computing systems (e.g., 100 teraFLOP/s by 2004).
- Software tools and algorithms for achieving terascale performance through distance computing in the form of heterogeneous distributed computing systems with thousands of commodity SMP's (with 8-256 processors per node) and commodity high-speed interconnects (SAN, LAN, WAN).[10]

The RFP also suggested scientific topics of interest, including energetic materials, condensed matter and material physics, and computational physics and computational mechanics. After a streamlined proposal and selection process in FY 1999, creation began of 13 Level 2 projects which were expected to be executed over three years. Unfortunately, due to budget pressure, funding for the Level 2 projects ended early; some were eliminated in FY 2001, while some were extended to the end of FY 2002.

The Level 3 Alliance activities were funded by the individual laboratories, usually through subcontracts. Each Laboratory devised its own selection process and employed its own separate ASCI funding.

Building Multi-Disciplinary Centers

After the proposal and selection process, the hard work of actually establishing the Level 1 Centers at the universities began. There were three main hurdles to surmount. First, a structure to coordinate the three Laboratories' interactions with the new Centers had to be established. Second, the universities had to deal with internal organizational issues to ensure a multi-disciplinary approach targeted on ASCI's needs. Finally, the actual researchers were confronted with technically challenging research projects and the building of suitable simulations.

To facilitate and oversee execution of Alliance Centers, the Laboratories continued to use the AST, which reported to the ASCI Executive Committee and functioned like the other Strategy Teams. For direct university interactions, the AST created a Technology Support Team (TST) for each Center; the TSTs had representatives from each Laboratory, including one expert on the application of interest, and one computer science expert. The TST's primary goal was to connect Laboratory scientists with university scientists. Each spring, a TST would meet formally with each Center to facilitate relations and to provide technical input to the scientific work. However, the TSTs and their respective Centers actively engaged in informal dialogue throughout the year. Each fall, the Laboratories organized peer reviews of the work at each of the Centers. The review panels each consisted of seven people, with five coming from outside the National Laboratories. These reviews helped the Centers break down barriers to multi-disciplinary science.

Breaking down Barriers

The Alliance universities, like most scientific organizations, tended to be organized around intellectual disciplines into departments, such as Physics, Chemistry, or Computer Science. That these departments were highly independent and rarely collaborated with other departments is not surprising. The goal of any academic department is to conduct world class research in a particular domain. Biologists, for example, focus intently on the problems within their expertise and generally do not have the resources or motivation to venture deeply into realm of the Chemistry or Physics departments. This phenomenon is known as intellectual "stovepiping."

ASCI had to change the culture of stovepiping. If the Alliance Centers were to help create new simulation capabilities, it was critical that their computer science people reach out to their physics, chemistry, and engineering people, and vice versa. The peer reviews helped expose relevant issues and break down an organization-centric perspective at all five centers. In a 2005 interview, Watson explained that peer reviews "have been crucial to keeping the Centers working. There's a strong temptation at the universities to slip back into their normal, single-discipline culture, but the reviews pull them back."[11]

On April 26 and 27, 2005, the first Computational Engineering and Science Conference was held in Washington, D.C. At that conference, three of the five Level 1 Centers (Utah, Illinois Urbana-Champaign, and Stanford) presented their work. Though the subjects were very different—simulations of explosives, rocket engines, and gas turbines respectively—a common theme emerged. All of the presenters talked about how difficult it was to make progress until their universities adopted a true multi-disciplinary approach, collaborating across domains to create the simulations.

Building the Technologies

The Level 1 Centers faced an enormous technical challenge, that of figuring out *how* to develop the simulation capabilities. Two overarching and interrelated problems had to be solved. First, it was essential that software integration frameworks be developed to integrate all the sub-parts of the large simulations. Second, it was necessary to develop and apply sound software-engineering methodologies for maintaining and testing the emerging, large pieces of software. It was swiftly discovered that faculty and students alone would not be able to develop, test, and maintain handle these massive software frameworks and packages—dedicated full-time university scientific and computer science staff were required.

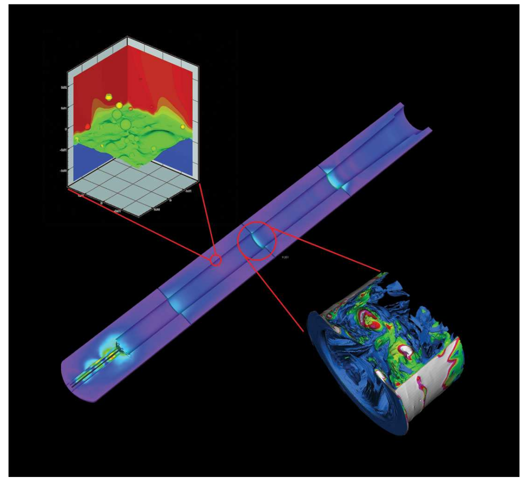

Figure 6-1. University of Illinois Urbana-Champaign simulation of solid rocket operation.

Figure 6-2. University of Utah simulation of explosive reactions in a fire.

Particular applications chosen by each Center were very different from each other. The University of Chicago explored the use of simulations to understand several issues associated with Type 1A supernovae. The University of Illinois Urbana-Champaign created simulations of full-scale solid rocket motor operation. The University of Utah simulated how explosives react in fires, and the California Institute of Technology simulated how energetic materials react to shockwaves. The Stanford University Alliance Center built simulations of the operation of full-scale gas turbine engines.

Despite the diversity of their research questions, the Centers faced a number of similar technical and scientific problems. Among the most challenging was that all of the simulations dealt with processes that occurred over widely ranging scales of space and time, a phenomenon also faced by the Laboratories in generating the simulations of nuclear weapons. There were many other scientific issues (e.g., turbulence, materials properties, and combustion) that were similar to those faced by Laboratory researchers. Because all the Centers were building end-to-end computational simulation systems, they also had to address integration and scalability issues, devising mathematical representations of physical phenomena that were accurate across the wide range of applications and scales. Furthermore, the Centers had to attend to parallel computing issues in order for their simulations to run on the large ASCI computing platforms. Just as did the Laboratories, the Alliance Centers used high-end visualizations to understand simulation results. Finally, the Alliance Centers also used experimental data to validate that their simulations were indeed true to the real world.

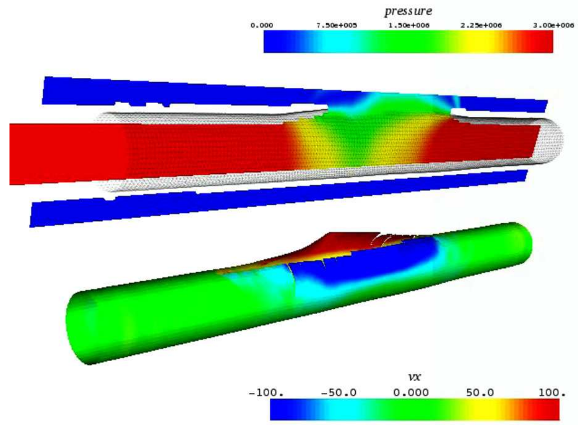

Figure 6-3. California Institute of Technology simulation of energetic material reaction to shockwaves.

New Tools for Science

ASCI's commitment to the Alliance Centers was well rewarded. Each of the five Level 1 Centers met ASCI's expectations. By 2005, all of the Centers had successfully developed comprehensive large-scale, integrated simulation capabilities for their respective areas of focus. The Centers helped bring about a change in the way the scientific community thinks about the use of simulation. Consider the image on the February 2002 cover of *Physics Today*. It is a beautiful image but not an artist's rendering. Rather, this visualization of how materials can build up on the surface of a star, ignite, and generate a runaway thermonuclear reaction was created using the University of Chicago's FLASH simulation code. This breakthrough in computational simulation also represented a breakthrough in astrophysics. Similar advances were made in different scientific areas at the other centers.

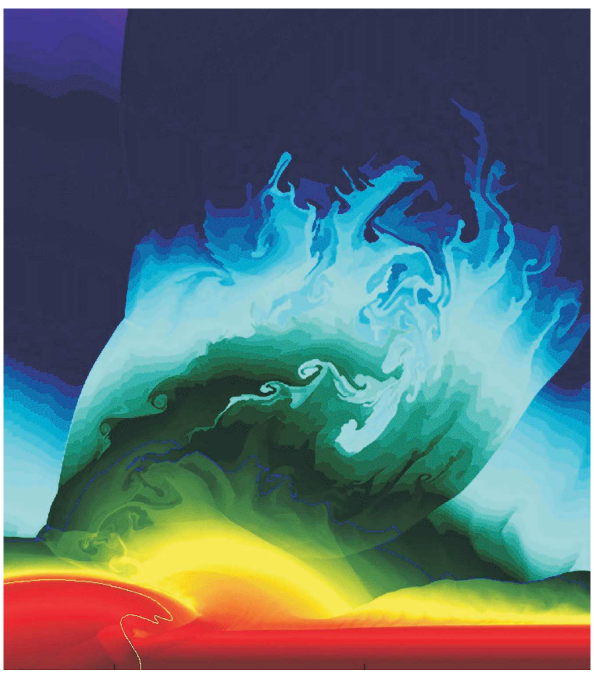

Figure 6-4. This University of Chicago FLASH simulation result appeared on the cover of the February, 2002 *Physics Today*.

Among the most significant achievements of the ASCI Alliance Centers was the fact that all five universities involved created programs, in some cases academic major degree programs, focused on computational simulation. Stanford established the Institute for Computational and Mathematical Engineering in 2004. In Stanford's technical progress report for that year, Parviz Moin, director of Stanford's ASC Center said, "In line with the multidisciplinary nature of CITS [Center for Integrated Turbulence Simulations] effort, ICME [Institute for Computational and Mathematical Engineering] is building bridges to all nine engineering departments at Stanford in addition to the Departments of Mathematics and Statistics."[12]

A good way to appreciate the accomplishments of the Alliances is to examine them in light of the original expectations. On December 5, 1996, Weigand presented the Alliances Strategy goals; they were to:

- establish and validate the practices of large-scale modeling, simulation, and computation as a viable scientific methodology in key scientific and engineering applications that support DOE science-based stockpile stewardship goals and objectives;
- accelerate advances in critical basic sciences, mathematics, and computer science areas, in computational science and engineering, in high-performance computing systems, and in problem solving environments that support long-term ASCI needs;
- establish technical coupling of Strategic Alliances efforts with ongoing ASCI projects in DOE laboratories;
- leverage other basic science, high-performance computing systems, and problem solving environments research in the academic community; and
- strengthen training and research in areas of interest to ASCI and SBSS and strengthen the ties among LLNL, LANL, SNL and Universities.[13]

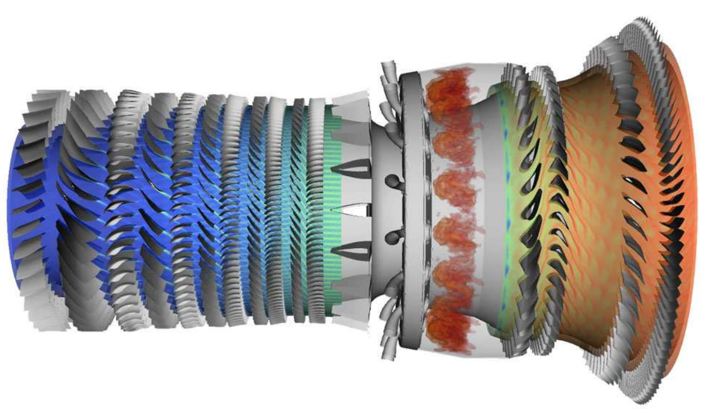

Figure 6-5. Stanford University simulation of turbine engine performance.

At ASCI's inception, the existing special relationship between the National Laboratories and universities provided a firm basis for these expectations. On the other hand, the ambitious strategy laid out by Weigand required new levels of intimacy, innovative organization, and facilitation, as well as a good deal of trust on both sides. Fortunately for ASCI, and thus the nation, the universities matched the hard work of the National Laboratories and DP's HQ staff to make these goals a reality, clearly establishing computational simulation as a full partner to theory and experiment for scientific investigation.

The relationships created in the Alliance program encouraged ASCI to extend and expand on the initial efforts. In 2005, the Initiative decided to hold another competition for Level 1 Centers, this one called the Predictive Science Academic Alliance Program (PSAAP). In an April 2006 press release, Dimitri Kusnevoz said, "The success of our centers at Caltech, Stanford, University of Chicago, University of Illinois, and University of Utah are examples of what such focused efforts can deliver. Through the PSAAP, we welcome the participation of our academic partners to help us develop the necessary, unclassified science and engineering applications and uncertainty quantification methodologies that will further establish viability of predictive science in multi-scale simulations."[14]

IBM and LLNL – *A Sense of Shared Mission*

At the beginning of ASCI, both IBM and LLNL were lagging in the world of high-performance computing. However, both organizations were laying the groundwork so that, in partnership, over the course of ASCI, each would “break out” by delivering and using massive amounts of computing power. LLNL and IBM developed a collaborative partnership that allowed them both to exceed all expectations.

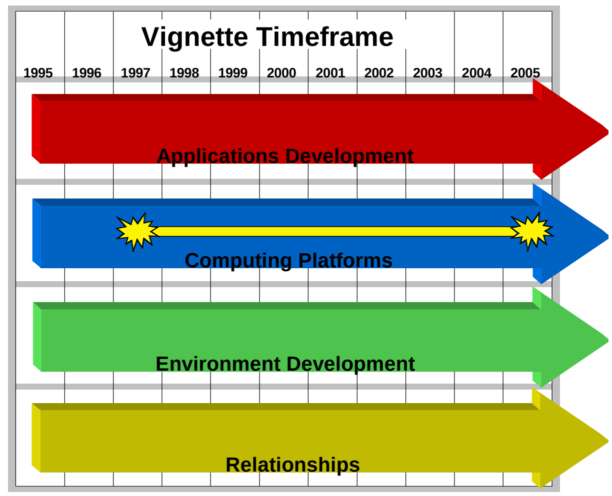

In 1954, LLNL procured a brand new IBM 701—a vacuum tube-based machine that was, at the time, one of the world’s most powerful. Its predecessor at LLNL was a Remington Rand Corporation UNIVAC computer, based on the earlier ENIAC design. The IBM 701 ran 12 times faster than the UNIVAC and was considered to be the company’s first commercially successful scientific computer. Its installation marked the beginning of a long and important relationship between LLNL and IBM that over the following decades included the design and delivery of many of the world’s most powerful computers.

The Top500 List Standings

ASCI sounded an alarm for both LLNL and IBM; by the Initiative’s start both institutions had fallen behind as the users and producers of leading edge, high-performance computer technology. One way of judging how institutions and companies stood in this regard was with the semi-annual Top500 list complied by Hans Meuer of the University of Mannheim, Erich Strohmaier and Horst Simon of NERSC/Lawrence Berkeley National Laboratory, and Jack Dongarra of the University of Tennessee.

The list is based on the amount of computational power a computer delivers to complete the Linpack benchmark, testing how many FLOP/s a computer uses to solve a dense system of linear equations. The Linpack benchmark and the Top500 have tracked the performance of high-performance computing systems since 1993.

The teraFLOP/s Decision

In January 1995, when Weigand announced at the Bishop's Lodge meeting that ASCI would buy the world's first teraFLOP/s computing platform, the latest Top500 list had just been released at the November 1994 Supercomputing Conference. A Fujitsu computer owned by the National Aerospace Laboratory of Japan claimed the top spot at 170 gigaFLOP/s. An Intel machine at SNL was second, coming in at 143.4 gigaFLOP/s. The third place machine, built by Thinking Machines, was at LANL and was rated at 59.7 gigaFLOP/s.

LLNL did not have a machine appear on the list until number 41, with a Cray Research supercomputer measured at13.7 gigaFLOP/s; LLNL's most capable computer was more than 10 times less powerful than the world's top machine. The situation for IBM was even worse. Its highest-rated computer appeared on the list at number 54, a 512-processor SP-2 system delivering only 12.1 gigaFLOP/s. Of the 500 computers on the list, only 32 wore the IBM nameplate.[15] Fortunately for both LLNL and IBM, in 1994 both institutions were already taking actions that would turn the situation around.

Laying a Foundation

Earlier in the decade, LLNL had commissioned the Massively Parallel Computing Initiative (MPCI) study led by Eugene Brooks, which had produced the famous "Attack of the Killer Micros" report recognizing that the fundamental nature of scientific computing was changing from systems with a few processors connected to a common memory bank to machines with hundreds and thousands of microprocessors, each with its own memory. The MPCI report gave LLNL confidence that the coming computer architectures could be applied to nuclear weapons simulations and helped the Laboratory develop strategies for how the new systems could be designed and used.

IBM was also making advances in parallel computing. During the late 1980s and early 1990s, architectures for parallel computers were emerging from companies like Intel, Thinking Machines, nCube, and Kendall Square. These companies pioneered the use of lower-power microprocessors grouped together on fast networks. Traditional computer

Figure 6-6. The IBM 701 at Livermore.

companies like IBM, Digital Equipment Corporation, and Cray Research were not so quick to embrace parallel architecture. It seemed that they were trying to preserve the mainframe approach to high-performance computing.

As parallel computers began to gain laboratory and industry notice, IBM was experiencing a multi-front crisis. The company was struggling to adapt to the shift to commodity microprocessors. In the 1980s, the IBM XT (and later the AT) became wildly popular, driving to a dominant position in the personal computer market. But IBM could not keep exclusive control of the technology and it rapidly became possible to buy clones of the IBM machines at much lower prices. This left the company reeling, and there was even some debate about whether IBM would survive. Fortunately, as the company entered the 1990s, they started to explore how they could use many of their high-end RS6000 workstations in parallel on a single problem. This approach resulted in the Scalable POWER parallel system of 1993, which became known as the SP1. Further development led to the more widely known SP2 computers, and these systems gave IBM parallel computing experience that would greatly benefit them (and ASCI) later in the decade.

ASCI Blue Pacific

The design of ASCI's first teraFLOP/s system (Red) was an outgrowth of SNL's earlier work with MPP computers, based on having just one or two microprocessors tied to each bank of memory. In March 1995, two months after Gil Weigand announced the machine would be at SNL, another approach emerged.

In March 1995, at a "teraFLOP/s summit" held in Livermore, LANL and LLNL proposed an alternative architecture, the clustered Shared Memory Processor (SMP) architecture (putting four or more processors into nodes that shared a single bank of memory), with a network connecting the nodes. ASCI decided to proceed with both the MPP and SMP architectures, and LLNL and LANL began the procurement process for the 3 teraFLOP/s ASCI Blue machine. In the end, two proposals were selected, one from SGI and the other from IBM. This left the problem of determining which Laboratory would get which computer.

In 2005, Bruce Tarter, LLNL's director during the ASCI Blue selections, observed that most at LLNL had wanted to get the SGI computer. The SGI proposal had been technically more interesting to the LLNL personnel, and envisioned a design offering shared memory across the entire computer. This could greatly simplify the creation of the simulation applications.

However, David Cooper, who had been recently hired as the Associated Director for Computation at LLNL, told Tarter that his Laboratory needed the IBM system. Cooper observed that LLNL trailed Finland in computing power and that it was crucial to ensure that LLNL's Blue machine would be successful. IBM, by far the more conservative company, had decades of solid engineering experience. After discussions between the laboratory directors, LANL chose SGI, and LLNL got the IBM. This put the dormant IBM-LLNL partnership officially back in action. In the following decade, this relationship delivered several platforms that propelled LLNL and IBM to the head of a number of Top500 lists.

Renewing the IBM Partnership

The IBM partnership with ASCI involved conversations at the very highest levels. During 1995, recognizing that the Initiative would ask computer companies to do things that may not be in their business plans, Vic Reis and Gil Weigand visited as many computer company CEOs as possible.

When Reis and Weigand met with IBM CEO Lou Gerstner, it was not clear the company would participate in ASCI. The market for high-end computing systems was shifting away from IBM and the traditional mainframes that formed its core business. The company was only just starting to build and deliver the new parallel computers, and demand was soft. Despite a great deal of activity in parallel computing in the late 1980s and early 1990s, there did not seem to be a viable market. A number of parallel computer companies, such as Thinking Machines and Kendall Square, had recently declared bankruptcy.

Reis, Gerstner, and their respective staffs met at IBM headquarters in Armonk, New York. Reis explained how the CTBT would impact the National Laboratories' ability to assess the safety, performance and reliability of the nuclear weapons, and described the stockpile stewardship program. He emphasized the role that computer companies would necessarily play in providing the platforms ASCI needed. The IBM executive listened intently.

Gerstner understood how crucial SBSS was to the nation and recognized the contribution IBM would have to make. He reminded Reis of IBM's long tradition of supporting efforts critical to the national interest and then assured him that IBM would be there when needed. According to Nick Donofrio, an IBM Executive Vice President, Gerstner asked him after the meeting, "Nick, can we do this?"[16] Donofrio reassured him, stating that he thought IBM was up to the challenge.

The Results

IBM's commitment to ASCI was a key event in the Initiative's history. The procurement of the 3 teraFLOP/s IBM system, later named ASCI Blue Pacific, was only the first of a long series of IBM contributions. Blue Pacific was followed by the 12 teraFLOP/s ASCI White computer, the 9.2 teraFLOP/s ASCI Linux Cluster, the 100 teraFLOP/s ASC Purple system, and the 360 teraFLOP/s BlueGene/L computer. Because ASCI used the R&D approach to contracting for these systems, IBM also produced early "initial delivery" versions as part of each contract. These early systems enabled the Laboratories to understand the performance characteristics of the eventual full-scale systems and to have software prepared for use immediately upon delivery of the hardware.

IBM not only delivered the large platforms, it was also involved in several PathForward and VIEWS technology development projects, including the final development of the Bertha Liquid Crystal Diode (LCD) computer display. This 22-inch diagonal display had a resolution of 3,840 by 2,400 pixels, and later became a successful commercial product called the IBM T220, along with its improved successor, the T221. Bertha was intended to display visualizations of simulation results and was so fine-grained that it was impossible to see individual pixels—even with a magnifying glass. IBM also worked on the PathForward development of the Colony interconnection network switch, which led to improvements in the communication efficiency of both the ASCI White and ASC Purple platforms.

Sharing Risks

A productive relationship must be based on trust and a sense of shared mission. Spurred by ASCI's ambitious goals, LLNL and IBM developed a tight partnership nurtured by (and essential to) their R&D projects. This sense of partnership started with the initial negotiation of contracts. Standard procurement practice for the Laboratories had always been that contract negotiations began once the competitive selection process was completed. Representatives from the Laboratories and the companies would work out the terms and conditions of how the work would be conducted, when it would be accomplished, and how much it would cost. For ASCI, contract negotiation was tricky because ASCI's needs entailed considerable risk for vendors. The computing platforms described in the procurement were often based on technologies and components that did not yet exist. This made it almost impossible for the contracts to have stable prices.

IBM and LLNL responded to these business uncertainties with an innovative approach. In a 2005 interview, Livermore's Mark Seager talked about the partnership strategy and the spirit that was written into the contract. Seager related that during the negotiations for ASCI Blue Pacific, LLNL and IBM committed to "renegotiate a mutually acceptable solution in the future, if IBM had problems delivering on something."[17] This understanding was turned into a formal clause that would appear in all of LLNL's ASCI contracts with IBM. For example, the following clause appeared in the sample subcontract contained in the ASC Purple Request for Proposals:

> The University and the Subcontractor agree that this Subcontract involves the development of cutting-edge technology under aggressive schedules. The University and the Subcontractor agree (i) that the Subcontractor shall use reasonable efforts to deliver in accordance with the requirements and schedules set forth in this Subcontract; and (ii) to reasonably consider limitations that may occur in meeting obligations under this Subcontract. If the Subcontractor is unable to meet its performance obligations, then the University and the Subcontractor hereby agree to negotiate the Statement of Work and the Subcontract price, if necessary, to reflect changes to the Subcontractor's performance obligations.[18]

Having this clause in place allowed both sides to sign contracts that might otherwise have seemed unacceptable. The acknowledgment of risk, and the commitment to a solution which respected both parties' needs, allowed delivery of systems that might otherwise never have been built.

Working the Partnership

With a signed contract in place, the hard work of R&D began. That both IBM and LLNL were able to approach this work as a partnership proved invaluable. The strategy facilitated open communication and allowed problems to surface earlier in the process. Technology issues were discussed fully as they arose and solutions were defined that were acceptable to both sides.

Over the course of the Initiative, the way LLNL operated with IBM was refined into a standard set of practices for building large-scale advanced computing systems. Mark Seager outlined these practices on a slide that has subsequently been used in many ASCI presentations. In it, Seager described a recursive process that began with fundamental "computer science research and developments" done at universities, industry, and the full complement of National Laboratories. This was followed by development of "innovative or evolutionary architecture ideas" that were then subjected to "rigorous reviews." Next, those ideas that passed reviews were incorporated into "research and development contracts." After a contract was agreed to, system design and production would commence in a spirit of partnership. The process then entered a cycle that used "flexible contracts with targets or requirements" followed by "milestone progress" and periodic reviews. The result was "full-scale delivery and integration" of the high-end computing system, which led back into fundamental computer science research and development, starting the process anew.[19] To make this process work, the partnership needed to be embraced at all levels within both organizations.

LLNL's Robin Goldstone knows the value of the partnership. She spent time with IBM developers in Poughkeepsie, working on the ASCI Blue Pacific system, and building face-to-face relationships. Later, when the inevitable technical problems arose, she not only knew whom to call, she knew the person she was calling. She wrote in 2005, "Having direct access to these individuals and having built a relationship with them over several years made these interactions extremely productive and led to faster turnaround and ultimately better solutions."[20] The

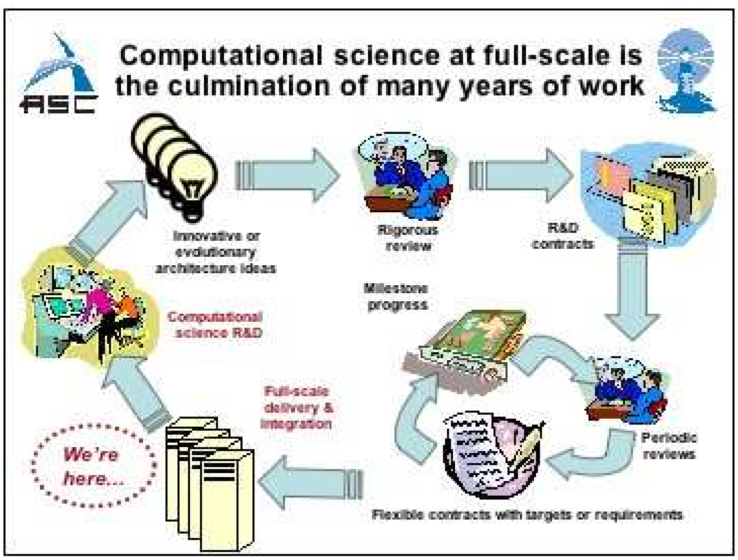

Figure 6-7. Practices for building large-scale advanced computing systems.

same well-established relationships between LLNL and IBM would be used again and again to solve problems and integrate solutions during development and deployment of the ASCI's White, Linux Cluster, Purple, and BlueGene/L systems.

The success of the LLNL-IBM partnership was a model that influenced all of ASCI. The three Laboratories jointly developed the Option Purple RFP for the 100 teraFLOP/s system. In it, they described important lessons they wanted to apply to the procurement that was open to all computer companies. The RFP authors eloquently recognized the value of partnerships:

> The most important lesson we have learned is that this effort, if it is to succeed, must truly be a "partnership" among all involved. While careful mutual planning on the part of the Laboratory and the Industry partner is essential to meeting requirements, unforeseen events and changes are likely. These events can only be successfully dealt with by a partnership that goes beyond an ordinary vendor-customer relationship. It must be one in which teaming, mutual respect, and an honest desire to achieve success is present on the part of everyone involved.[21]

Moving Up the List

Strong evidence of the partnership's efficacy is how quickly, and how completely, LLNL and IBM turned around their standings in the HPC world. Just under four years after the Bishop's Lodge meeting launched the era of teraFLOP/s computing, LLNL and IBM had risen from 41st and 54th, respectively, on the Top500 list, to numbers 8 and 6, each with versions of Blue Pacific. By June 2000, Blue Pacific stood second on the list behind ASCI Red, and six months later the IBM-LLNL partnership topped the list with ASCI White. By November 2005, they again together stood at the head of the Top500 list, this time with BlueGene/L. Even more startling is the increase in power of the computers that appeared on the list. In November 1994, Livermore's most powerful computer (41st place) weighed in at an anemic 13.7 gigaFLOP/s while IBM's most powerful system (54th) delivered only 12.1 gigaFLOP/s. In November 2005, the IBM BlueGene/L system installed at LLNL (1st on the list) was the most powerful computer in the world at 280 teraFLOP/s. In 11 years, computing power from the top machine at LLNL had increased by a factor over 23,000.[22]

The LLNL-IBM partnership for ASCI produced computing platforms, simulation environment, and applications that have had a global impact on how simulations are used in conducting modern scientific investigation. The partnership also influenced the direction of the entire computer industry. In 2005, Seager commented, "One of the reasons we were successful was that we were very aggressive about evaluating the vendor partner plans from both a technical and risk perspective and very adamant about having solid risk mitigation plans, with hard decision dates and persons identified as responsible for each risk. We then tracked progress against this plan, weekly, monthly and quarterly at different levels within the company and partnership. This allowed us to identify problems as they arose and make changes, sometimes major changes, early enough to be able to keep the schedule."[23]

According to IBM's Donofrio, service in the interest of the country was intrinsic to IBM. But it was the very challenge of working with ASCI and LLNL that attracted his company. IBM researchers were excited by the technical challenges posed by ASCI, and eager to work with their colleagues at the National Laboratories. In a 2005 interview, Donofrio captured the relationship thus: "We never stopped talking. I think what went right, at least from the Livermore partnership, was that we were one-minded. We stuck to our guns and delivered more than either one expected."[24]

ASCI Organization – *Synergy at the Edge of Chaos*

The program management for ASCI needed to be innovative. Not only did the Initiative introduce new ways of doing science, it also created new ways of doing business. Part of ASCI's success was in building these new approaches. This is evidenced by the fact that these business models have spread to other parts of the NNSA and throughout the DOE.

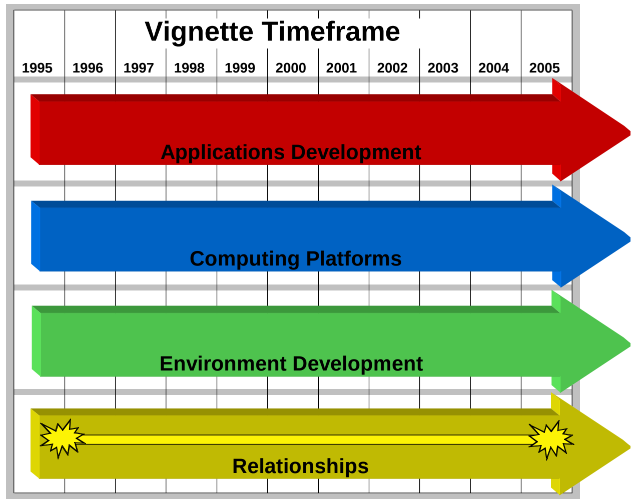

ASCI's technical accomplishments—from accelerating high-performance computer development to creating massive simulation codes running on parallel platforms—did not happen spontaneously. For every innovation, for every technological breakthrough, there was a corresponding organizational vision. The program frequently seemed on the brink of chaos, as the task of building the unprecedented capabilities at an extraordinarily rapid pace produced exceptionally challenging technical issues. Complicating the mission was the need to create synergy out of the complex relationships among the Laboratories, DP, and outside academic and industrial partners. ASCI's program management had to be as aggressively innovative as its technology development. Ultimately, the planners and administrators of the Initiative knew that the incredible demands posed by the CTBT and stockpile stewardship would require a whole new way of doing business.

One Program – Three Laboratories

The most significant aspect of the way ASCI did business was the "One Program – Three Laboratory" Strategy. On April 26, 1995, the first edition of the "ASCI Team Report" was distributed. The introduction provided an excellent overview of the collaboration that would be required for the Initiative to be successful:

> As you all know ASCI is a different kind of Defense Program initiative. It is one of the first major initiatives that has a strategy of 'One Program – Three Laboratories' from its inception. This strategy is not just a slogan. The technical problems in making the shift from test-based to compute-based stockpile stewardship are so large and so complex that no one laboratory can solve them. True collaborations, where PIs [Principal Investigators] at one lab will be relying on the work of PIs at another lab, are absolutely required.[25]

Fortunately, DP had experience with two other activities that provided good examples of how HQ and its three Laboratories could work together: SuperLab and the DP Technology Transfer Initiative (TTI).

The SuperLab initiative, begun in 1994, explored how emerging Internet technologies could be used to enhance collaboration between the Laboratories. TTI emerged as the Cold War was winding down and was seen as a possible future mission for the Laboratories. ASCI found TTI interesting because it relied on a number of teams comprising representatives of the Laboratories, DP manufacturing and processing plants, and HQ to evaluate and select technology transfer projects. The most valuable contribution that TTI provided, however, was its people, many of whom eventually worked for ASCI. Some TTI projects, such as the IBM-Tri-Lab HPSS effort on storage systems, came over to ASCI with their technical laboratory teams almost intact.

ASCI built on the lessons from SuperLab and TTI, establishing a management structure employing teams of representatives focused on particular technology areas, later specified as "strategies." From the Initiative's early planning stages onward, teams consisting of representatives from each of the three Laboratories, along with a representative from HQ, led efforts on the Applications, Platforms, and Environments strategies. An Executive Committee, consisting of the ASCI leader from each Laboratory and from HQ, provided overall direction for the Initiative.

Program Office Staff

Because ASCI was a new element in Defense Programs, it was required to form a staff office within the federal structure, yet another task that had to be completed on an accelerated basis. In an e-mail dated May 29, 1996, Steve Berggren captured how hard the staff had worked, and how far the Initiative had come in just one year:

> I would point out that just over a year ago, the "organization" consisted entirely of a brand new advisor on Reis' staff, a secretary, and two detailees from DP-10 and DP-14 [existing offices in DP]. We have journeyed through a year of "stealth staff," late hours, airline miles, airline meals, rental cars, hotel rooms and more work than we could have imagined possible. That journey has led us all to a home in the DP organization and the DP budget. Welcome home.[26]

At ASCI's inception, staffing was very tight in DP. As it turned out, a reduction in the DP-14 budget for TTI provided an opportunity for federal employees to transfer to ASCI. For example, Mike Michaelis and Tom D'Agostino, who were used to operating in the fast-paced, innovative TTI environment, came to ASCI eager for the new challenge. (Editor's Note: D'Agostino eventually became the head of NNSA, replacing Linton Brooks in 2007.) Others, including Sean Headrick, Brooks Hooper, and Thuc Hoang, came via TTI from the DP intern program, bringing with them the energy and drive characteristic of interns. While Berggren's e-mail celebrated an important milestone, it did not indicate an end to the late hours, the airline flights, or the volume of work. ASCI's success would require still further sacrifice—a common, if rarely celebrated, characteristic of federal service.

The ASCI HQ team, located in Washington, D.C., was critically supplemented by temporarily assigned staff, or detailees, from each of the Laboratories (and later from the manufacturing plants). Early ASCI detailees, such Christian Mailhiot from LLNL, Steve Turpin from LANL, and Jim Ang from SNL, were followed by many others in providing crucial technical resources for the federal government staff. ASCI HQ also looked outside federal government to recruit support, employing contractors and staff working under the Intergovernmental Personnel Act (IPA) agreements. It was the IPA, in fact, that allowed Paul Messina to take a multi-year leave of absence from the California Institute of Technology to go to Washington as ASCI's second HQ leader.

Planning Documents and Milestones

Much of ASCI's success came from a sustained focus on planning, captured in a hierarchical organization of planning documents. This yielded a full spectrum of ASCI plans, from the articulation of a long-term vision down to specific plans for the activities to be accomplished during each fiscal year. At the highest level, the vision and strategy were documented each year in the ASCI Program Plan, described thusly: "[It provides] the overall direction and policy for ASCI. This plan serves as the strategic plan for the program and identifies the key issues and work areas for ASCI."[27]

The next level of documentation comprised the Implementation Plans, or IPs. While the Program Plans were long-term strategic documents, the IPs were tactical and each focused only on work to be done in the upcoming fiscal year. The IP process started in the spring after the President's budget was submitted. By then, ASCI would have a good idea how to structure its budget for the coming fiscal year, which would begin on October 1. This timetable often had to be modified based on the evolving congressional budget process. The strategy teams worked jointly on the IPs, developing a Tri-Lab consensus on what would be accomplished each year. A primary objective of the IPs was to define how the Laboratories would collaborate on upcoming projects. During the summer, usually in July, the Strategy Committees, along with the Executive Committee, met to review the IPs. This ensured that all parties were aware of the upcoming fiscal year's work and afforded an opportunity for comment.

ASCI needed tools to ensure that it was on schedule to meet the program's ambitious goals. What's more, it needed a mechanism to demonstrate to the rest of the world that the new computational capabilities were available and could be used on problems important to stockpile stewardship. The "ASCI milestone" became the needed tool, and the milestone descriptions formed another set of important ASCI planning documents. The establishment of milestones could be used to measure and highlight accomplishments, as well as to plan for new capabilities.

The ASCI milestones were quite aggressive and often represented breakthroughs in the use of computational simulations. They also targeted work that was important to ensuring safety and reliability of the stockpile in the absence of underground testing. For example, the FY 2001 Program Plan included this milestone: "By December 31, 2001, the ASCI project will make a prototype calculation of the explosion of a full weapon system (primary + secondary) with three-dimensional features. The simulation will produce relevant information, including the primary and secondary yields that will be compared to a nuclear test."[28] This milestone was achieved and demonstrated a new, unprecedented level of simulation that had a huge impact on the Weapons Program.

The milestones have been cited as a major reason for ASCI's success. As well as being a way to communicate progress to the outside world, they also served as focal points and motivators for the people of the Initiative. The milestones served as technology targets and also as checkpoints in time. To accomplish the milestone cited above, for example, different teams within the Laboratories had to deliver the applications, the computers, and the environments and then tie all the components together into a computational system. James Peery, a former LANL ASCI Executive, commented, "The use of milestones went particularly well for ASCI. They generated an incredible amount of work and commitment to the project."[29] As an ASCI milestone deadline approached, it was common to find people working long into the night and on the weekends. Each milestone created a crucible in which disparate technology advances were melded, resulting in one more step toward the working system that enabled ASCI scientists to achieve insight.

PI Meetings

ASCI brought to DP another tool for ongoing peer monitoring and managerial support. The PI meeting, in which principal investigators shared their progress with their fellow scientists and ASCI management, was a concept borrowed from DARPA. There Vic Reis and Gil Weigand had seen the PI meeting used successfully to promulgate technical progress, as well as to provide a tool whereby management could adjust the direction of ongoing work. At semi-annual (and later annual) ASCI PI meetings, the principal researchers for each project would report on the status of their work. Fellow researchers and HQ representatives then asked questions (sometimes quite pointed), provided criticisms, and made recommendations on how to proceed.

The PI meetings had much in common with the scientific practice of peer review. One important difference was that the PI meetings exposed work that was currently underway, unlike the academic peer review process, which is generally applied to finished products. For ASCI, the PI meetings ensured that results, both positive and negative, would be available to ASCI researchers quickly. This proved a huge benefit to Lab researchers. Also, because these meetings were held in secure facilities, the ASCI researchers could freely discuss classified matters. For HQ and Laboratory managers, the PI meetings provided excellent snapshots of technical progress, and gave an opportunity to provide appropriate oversight. The meetings usually lasted three days and, as much as possible, provided a comprehensive view of the state of the technology development.

Outside Reviews

As valuable as the PI meetings were, they were conducted by ASCI insiders. The leadership also commissioned a number of outside critiques, necessary to ensure that ASCI did not become afflicted with tunnel vision and miss something critical. Some of the first external reviews were done by JASON in 1994 and 1995. JASON (the term applies to both the group as well as individual members) is an independent group of eminent scientists who gather regularly to review scientific projects related to national security. JASON operates under the auspices of MITRE, a federally funded research and development corporation. JASON provided important early recommendations that helped set ASCI on a path to realize many of its technical goals.

After the Initiative was well under way, outside reviews helped affirm its direction. In October of 1998, Reis and DP commissioned a Blue Ribbon Panel, led by Harvard University's Venkatesh Narayanamurti, to review ASCI. Reis charged the panel, asking, "Are the ASCI technical program objectives properly aligned to support our nation's rapid shift from test-based to science-based stockpile stewardship?"[30] The panel was comprehensively briefed on the Initiative between late 1998 and early 1999. On March 2, 1999, Narayanamurti replied to Reis that the panel had concluded that ASCI's technical objectives were indeed properly aligned. Although not entirely satisfied with the relative emphases ASCI placed on the various research areas, the panel conceded that they had only a limited time for the review. The panel also offered a number of recommendations that helped ASCI managers refine and improve the Initiative.[31]

Not all of the outside reviews were as friendly as that of the Blue Ribbon panel. Starting in 1998, ASCI was reviewed several times by the Government Accounting Office (GAO) of the United States Congress. The GAO (whose name was later changed to the Government Accountability Office) reviews federal programs to ensure they meet the expectations of the congressional committees that authorized them and appropriated their funds. Congressman Tom Bliley, chairman of the Commerce Committee, requested the first review, which focused on the alignment of the ASCI plans with higher-level DOE strategic plans. The GAO report, published in April 1998, found that there was not, in fact, good alignment between the plans; however, the report did acknowledge that part of the problem was that the DOE science organization was not always well-aligned with its own business lines.

Another review took place in early 1998 at the request of the Chairman of the House Budget Committee. The GAO concentrated on procurements and utilization of HPC platforms at both DP and the DOE Office of Science. The report focused on the number and cost of the supercomputers DOE had acquired from FY 1994 through 1997 and the HPC platform acquisitions planned for FY 1998 through FY 2000. The report, entitled, "Department of Energy Does Not Effectively Manage Its Supercomputers," criticized ASCI primarily for not using the process defined by the Clinger-Cohen act to document and plan for computer procurements. The report criticized the overall DOE utilization of HPC platforms, contending that major new systems were being acquired despite underutilization of the existing machines.[32]

These judgments were not without basis. It was true that the ASCI planning process differed from normal bureaucratic processes, and platform utilization rates did differ from industry standards. However, the normal bureaucratic processes and industry standards would never be able to produce the machines necessary to meet ASCI's mission. The reality was that ASCI needed the platforms up and running, with the user environment in place, in order to develop and test the applications that would ensure the safety of the nuclear stockpile. The only way ASCI could learn how to program and use teraFLOP/s computers was by having teraFLOP/s computers to use. The second criticism, that overall platform utilization was low, came mostly from a part of the computing community that resisted the move to massively parallel systems. There was a learning curve associated with effective utilization of the new systems, but ASCI was able to document that system utilization was as high as, or higher than, systems installed at other organizations. It was common that computing time on the ASCI systems were fully subscribed as soon as the machines were released for general use.

Over the course of the Initiative, ASCI was reviewed several more times. Some reviews were unfavorable, such as another GAO review entitled, "DOE Needs to Improve Oversight of $5 Billion Strategic Computing Initiative." All the reviews, favorable or not, helped ASCI sharpen its operations or clarify how its operations were perceived.

One such clarifying review was conducted by JASON in October 2003, assessing ASCI's need for platform capability versus platform capacity. JASON defined capability as the ability to run a very large single simulation, with capacity defined as the ability to run many smaller jobs. Either type of simulation could require huge enormous computing power, but each would employ that power in different forms. JASON determined that ASCI required both types of computing and recommended instituting plans to ensure that both were provided.[33]

Bob Meisner of ASCI's HQ staff was asked in 2005 what he thought about all of the reviews the Initiative had to undergo. He said, "Having served on the Air Force Systems Command's Inspector General team, I have always found reviews to be helpful. Likewise, the ASCI reviews identified areas where we could have done better."[34] ASCI reviews always resulted in a solid set of recommendations and ASCI leadership implemented them, as long as they did not cause the Initiative to lose focus on its primary mission.

Broader Impacts

ASCI's program management approach had a far-reaching impact on how DP and the Laboratories conducted business. The shift from underground testing to science-based stewardship demanded the embrace of new ideas by, and true collaboration among, all ASCI program contributors. As the first "ASCI Team Report" noted, without testing, researchers at one Laboratory would be dependent on the work of researchers at the other Laboratories.

In 1998, Weigand became Deputy Assistant Secretary for Research and Development in Defense Programs. As such, he assumed management responsibility for all of the research activities at the DP Laboratories. Based on his ASCI experience, he implemented similar management approaches into activities he called Campaigns. The idea behind Campaigns was that R&D could be focused on delivering specific new and important scientific capabilities. Milestones would be established to mark progress.

Tom D'Agostino was Weigand's deputy at that time and was responsible for implementing the Campaign management structure. Later, in 2005, D'Agostino was promoted to be the Deputy Administrator for DP (a position similar to the Assistant Secretary post held by Vic Reis). His many years of service in an accelerated, high-demand environment grant him a certain authority to comment on the impact of ASCI's organizational approach. In a 2005 interview he said, "The impact of the management approach used by ASCI on Defense Programs has been complete and comprehensive. ASCI introduced a planning process that required 5-year plans and the creation of milestones. That has been spread to the other parts of DP and now to the larger NNSA. It is now going to become part of how the DOE manages their programs. The impact of the management approach that ASCI developed and implemented should not be underestimated."[35]

---

[1] U.S. Dept. of Energy, *Nevada Test Site: Post Cold War Initiatives and Facilities*, 14.
[2] U.S. Dept. of Energy Defense Programs, *Accelerated Strategic Computing Initiative : Program Plan*, 1996, 18.
[3] Ibid., 18.
[4] Dick Watson, telephone interview by author, July 5, 2005.
[5] U.S. Dept. of Energy, *ASCI Academic Strategic Alliances Program*.
[6] U.S. Dept. of Energy, LANL Public Affairs Office, *Department of Energy, national labs select five universities*, http://www.lanl.gov/news/index.php?fuseaction=home.story&story_id=1630.
[7] Ibid.
[8] Ibid.
[9] U.S. Dept of Energy, LLNL, ASCI Academic Strategic Alliances Program (ASAP) Request for Proposal ASAP-02 Level 2 - Strategic Investigations, 1.
[10] Ibid., 2-3.
[11] Watson, interview.
[12] Moin, *Center for Integrated Turbulence Studies 2004 Annual Technical Repor*t, ii.
[13] Weigand, "Accelerated Strategic Computing Initiative: Alliance Pre-Proposal Conference" (unpublished presentation on December 5, 1996), 18.
[14] U.S. Dept. of Energy, NNSA Public Affairs Office, "NNSA Announces a New Phase of its Academic Computational Science Partnership Program," http://www.llnl.gov/asci/alliances/psaap/PSAAPPressRelease.pdf.
[15] Dongarra, Meuer, and Stohmaier, "Top 500 Supercomputer Sites," November 9, 1994.
[16] Nick Donofrio, telephone interview by author, August 23, 2005.
[17] Mark Seager, telephone interview by author, June 21, 2005.
[18] Ward, "Sample Subcontract," in *Option Purple RFP*, 9.
[19] Louis, *Computational science at full-scale*, 1.
[20] Robin Goldstone, e-mail communication to author, *IBM Working Level Story*, June 28, 2006.
[21] Ward, "Statement of Work," in *Option Purple RFP*, 70.
[22] Dongarra, Meuer, and Stohmaier, *November 2005 | TOP500 Supercomputing Sites*, [Online].
[23] Seager, interview.
[24] Donofrio, interview.
[25] Larzelere, "The ASCI Team Report (1st Ed.)," April 26, 1995, e-mail.
[26] Berggren, "Musings on the birthday of a bureaucracy," May 29, 1996, e-mail.
[27] *Accelerated Strategic Computing Initiative: Program Plan*, 1996, 21.
[28] *ASCI Programming Plan 2001*,40.
[29] James Peery, interview by author, June 20, 2005, Los Alamos National Laboratory.

[30] Narayanamurti, *Blue Ribbon Panel Report,* 1.
[31] Narayanamurti, 7.
[32] U.S. General Accounting Office, *Information Technology: Department of Energy Does Not Effectively Manage Its Supercomputers*, RCED-98-208, 2.
[33] Schwitters et al., *JASON Study – Requirements for ASCI*, 1-4.
[34] Bob Meisner, interview by author, April 21, 2005, NNSA Headquarters.
[35] Tom D'Agostino, interview by author, June 29, 2005, NNSA Headquarters.

Chapter Seven
Impact and Lessons Learned

The end of underground nuclear testing changed the questions the Laboratories needed to answer. It was no longer enough simply to know that the weapons worked, knowledge formerly obtained from a full-scale test. In an era without testing, the Laboratories needed a much better understanding of *how* the weapons worked, or, crucially, how they might fail. Would they become unreliable or unsafe as they aged? Answering those questions demands a full understanding of the fundamental physics involved. Computational simulation could deliver the answers, if the applications were detailed enough, and if the computing platforms were powerful enough, and if user environments existed to allow them to be used. If all these elements could come together, they would produce new methods of conducting scientific investigation. Computational simulation would enable scientists to attain new insight about the nuclear weapons. Moreover, scientists would have the new tools for studying the rest of the physical universe as well.

To put simulation on a par with theory and experiment, its capability had to be raised to new, unprecedented levels of power, detail, and accuracy. Rough 1D and 2D approximations, refined by experimental data, would no longer suffice. Full simulations would have to capture and predict phenomena like the weapons aging process, the behavior of aged materials under unexpected circumstances, and details of the weapons initiation sequence. ASCI-level simulations would have to capture the nature of the physical components, the correct physics governing their interaction, and do it in 3D at meaningful levels of resolution.

Mission Accomplished – Science Through Simulation

ASCI completed its first decade in 2005. Since ASCI's early planners had selected 10 years as a target to achieve a certain level of simulation capability, it was a good time to assess the progress of the Initiative. Ten years had been selected as the target date because that was when the majority of scientists with actual test experience were expected to begin retiring and because it was expected that if aging would affect the weapons, it would begin about then.

ASCI succeeded. That is clearly the most important impact of the first decade. The Initiative provided the simulation capabilities required for Science Based Stockpile Stewardship. ASCI's mission was to provide those capabilities to the U.S. nuclear weapons program to ensure that those weapons were safe and ready to deploy if necessary. In accomplishing the mission, ASCI instigated extraordinary change in the science of computational simulation, the design, acquisition, and use of HPC platforms, and the development of user environments that facilitated the use of the platforms. ASCI established new relationships among the Laboratories and forged new partnerships between the Laboratories, academia, and industry.

Beyond the ASCI mission, the initiative also was instrumental in fostering computational simulation as a fundamental tool for scientific investigation in the broader community. Naturally, having HPC platforms running large, complex simulations in a robust user environment would have been scientifically interesting in any case, but ASCI's emphasis on collaboration and partnership meant that the universities and industrial partners were first-hand participants in realizing the goal of putting simulation on a par with theory and experiment. The broader repercussions of ASCI will be felt for a long time.

Applications, Platforms, Environments

By 2005, the simulations developed by ASCI had provided weapons scientists with new ways to investigate how their systems worked. Platforms were powerful enough, had enough memory, and were fast enough that the physics models could be run at resolutions so fine that numerical error was small compared to actual physics effects. The power and size of the systems also permitted calculating the evolution of detail-level physics that had always been approximated before. The results were impressive. Simulations have led to several important discoveries. Sometimes the results confirm theories long held. Sometimes new insight has resulted. More than once, after viewing a new ASCI simulation result, weapons scientists have remarked, "Oh, so that is how it works."

ASCI succeeded despite the instability of the computer industry when the Initiative began. With some companies were entering bankruptcy and others merging, it wasn't clear that anyone could build the computers that would be needed. ASCI made strategic investments, created new procurement mechanisms, and formed partnerships with vendors. Initiative planners bet on two approaches: the MPP and the clustered SMP architectures. By 2005, the needed computers were readily available and were being used for a wide range of scientific applications. Both the MPP and SMP architectures were in place and being used effectively, as demonstrated by SNL's Red Storm and LLNL's Purple machines.

ASCI also played a critical role in developing two other high-end architectures. Clustering large numbers of PC-based computers using the Linux operating system proved to be cost effective, especially for capacity computing. The BlueGene/L platform, built in response to space, energy consumption, and cooling issues associated with other architectures, networked huge numbers of low-power "system-on-a-chip" processors.

Without an effective user environment, no computer will be useable. At the start of ASCI, like the machines themselves, a user environment for teraFLOP/s systems did not exist. ASCI responded by developing and scaling programming tools, libraries, data storage, schedulers, and a host of other elements necessary to make the massively parallel machines useful. An example the fundamental change this brought about is seen in how simulation results were visualized. The Laboratories worked with universities and industry to develop new, more scalable ways to visualize data, taking advantage of the graphic cards developed for PC video games. That feat changed—far beyond the Laboratories—how terascale visualization is done.

Management and Relationships

The Initiative's "One Program – Three Laboratories" strategy was designed to foster creative research through collaboration. ASCI's milestone approach set clear goals and demonstrated new capabilities regularly. The ASCI management techniques, with the emphases on milestones, peer review, and relationship building, spread to other parts of the NNSA, DOE, and the National Laboratories.

All of this required changing the relationships among the Laboratories, DP, other government programs, the computer industry, and universities. There was no way to build the new simulation capabilities without harnessing exceptional talent; to get it, ASCI's management had to look both inside and outside DP and its Laboratories. The HPC industry had to be induced to create new architectures, achieving new performance levels, much more quickly than it would do on its own. ASCI accomplished its accelerated goals, and helped the HPC industry survive, while letting the computer companies build systems the industry could make commercially viable. Keeping the industry in business was in fact a central, if unstated, feature of ASCI's mission.

SBSS also meant the Laboratories and the universities would have to interact in new ways. ASCI created academic centers with sufficient funding and time that universities could suggest solutions to many of ASCI's difficult technical problems. Five university Centers were created; all successfully developed simulation capabilities that helped address questions ranging from the very large (learning how neutron flashes on stars initiate) to the very small (determining how atoms of energetic materials react to shockwaves).

What Went Right

ASCI was necessary. It was also ambitious, daring, and controversial. The research for this history required more than 40 interviews, conducted with people from the National Laboratories, from DP Headquarters, from other government HPC programs, from computer companies, and from universities. Every person was asked what important lessons, positive and negative, should be taken from ASCI.

Respondents were eager to talk about what went right as well as about what could have been done better. Naturally, there was no consensus on either condition. Many thought ASCI served its mission well, that it took the right approach. Vic Reis, for instance, when asked what he would have changed, said, "Nothing."[1] Although others offered numerous ideas about what could have been done differently, everyone agreed that ASCI's accomplishments were substantial and would have an enduring impact on the weapons program, the Laboratories, and science in general. The Initiative existed in an era where many large, expensive science projects do not live up to their expectations. It is useful, then, to try to summarize what went right with ASCI. Several themes emerge from the many interviews: vision, leadership, endurance, and partnership.

Vision

The most commonly cited aspect that "went right" involved the vision put forth by the ASCI leadership. Although not everyone agreed with that vision, all acknowledged that the Initiative clearly communication it. The first ASCI Program Plan started with a clear statement that ASCI's mission was to "shift promptly from nuclear test-based methods to computation-based methods."[2] The reason for ASCI was made abundantly clear: with no underground testing, the Laboratories needed vastly improved simulation capabilities.

A succinct "elevator speech" for ASCI was not, however, sufficient to guarantee its success. The vision had to extend beyond *why* ASCI was needed, to encompass *what* the Initiative had to do. Most of the interviewees observed that ASCI's success was due in large part to the comprehensive approach taken to build simulation capabilities. The ASCI vision was broad enough to foresee the need for new, advanced technologies across all of the areas needed to support simulation-based science.

In addition, ASCI's leadership created a vision that, in addition to why and what, also anticipated *how* to do it: through the "One Program – Three Laboratories" strategy. Though it was not, at first, the easiest way for the Initiative to operate, the powerful collaborations effected through the strategy paid off handsomely in the long run. The united effort quickly forged technically credible plans that were defensible within each Laboratory, with DP, and before Congress. As the Initiative progressed, the collaborations grew to include the computer companies and to universities. ASCI's vision of why it had to be done, what had to be done, and how it would be done was the foundation upon which the Initiative was built.

Leadership

Communicating vision is one responsibility of leadership. The importance of ASCI's leadership, especially the leadership of Vic Reis and Gil Weigand, was a common theme of the interviews. Reis came to DP at a significant turning point for the nuclear weapons Laboratories. He played a key role in establishing the Science Based Stockpile Stewardship concept and knew that it would be crucial for the Laboratories to use advanced simulations to execute it. Reis helped craft the SBSS vision. Weigand served as the keeper of the vision and is credited with having the energy and tenacity to refine the vision and execute it. Weigand would not tolerate compromises that threatened the Initiative's execution. His capability to organize, fund, and execute federal research and development programs was pivotal to making it all happen. Though their roles were quite different, both Reis and Weigand employed leadership critical to ASCI's success.

Reis and Weigand are the most visible examples of "what went right" with ASCI leadership. There were, however, many other people at DP and at the Laboratories who also demonstrated outstanding leadership. ASCI created an empowering atmosphere in which participants knew they were working on something special. At every level, ASCI's people were not only allowed, but expected, to show leadership in technology areas as well as in program management. Leadership became part of the Initiative's culture, part of how business was done on a day-to-day basis.

Endurance

The people of ASCI worked under great pressure, seemingly all the time. The constant deadlines, especially those posed by the milestones, together with the frequent necessity to do something never done before, made an air of frenetic activity seem the norm. It was as though ASCI always seemed to be sprinting but had to keep the pace over the course of a marathon. The challenge was to maintain an eleventh-hour mentality while simultaneously acknowledging that fulfillment of the vision remained years away. Many of those interviewed for this history attributed ASCI's success in large part to its endurance—the steadfast commitment to its vision. The Initiative "stuck with it" long enough to achieve success.

The Initiative demonstrated its endurance across all major fields of the endeavor. Applications scientists and engineers often had to invent entirely new approaches to make codes parallel and to achieve full-physics 3D representations accurately. ASCI created a curve of increasing platform capabilities and stuck to it, keeping the computers on the very edge of the technologically possible. At the same time, the arrivals of the machines were spaced in time, giving scientists sufficient experience to understand the issues of each system and preparing them to design the next-generation platform. This step-wise approach followed the ASCI strategy of having industry build systems as large as they thought were commercially viable, and then paying to have the systems scaled to ASCI levels. With constantly evolving applications to be run on successive state-of-the-art hardware, environment developers were under constant pressure to perfect tools enabling users to work effectively.

ASCI was steadfast in pursuing its vision, but it was not hidebound. Part of the Initiative's vision, deeply imbedded in its culture, was an openness to change. ASCI freely adapted to evolving technologies and shifting circumstances. Two of the best examples of this were the development of Linux Clusters and the novel architecture of BlueGene/L. When ASCI started, nobody could have predicted that inexpensive Linux clusters could provide the power required for nuclear weapon simulations, or that machines like BlueGene/L, with processors numbering in the hundreds of thousands, could be practical. But ASCI showed the flexibility that embraced these developments, along with the stamina to see them to completion.

Partnership

For ASCI, partnership ranked high on the list of key attributes. During the interviews, several people noted the efficacy of the "One Program – Three Laboratories" approach. The strategy, not easy and not always successful, was definitely effective. It established that ASCI clearly was more than any one Laboratory and more than DP. It is an effective mechanism to create a sense of shared mission. The ASCI partners agreed that the mission was necessary, and while there were times when partners disagreed about details, all parties always agreed on the larger goals. The "One Program – Three Laboratories" partnership also allowed the Laboratories to speak with one voice regarding the Initiative.

The spirit of partnership extended throughout the Initiative, within DP and the Laboratories, certainly, but also with the external participants. ASCI formed true partnerships with the computer companies for the R&D needed to deliver the platforms, which were well beyond what could be offered commercially. In almost every case, the platforms did not turn out exactly as proposed, but research partnerships recognize that research rarely turns out as expected.

The ASCI Alliance Centers demonstrated that the ASCI approach to partnership extended into academia, as well. The Centers combined a great deal of effort by university participants and on the part Laboratory scientists to create a new simulation-based approach to science. These partnerships included exchanges of people, the use of the ASCI computing platforms, and direct financial support. These university partnerships enabled ASCI, by 2005, to declare that simulation was a sound way to apply science to nuclear weapons research, as its Alliance Centers had publicly shown on unclassified, peer-reviewed research questions.

Ten Years After

In 2005, ASCI faced its year of reckoning. Was it a success? It would be absurd to believe that an Initiative as large and complicated as ASCI was always executed perfectly. It was not. A number of reviews by the GAO leveled legitimate criticisms at the Initiative. ASCI leadership took those criticisms very seriously and sought to make improvements.

Like any research program, ASCI found that research does not always go as expected. Sometimes promising approaches failed. Sometimes the research resulted in surprises and new approaches were discovered that provided even better results. But during its 10-year window, the Initiative completed its primary mission of providing the simulation capabilities needed to support Science Based Stockpile Stewardship.

In doing so, ASCI changed many things. It changed how the National Laboratories operate and it helped revive the HPC industry. It changed change how high-end computing systems are created, and the environment in which these systems are used. Most of all, it changed how science is done by proving the use of simulations for scientific discovery.

The Initiative's success came because ASCI did the right things, at the right times, in the most important areas. In this way, ASCI represented the best tradition of American national science—face the biggest problems, assemble the best minds, lend supportive management, and drive it all with a challenge that has global implications. This recipe, which worked so well in the Manhattan Project, the Space Race, and the Cold War, was proven yet again by ASCI.

---

[1] Vic Reis, interviews by author, April 21, and June 20, 2005, telephone interview and at the DOE Forrestall Building, Washington, DC.

[2] U.S. Dept. of Energy Defense Programs, *Accelerated Strategic Computing Initiative : Program Plan*, 1996, v.

Chapter 8
Looking Toward the Future

During its first decade, ASCI accomplished all of its major objectives. Taken together, the accomplishments add up to ASCI's most important achievement, the delivery of a new means to obtain scientific insight about the performance, safety, and surety of the nation's nuclear weapons arsenal. But at its 10-year mark, ASCI's success was only a starting point. The work to build even better simulation capabilities did not end; in fact, the need for advanced simulation capabilities at the National Laboratories and DP has only continued to grow.

From the onset of the Initiative, the year 2005 was important to ASCI planners. It was estimated that the last group of scientists and engineers with underground test experience would be nearing retirement. However, in 2005 they would still have enough time, before that retirement, to validate the new simulation-based approach to stockpile stewardship. But for that to happen, the early planners reasoned, ASCI would have to cross a performance threshold by 2005. Specifically, ASCI simulations would have to be large enough, and with sufficient resolution, that errors in the simulations caused by the numerical methods would be insignificant. Then scientists and engineers could focus on correcting the errors caused by limited understanding of the underlying physics, rather than errors introduced by the way the results were calculated. Moreover, the simulations had to be computable in a reasonable amount of time—the nation could not afford to wait too long. If ASCI could build and deploy sufficiently powerful simulation systems in time, Science-Based Stockpile Stewardship could become a reality.

ASCI succeeded. By 2005, the applications, computing platforms, and user environments were in place to implement SBSS. The actual process of validating the results, correcting the physics, and refining the simulations, was only beginning, however. The great hand-off, from the generation of scientists with test experience to a new generation reliant entirely on SBSS, was just starting. DP and the Laboratories needed yet greater capabilities to allow detailed predictions of how nuclear weapons would work under a wide range of conditions. Moreover, that need continues beyond the present day.

It seems reasonable, then, to end this history of ASCI with the accomplishments of 2005. But it is worthwhile to close by examining what ASCI's future looked like at that time; to glimpse, as it were, where the Initiative was heading as it rode off into the sunset.

ASCI to ASC

To be clear, ASCI did not end in 2005. But that fiscal year marked a major transition. ASCI had changed its name several years earlier. In 2000, under Paul Messina, ASCI had officially changed its name from the Accelerated Strategic Computing Initiative to the Advanced Simulation and Computing (ASC) program, however, Messina chose to maintain the acronym ASCI because of its wide recognition. As FY 2005 began, ASC program director Dimitri Kusnezov and his staff decided it was time for the program to assume a new identity which would emphasize its evolving objectives, and a new logo using the ASC acronym was introduced. The FY 2005 ASC Program Plan describes the changing mission:

> In its first decade, the ASC strategy focused on demonstrating simulation capabilities of unprecedented scale in three spatial dimensions. The next decade will focus on increasing the ASC predictive capabilities in a three-dimensional simulation environment while maintaining the support to stockpile stewardship. To achieve the goals, ASC must continue to meet three objectives:
>
> **Objective 1. Robust Tools**
> Develop robust models, codes, and computational techniques to support stockpile needs such as refurbishments, SFIs [Significant Finding Investigations], LEPs [Life Extension Projects], annual assessments, and evolving future requirements.
>
> **Objective 2. Simulation as a Predictive Tool**
> Deliver validated physics and engineering tools to enable simulations of nuclear-weapons performances in a variety of operational environments and physical regimes and to enable risk-informed decisions about the performance, safety, and reliability of the stockpile.
>
> **Objective 3. Balanced Operational Infrastructure**
> Implement a balanced computing platform acquisition strategy and operational infrastructure to meet Directed Stockpile Work (DSW) and SSP [Stockpile Stewardship Program] needs for capacity and high-end simulation capabilities.[1]

The program plan recognized that the simulation capabilities created by ASCI would now have to be maintained and improved by ASC to provide critical resources to the weapon scientists and engineers in support of the resolution of ongoing, important stockpile issues.

Ongoing Challenges

The Laboratories' work on SBSS and ASCI has created a capability crucial to maintaining the enduring stockpile. An important part of the Laboratories' contribution is called Quantification of Margins and Uncertainties (QMU). This process is intended to understand the uncertainties in the characteristics of a weapon due to variations in design, manufacturing, or maintenance parameters, and to ensure that the weapon's performance remains within pre-defined, acceptable margins. The QMU process depends heavily on ASCI simulations to complete sensitivity studies of the designs of the well-tested, well-understood weapons. Small changes are made to design parameters and a simulation is used to predict the resulting effects. Doing this over and over again with different changes allows the design margins of the weapons to be quantified. QMU is essential to certifying existing weapons and to implementing effective LEPs.

QMU places huge demands on the tools developed by ASCI and now supported by ASC. Where a single large simulation can be employed to understand a physical process, sensitivity studies require that many simulations be run. QMU requires that *capability* simulations be run as *capacity* jobs, stressing every component of the simulation system. Incontrovertibly, the work begun by ASCI is not over. To meet the continuing challenges of SBSS, ASC must continue to support capabilities created by ASCI, and it will have to further advance the field of predictive science.

The material simulations conducted on BlueGene/L, for instance, demonstrated that ASC had new opportunities and new challenges for groundbreaking work. Material modeling had always sought to capture the characteristics of materials based on the behaviors of individual atoms. Until BlueGene/L became operational, the problem was that not enough atoms could be simulated to capture the material behavior on the larger, macro level. LLNL's simulation of the solidification of molten tantalum showed that previous simulations had only revealed the solidification of one part of a large grain formation. BlueGene/L enabled the simulation of an entire grain as it interacted with other grains.

As ASC looks to the future, predictive science, as typified by the BlueGene/L material simulations, will become ever more important. Better and more finely resolved applications will be needed; still more powerful computing platforms will be required to execute them; and enhanced user environments will remain essential to making it all work. Without doubt, ASC will face future challenges every bit as daunting as those facing ASCI in 1994.

The ASCI Challenge Met

Vic Reis put the original ASCI challenge on the table at the September 25, 1994, ASCI Workshop in Santa Fe, New Mexico, asking,

> Is [science-based stockpile] stewardship a killer application for high-performance computing? Does the issue that production and testing have disappeared have enough substance to make this an obvious national priority? Is ASCI the correct vehicle for achieving this?[2]

In a 2005 interview, Tom D'Agostino, who would become head of DP (and later director of NNSA), was able to answer those questions. He said, "The fact that we are even talking about the Reliable Replacement Warhead or replacing components in the weapons without testing is a testament to how well the Stockpile Stewardship Program has succeeded. ASCI has played a key role in allowing this kind of thinking and will definitely need to continue to make contributions."[3]

Delivering Scientific Insight

ASCI sprang into being because SBSS was a critical national priority. As important as enabling stockpile stewardship was, however, the simulation capabilities developed by

ASCI may have broader and longer-term impact in how they are changing science itself. By the 1990s scientists, whose experiments for many hundreds of years had been suggested by theory, and used to test that theory, were becoming increasingly beleaguered trying to implement that methodology. Many experiments were becoming too intricate or too expensive, requiring exceedingly complex, large, or expensive equipment. Experiments to test theories at either the cosmic or the atomic scales, or occurring over time spans of millennia or of femtoseconds, entail machinery costing billions of dollars and take years to build. How to capture information about such processes was a persistently elusive question for scientists.

Since their invention, computers have become increasingly important supplemental tools in scientific discovery. Computer modeling is used both to interpret experimental results and to design experiments that address theories. New simulation capabilities were under development at Laboratories and academia even as ASCI began, but much of the success of this approach is directly attributable to the ASCI effort. It was the Initiative that provided the urgency, the money, and the unified approach to developing simulation by advancing applications, platforms, and environments together.

With the sophisticated simulation capabilities growing out of the ASCI era, scientists now can set up a simulation with the appropriate dimensionality, resolution, physics, time span, and size to replicate a physical event with great accuracy. The resulting data are then studied in detail by scientists, who can freeze the time variable and explore the physical realm in detail. Fleeting phenomena can be halted for examination, and the nearly invisible can be made visible. With accurate underlying physics inputs, much is revealed, sometimes including surprising events not predicted by theory.

Somewhat suddenly, scientists have a new approach to discovery, one that enables them to attain new insight into the world around them. This results in new theories to explain the simulation results and new experiments to validate the simulations. In a 2005 interview, LLNL's Dona Crawford said, "The world is starting to understand how we [the National Laboratories] go about, and what we do with, simulation. It has changed the scientific discovery process and at this point we couldn't do it [science] any other way."[4]

It is impossible to predict just how far computational simulation can go in predictive science. We do know that it is already providing weapons scientists with information at resolutions never before possible and facilitating insights about how physical systems behave under the hostile conditions of nuclear weapon operations. We also know from the ASCI Alliance Centers that simulations can lead to insights about how energetic materials react to shocks, how explosives operate in fires, how turbulence flows in jet engines, how solid rockets burn, and how neutron flashes are generated on the surface of stars.

Other government scientific programs are also using simulations to develop scientific insights, including efforts by the DOE's Office of Science, the National Science Foundation, and others such as NASA, NIST, NOAA, and DARPA. The challenge, posed by SBSS, of how to understand the complex physical systems of nuclear weapons without full-scale experimentation, was taken on and met by ASCI. The legacy of that work, however, will benefit people far beyond the nuclear weapons complex. Computational simulation has now been established as a fundamental part of science. We may not know what discoveries it will bring, but we do know that new scientific insights will be gained, and they, just like ASCI, will change the world.

---

[1] National Nuclear Security Administration, Office of the Administrator, *Department of Energy FY 2005 Congressional Budget Request*, 1.
[2] Reis et al., *Summary of ASCI Workshop*, 3.
[3] Tom D'Agostino, interview by author, July 29, 2005, NNSA Headquarters.
[4] Dona Crawford, telephone interview by author, May 6, 2005.

Epilogue

This history of ASCI was commissioned in 2005, at the tenth anniversary of the beginning of the Initiative. Throughout its writing, every effort was made to characterize the Initiative in terms of its considerable impact on the nation and on science itself. The many technical innovations spawned by ASCI have been enumerated and detailed. The ways in which ASCI reshaped the workings of the DP Laboratories, and reshaped their relationships with industry, academia, and other government agencies, have been noted. The fundamental change from test-based to science-based stockpile stewardship, in which ASCI played a key role, has been discussed. Above all, this history emphasizes how ASCI facilitated and promoted the use of computational simulation as a fundamental “third leg” of science, along with theory and experiment.

In 2008, the Department of Energy established the James R. Schlesinger Award, named after the first Secretary of Energy, and announced that it would be the highest award the Department could bestow. The official description of the award is:

James R. Schlesinger Award

This award represents the highest non-monetary level of recognition an employee or contractor can receive in the Department. It is named in honor of the first Secretary of Energy and is bestowed upon one individual each year whose outstanding performance is responsible for contributions of national importance or for affecting significant improvement to the successful implementation of the Department’s mission.

The award recipient, in the tradition of Dr. Schlesinger, should have a record of consistently demonstrating outstanding leadership in public service and should exhibit the highest levels of integrity, professionalism, and dedication throughout their service to DOE. All DOE employees and contractors who meet the criteria are eligible to receive this award.

In 2008, the inaugural Schlesinger Award was bestowed on Gil Weigand. In 2009, the Award was given to Vic Reis. These two events, occurring years after the period covered by this history, may, better than any other description, place ASCI in its proper historical context, while simultaneously celebrating the two men most responsible for ASCI's existence and success. Appended here are the citations describing the accomplishments leading to these two awards, and the acceptance speeches given by Drs. Weigand and Reis. In their own words, they commemorate the accomplishments and meaning of ASCI and provide their unique perspectives of the future.

Dr. Gilbert G. Weigand

National Nuclear Security Administration

Dr. Weigand has distinguished himself with his passion for excellence along with his ability to foster and implement the practices and values that are necessary for the protection of our nation. It is because of his vision and determination that the United States is the world leader in high performance computing, that the Department of Energy leads the country with the best scientific and computing tools and scientists, and that we are able today to certify our nuclear weapons stockpile without underground nuclear testing.

Throughout his tenure at the Department of Energy, he has led the development and use of next-generation scientific and technical tools that provide the foundation of today's Stockpile Stewardship program. From this leadership position, he conceived and implemented the Department's most successful technical program to date, the Accelerated Strategic Computing Initiative (ASCI). Dr. Weigand successfully united the needs of the government programs and national laboratories with the knowledge of the U.S. computing industry, and the support of the U.S. Congress, to put together a ten year plan to build the world's best high performance supercomputers. His vision and ability to engage and organize the technical community were the driving forces behind the successes of ASCI and his implementation strategy assured rapid development and effective alignment in computer industry long-term goals and computing investments. More importantly, Dr. Weigand's contributions have impacted more than stockpile stewardship. High-performance computing and simulation, at the ASCI level, pervade all areas of science and engineering. In addition to the scientific accomplishments of the ASCI program, Dr. Weigand's efforts have provided reassurance in the safety and security of the weapons stockpile and protection of this Nation.

For his intellectual aptitude, drive, determination and unwavering commitment to upercomputing benefitting both the Department of Energy and the Nation, Dr. Weigand is presented the James R. Schlesinger Award.

Text of Dr. Weigand's acceptance speech:

Thank you Secretary Bodman and thank you Secretary Schlesinger. I am truly humbled by the enormity of this honor and I would like to acknowledge the true heroes upon whose shoulders I have had the honor to stand. They are the men and women of the simulation and modeling community. To them I would like to say, "Congratulations, we have made it. Simulation and modeling has arrived. It is wholly acknowledged today as a scientific methodology on an equal par with theory and experiment."

I would also like say to my two daughters, who are here with me today, "Thank you girls for understanding when so often you did not know why. I love you."

Recently, Secretary Schlesinger made a presentation at the energy summit sponsored by the National Academy of Sciences. During his talk, he painted a grim picture for the future of U.S. Energy Security and the possibility of energy independence for as long as we rely on oil-powered engines. However he did make one optimistic point: The solution to our national energy challenge is technology. In my own work on energy security today at Oak Ridge I have been seeking an "elevator speech" to describe the path forward in energy security and I now have it: in the precise words of Secretary Schlesinger "technology is the solution." Secretary Schlesinger went on to say that the high price of oil has stirred the country's entrepreneurial juices.

Well, given my experience at Time-Warner's AOL, I can tell you that those entrepreneur juices are a powerful brew. I believe they are part of the "magic" that makes America great. The free enterprise nature of this great country allows entrepreneurs with good ideas; to get access to capital, to put new products into the marketplace, to sometimes fail, but ultimately to succeed at solving extremely important problems, and getting rewarded for their risks and efforts. I saw this at AOL and I agree with Dr. Schlesinger that it will be part of the eventual solution to our energy challenges.

So while the Time-Warner part of me says that free enterprise will be part of the solution, the DARPA and DOE side of me says that there are some problems that the markets cannot solve on their own. These are problems that are too important – like fighting wars – or problems that are too big – like finding a cure for cancer or conquering space. They cannot be left completely to a *laissez-faire* marketplace. I believe meeting our future energy needs is in this category.

As I stand here, holding an award for leadership in technology, I have to look into a mirror and ask myself, so what now? Clearly my and my colleagues' computer and simulation accomplishments within the DOE nuclear weapons complex and at DARPA were important, but they are not the end of the story. Simulation and modeling was the driving and pacing force in Stockpile Stewardship via the ASCI program and more recently they have been the driving and pacing force in science via the SciDAC program. If we act, and act boldly, simulation and modeling likewise can be the driving and pacing force in Dr. Schlesinger's challenge, "the solution is technology."

I truly believe that the fearless application of advanced modeling and simulation, enabled by DOE's world-leading high performance computing, is precisely what our government needs to do to enable our free market entrepreneurs to discover new ways of meeting our energy needs. Whether the energy system is nuclear, solar, wind, or bio-mass we can defy the pessimists; I believe we can achieve a credible impact in 10 years. At DOE we have done so before; the existence proof exists for us.

Imagine in 10 years going into your garage and disconnecting your car from the electrical grid that's connected to a non-polluting nuclear energy power station and then traveling hundreds of miles on batteries supplemented by solar energy collectors on your roof. Also imagine when your batteries get low, a bio-fuel engine kicks-in to get you to a service station where you can exchange your batteries for a set of fully charged ones with about the same level of fuss we experience today when we fill up our car with gasoline.

I am truly grateful for this award today and for Secretary Bodman and Dr. Schlesinger being here to present it to me. I think the best way for me to justly honor what they have done and this great nation for giving me the opportunities it has is for me and the simulation and modeling community to continue our work to find new ways that "technology can be the solution."

I promise to do that and I am looking forward to continuing to work with many of you.

Thank you.

Dr. Victor H. Reis

Senior Advisor, Office of the Secretary
Former Assistant Secretary for Defense Programs

Dr. Victor H. Reis has distinguished himself over the past two decades with his creative and effective leadership in the Department of Energy and its National Laboratories, addressing vital national security interests and advancing scientific understanding. Throughout his public service at senior levels over many administrations, he has pursued his charge with unwavering dedication.

In his capacity as Assistant Secretary of Energy for Defense Programs from 1993 to 1999, Dr. Reis led the team that created science-based stockpile stewardship as an alternative to underground nuclear testing. This approach remains the keystone to ensuring the country's nuclear deterrent to this day. The formulation of the stockpile stewardship program also significantly advanced the science of simulation and has resulted in the United States being the world's leader in high performance computing.

More recently, as a Senior Advisor to the Department's leadership, his vision to use simulation and computing tools to develop advanced nuclear energy and security technologies will help ensure the clean, safe, secure expansion of nuclear power and serve the energy, security and environmental needs of this century. Dr. Reis played a key role in the establishment of an international partnership to provide nuclear fuel services and further develop advanced nuclear fuel recycling. This partnership will develop more proliferation-resistant commercial nuclear power systems and enable significant cooperation between governments, industry, laboratories, and universities.

For his lasting impact on nuclear energy, supercomputing, the nuclear weapons complex, and the national security capabilities of the Department of Energy and the Nation, Dr. Victor H. Reis is presented the James R. Schlesinger Award.

Text of Dr. Reis' acceptance speech:

Thank you, Secretary Bodman and Secretary Schlesinger.

It's been a pleasure and an honor to work with you Secretary Bodman, and Deputy Secretaries Sell and Kupfer for the past three years and it is a particular honor to accept an award named for James Schlesinger, who set the benchmark for government service: Assistant Director of the Bureau of the Budget, Chairman of the Atomic Energy Commission, Director of Central Intelligence, Secretary of Defense and the first Secretary of Energy. And when Secretary Gates needed to get to the bottom of the unscheduled flight of nuclear weapons from North Dakota to Louisiana, he called upon Jim Schlesinger.

When I joined the Department of Energy back in 1993, one of the first people I talked to was Jim Schlesinger, and we've talked at regular intervals since then. His advice is always frank and pungent.

I won't use my time today to thank friends and colleagues – though there are many who deserve thanks, and there is much to be thankful for - but to describe an integrated strategy for DOE, based upon DOE's demonstrated competence, and a vision of a global nuclear future.

But there is one person who I would like to recognize - my wife and life companion of over 52 years - Marilyn.

Marilyn was the first female Chairman of the Board of Selectmen of Marblehead Massachusetts, a member of the State of Massachusetts Transportation Advisory Committee, a producer and reporter for public radio, a long time daily volunteer at a soup kitchen, and she is now a docent at the Smithsonian American Art Museum and the National Building Museum and President of our Condominium Board of Directors.

But most important the she's the Mother of our 4 children, and Grandmother to our 10 grandchildren. I should add when she met Jim Schlesinger some years ago, they discussed the sighting of a Ross's Gull at the Parker River National Wildlife Refuge. Serious birdwatchers do that sort of thing.

Let me return to the strategy: How do you integrate DOE to solve today's energy and national security challenges?

It begins with J. Robert Oppenheimer who said:

"The atomic bomb made the prospect of future war unendurable. It has led us up those last few steps to the mountain pass; and beyond there is a different country."

The atomic bomb did not make war unendurable, but it most certainly helped to avoid the war many felt inevitable – a cataclysmic struggle with the Soviet Union. We had instead, the Cold War. The Cold War dominated international relations and our world-view throughout the latter half of the 20th Century, and its effects linger today.

We are again at a mountain pass, and I believe that the science, technology and people vested within the Department of Energy hold the key to making that different country – the world - a peaceful, prosperous and beautiful place.

Let me start with a few examples of how nuclear weapons have changed the way we live. To understand the effects of radiation on the survivors of the atomic bombs dropped upon Hiroshima and Nagasaki, DOE labs started and continue to play a critical role in the human genome project.

The U.S. space satellite program was developed to spy on the Soviet nuclear arsenal, and to provide strategic communications in the event of a Soviet nuclear attack.

Packet switching, the technology that lets computers communicate with other computers was developed to ensure that the communications to our nuclear forces survive a surprise Soviet nuclear attack. The first major application of packet switching, the Arpanet, grew into the Internet.

Packet switching was also central to the development of parallel computer processing, which driven by DOE's stockpile stewardship program has led to a factor of 10,000 improvement in supercomputing over the last 15 years.

Thus while the Genome project, the Internet, Petaflop computing and the space program might have happened without the atomic bomb and the ensuing Cold War, certainly the urgency and shape of those programs were driven in large part by nuclear weapons and the Cold War.

But Oppenheimer, Fermi, Bethe, Teller and the early nuclear pioneers also recognized the enormous potential for nuclear energy to generate power. Beginning with the nuclear submarine, much of that potential has been realized. But they also recognized that the link between civilian nuclear power and nuclear weapons - enrichment and reprocessing of fissionable material - would eventually require a global political solution.

It was President Dwight Eisenhower who saw in this conundrum a potential framework to help solve the world's economic development problems and provide a peaceful ending to the Cold War; hence his bold, prescient "Atoms for Peace" speech to the United Nations in December of 1953.

Much of Eisenhower's vision has been fulfilled: The Cold War ended without nuclear holocaust, nuclear weapon stockpiles are coming down from their Cold War high of an estimated 70,000, and all but a few nations have signed the Nuclear Nonproliferation Treaty. Nuclear energy produces some 16% of the world's electricity and the International Atomic Energy Agency helps provide a barrier between civilian and military applications.

But to continue Oppenheimer's metaphor we have reached a new mountain pass. There is continuing debate on deterrence and how to reduce the global nuclear weapons stockpile - perhaps going to zero.
There remains the enormous task to clean up the excess Cold War nuclear weapons production complex. There are new concerns about proliferation and the specter of nuclear terrorism. Finally there is the challenge of improving domestic energy security, supporting global economic development while at the same time reducing global emissions of greenhouse gasses. A landscape fraught with danger – but, I would argue - rich with opportunity.

Consider the opportunity to increase energy supply in developing nations and improving our own energy security while mitigating climate change. A strategy for doing this is to electrify the transportation sector, decarbonize the electric sector; stabilize the forests; use energy efficiently and conserve. This will mean much better batteries, more wind, solar and hydro where feasible, but the backbone of this strategy is a steady and significant replacement of coal by nuclear for baseload power.

A large increase in nuclear power will require the highest international standards for safety and security, an effective system of material control and environmental management and an acceptable solution for the disposition of high-level radioactive waste - all at an affordable price.

This must fit within a proliferation regime built upon a stable international world order that can prudently reduce the number of nuclear weapons. Think of economic development, energy, climate, strategic deterrence, cleanup and proliferation as all part of the same, interlocked global issue. No small challenge.

Let me suggest that the DOE - and its labs - have in many respects demonstrated the capability meet this challenge.

During the past 15 years, since the end of the Cold War, the DOE and its labs have:

> Completed – or almost completed – some seven world leading science installations.
>
> Created a stockpile stewardship program that has deepened our knowledge of the nuclear explosive process. We have gone through some 113 certifications without nuclear underground testing, and stockpile stewardship has driven the world's supercomputing capability.
>
> Safely dismantled over 13,000 U.S. nuclear weapons and converted some 350 tons of Russian weapons grade uranium into nuclear fuel to produce electricity in the U.S.

> Created with our Russian colleagues a first class materials control and accounting program in Russia, and is involved in material protection and security in some 108 nations.
>
> Become the major supplier of safeguard technology to the IAEA, and leads international efforts to safely expand nuclear power world wide with minimal proliferation risk.
>
> Transformed the Rocky Flats plutonium pit production complex into a National Wildlife Refuge.
>
> Submitted the Yucca Mountain Repository to the NRC for a license, and have been permanently storing actinide waste at the Waste Isolation Pilot Plant in Carlsbad New Mexico for almost 10 years.

With over 12,000 PhD's and some 25,000 visiting scientists each year, the DOE lab system is probably the world's largest collection of scientific talent - maybe ever.

And to top it all off since 1992 the DOE has been directly associated with 29 Nobel Prize winners – one of whom - Steve Chu - has been nominated to be the next Secretary of Energy.

All and all an astonishing record of accomplishment!

So what could be a DOE strategy for the new administration that builds upon DOE strengths, and has a major impact on the future of United States and the world?

Here are my suggestions.

First, reinvigorate the stockpile stewardship approach – scientific understanding and detailed simulation – and extend it to prediction and risk analysis. This would enable a more confident reduction of the U.S. nuclear stockpile, and with a concurrent improved monitoring capability, a more confident entry into a Comprehensive Test Ban Treaty

Vigorously apply this deep understanding and simulation approach to the U.S. civilian nuclear power enterprise: extend the life of current reactors and develop new fast spectrum reactors and closed fuel cycles, including fission/fusion hybrids. This means robust experiments, demonstrations and concomitant improvements of simulation capability.

And while the emphasis here is nuclear, this same approach is applicable to all aspects of the energy problem – renewables, energy storage, transportation and electricity transmission.

Deep understanding and simulation will be essential as the U.S. and the international community develops new systems to monitor, analyze and predict climate change, the effects of climate change and the efficacy of climate change mitigation measures. This increased understanding will directly feedback into the DOE's – and others – energy development efforts: an integrated climate and energy system.

Second, DOE's Environmental Management Program has suggested the installation of "energy parks" at major remaining clean-up sites; applying EM's infrastructure expertise to new private nuclear power and fuel cycle systems during continued clean up and restoration.

Such installations would provide integrating centers for DOE's widespread operations. This could include:

> Interim storage sites for spent or "used" fuel from current and new reactors while they await recycling and ultimate disposition
>
> Sites for demonstration programs (non-nuclear as well as nuclear)
>
> Integration with weapons production complex transformation
>
> "Take-back" sites for leased spent fuel from foreign entities that abjure their own enrichment and reprocessing.

The concept of fuel leasing dates back to Oppenheimer. It has been reinvigorated with the Global Nuclear Energy Partnership and endorsed by much of the international community. Every nation should have the opportunity for carbon free nuclear power, but every nation does not need the expense of indigenous enrichment and reprocessing.

Finally, the DOE should start a salt repository similar to the Waste Isolation Pilot Plant for commercial high-level radioactive waste. The Permian Salt Formation has been there for 250 million years, it is safe, secure, clean as a whistle, and the people of Carlsbad want it.

These actions would place the U.S. Department of Energy at the heart of solving some of the most difficult, complex problems the nation and the world is facing. It would provide the coherent vision to drive the DOE to become the organization that I'm sure that Jim Schlesinger anticipated when he accepted President Carter's call to create a Department of Energy over 30 years ago. It would help shape the topography of Oppenheimer's new country, and would provide the scientific and technical underpinning to redeem President Eisenhower's pledge to the United Nations on December 8 1953:

"To the making of these fateful decisions, the United States pledges before you – and therefore before the world – its determination to help solve the fearful atomic dilemma –to devote its entire heart and mind to find the way by which the miraculous inventiveness of man shall not be dedicated to his death, but consecrated to his life."

I think that best captures what the DOE and the people of DOE are all about. Thank you.

Chronology of Events

Date	Event	Location
11/89	Berlin Wall torn down	Berlin, Germany
12/91	Fall of Soviet Union	Europe and Asia
9/92	US conducts its last underground nuclear test	Nevada Test Site
10/92	President George H. W. Bush institutes unilateral moratorium on underground nuclear testing	Washington, DC
7/93	President William J. Clinton extends moratorium on testing	Washington, DC
8/93	Vic Reis becomes Assistant Secretary for Defense Programs	Washington, DC
8/94	JASON Review of Science Based Stockpile Stewardship	LaJolla, CA
9/94	ASCI workshop with computer industry	Santa Fe, NM
1/95	ASCI planning meeting and first teraFLOP/s decision	Bishops Lodge, Santa Fe, NM
2/95	ASCI briefing to Vic Reis	Washington, DC
3/95	Tri-Laboratory teraFLOP/s Summit	Livermore, CA
4/95	ASCI Red computing platform Request For Proposal released	Albuquerque, NM
4/95	First ASCI Principal Investigator Meeting	Arlington, VA
5/95	ASCI Blue computing platform Request For Information released	Livermore, CA
5/95	Option Red proposals submitted & reviewed	Albuquerque, NM
8/95	Christian Mailhiot, first Lab detailee, joins ASCI headquarters staff	Washington, DC
8/95	President Clinton Announces Support for CTBT (test ban treaty)	Washington, DC
9/95	ASCI Red selection announcement by Department of Energy Secretary Hazel O'Leary	Washington, DC
11/95	ASCI's first budget passed at $85M	Washington, DC
1/96	Last nuclear test by France	French Polynesia
2/96	ASCI Blue computing platform Request For Proposals Released	Livermore, CA & Los Alamos, NM
4/96	ASCI Blue proposals received and reviewed	Livermore, CA & Los Alamos, NM
5/96	First ASCI Headquarters Staff Meeting	Washington, DC
7/96	JASON Review of ASCI	LaJolla, CA
7/96	President Clinton announces Option Blue Awards	Washington, DC

Date	Event	Location
7/96	China conducts its last underground nuclear test	Lop Nor, China
7/96	President W. J. Clinton extends testing moratorium	Washington, DC
9/96	United Nations General Assembly approves CTBT and five declared nuclear powers sign treaty	New York, NY
12/96	Academic Alliance Level 1 Centers Request For Proposals released.	Washington, DC
12/96	ASCI Red demonstrated teraFLOP/s performance on Linpack benchmark	Beaverton, OR
12/96	Academic Alliances Level 1 Centers pre-proposal meeting	Dallas, TX
3/97	Alliances Level 1 Centers proposals received and reviewed	Washington, DC and Laboratories
5/97	Level 1 Centers site visits	Various locations
5/97	ASCI zero-based budget planning meeting to restructure the Initiative	O'Hare Airport, IL
6/97	ASCI Red delivered to Sandia	Albuquerque, NM
6/97	Fifth ASCI Principal Investigator meeting	Livermore, CA
7/97	Secretary of Energy Federico Peña announces Alliance Center 1 selections	Washington, DC
8/97	First PathForward Proposals received	
10/97	Unclassified Principal Investigator meeting for Alliance Centers	Snowbird, UT
2/98	Secretary of Energy Federico Peña announces selection of IBM to build ASCI White computing platform	Washington, DC
10/98	Blue Ribbon Panel Review of ASCI	Washington, DC
12/99	Completion of milestone for first-ever 3D simulation and visualization of nuclear weapon primary explosion	Livermore, CA
3/00	ASCI successfully demonstrates 3D hostile environments simulation.	Albuquerque, NM
4/00	ASCI completes the first-ever 3D simulation of a nuclear weapon secondary explosion.	Los Alamos, NM
6/00	Announcement of the selection of Compaq to build ASCI Q computing platform	Washington, DC
9/00	ASCI accepts delivery of ASCI White.	Livermore, CA
2/01	BlueGene/L external review	Berkeley, CA
3/01	Tri-Laboratories meet Level 1 ASCI milestones by delivering an essential application and distance-computing environment for use on ASCI White.	Albuquerque & Los Alamos, NM, and Livermore, CA

Date	Event	Location
3/01	Delivered hardware and software visualization capabilities to display simulation results at over 60 million pixels.	Albuquerque, NM
9/01	Completion of milestone for a 3D analysis for a stockpile-to-target sequence (STS) for normal environments.	
11/01	BlueGene/L R&D contract signed	Livermore, CA
4/02	ASC Purple computing platform Request For Proposals released	
6/02	Sandia announces partnership with Cray Inc. to build the Red Storm computing platform	Albuquerque, NM
9/02	Completion of milestone to conduct a 3-D system simulation of a full-system (primary and secondary) thermonuclear weapon explosion.	Livermore, CA
9/02	Completed the 3-D analysis for a nuclear weapon in a crash-and-burn accident.	Albuquerque, NM
10/02	Announcement of Cray Inc. to build Red Storm computing platform	Albuquerque, NM
11/02	ASCI Q installation complete	Los Alamos, NM
11/02	Secretary of Energy Spencer Abraham announces selection of IBM to build ASC Purple & BlueGene/L	Baltimore, MD
5/03	Completion of the ASCI milestone to simulate the stockpile to target sequence for a nuclear weapon in a hostile environment.	
7/03	JASON review of ASCI requirements	LaJolla, CA
9/03	ASCI delivers a nuclear safety simulation of a complex, abnormal, explosive initiation scenario.	
9/03	ASCI demonstrates the capability of computing electrical responses of a weapons system in a hostile (nuclear) environment	
10/05	NNSA Director Ambassador Linton Brooks commissions ASC Purple and BlueGene/L computing platforms	Livermore, CA

Appendix B

ASCI teraFLOP/s Computing Platforms

Final (2005) Configurations

Features	ASCI Red*	ASCI Blue Mountain	ASCI Blue Pacific	ASCI White	CPlant	ASCI Q
Installation Year	1997	1998	1999	2000	2001	2002
Location	Sandia	Los Alamos	Lawrence Livermore	Lawrence Livermore	Sandia	Los Alamos
Manufacturer	Intel	SGI	IBM	IBM	Self Made	HP
Peak Performance (teraFLOP/s)	3.2	3.1	3.8	12.2	1.3	20.5
LINPACK Performance (teraFLOP/s)	2.4	1.6	2.1	7.3	0.7	13.9
Compute Nodes	4,510	48	1,464	512	1,536	256
Number of Processors	9,632	6,144	5,856	8,192	1,536	2,048
Processor Type	Pentium II Xeon	MIPS R10000	PowerPC 604e	Power3 375 MHz	Alpha EV6	Alpha EV68
Memory (terabytes)	1.1	1.5	2.6	8.1	0.4	22.0
Interconnection Network	Intel 3D Mesh	ccNUMA & HiPPi	HPGN	SP Switch 2	Myrinet	Quadrics
Operating System	Unix & Cougar LWK	Irix	AIX	AIX	Linux & Puma	HP-UX

* Note that ASCI Red processors and memory were upgraded in 1999 raising the peak performance from 1.8 to 3.2 teraFLOP/s

Features	ASCI Linux Cluster (ALC)	Lightning	Red Storm	ASC Purple	BlueGene/L
Installation Year	2003	2003	2005	2005	2005
Location	Lawrence Livermore	Los Alamos	Sandia	Lawrence Livermore	Lawrence Livermore
Manufacturer	IBM	Linux Networx	Cray Inc.	IBM	IBM
Peak Performance (teraFLOP/s)	9.2	11.3	43.5	92.8	367.0
LINPACK Performance (teraFLOP/s)	6.6	8.1	36.2	75.8	280.6
Compute Nodes	960	1,408	10,368	1,536	65,536
Processors	1,920	2,816	10,368	12,288	131,072
Processor Type	Xeon 2.4Ghz	Opteron 2 GHz	AMD Opteron	Power5 1.9GHz	Power4 System on a Chip
Memory (terabytes)	3.8	5.6	31.2	49.2	32.8
Interconnection Network	Quadrics	Myrinet	Cray 3D Mesh	IBM Federated	IBM Proprietary
Operating System	Linux	Linux	Linux & Catamount	AIX	Linux

Appendix C

ASCI Leadership

Representatives to ASCI Executive Committee

<table>
<tr><th>Year</th><th>Defense Programs Headquarters</th><th>Lawrence Livermore</th><th>Los Alamos</th><th>Sandia</th></tr>
<tr><td>1995</td><td rowspan="4">Gil Weigand</td><td rowspan="7">Dave Nowak</td><td rowspan="2">John Hobson</td><td>Ed Barsis</td></tr>
<tr><td>1996</td><td rowspan="4">Bill Camp</td></tr>
<tr><td>1997</td><td rowspan="4">Don McCoy</td></tr>
<tr><td>1998</td></tr>
<tr><td>1999</td><td rowspan="2">Paul Messina</td></tr>
<tr><td>2000</td><td rowspan="3">Dona Crawford</td></tr>
<tr><td>2001</td><td rowspan="2">Bill Reed</td><td rowspan="2">Don Shirk</td></tr>
<tr><td>2002</td><td rowspan="4">Mike McCoy</td></tr>
<tr><td>2003</td><td rowspan="3">Dimitri Kusnezov</td><td rowspan="3">James Peery</td><td rowspan="2">Mike Vahle</td></tr>
<tr><td>2004</td></tr>
<tr><td>2005</td><td>Art Hale</td></tr>
</table>

About the Author[1]

Alex Larzelere II founded and served as president of Exagrid Engineering. Exagrid provided business engineering design services to create and implement architectures optimized to address customers' unique mission needs. Clients included: Sandia National Laboratories, Lawrence Livermore National Laboratory, the District of Columbia, Science Applications International Corporation (SAIC), Platform Computing, and Panta Systems. Exagrid's work was focused on understanding and analyzing impacts of various system design issues. This included providing advice on issues such as the use of high-performance computing technologies, user workflows, productivity metrics, and security.

Prior to starting Exagrid Engineering, Larzelere was the Executive Director for Advanced Computing at SAIC. There he ran a number of programs that developed architectures for high performance computing centers and grid computing research and development projects. Larzelere was also the program manager for the SAIC team that developed a concept design phase of the U.S. Coast Guard Integrated Deepwater System. This system included a collection of ships, aircraft, command and control, logistics and concepts of operations to execute 14 highly diverse Coast Guard missions. He managed a set of integrated product teams that provided a design to the Coast Guard maximizing operational effectiveness while minimizing the Coast Guard's total ownership cost.

Before joining SAIC, Larzelere was the Director of the Office of Strategic Computing at the U.S. Department of Energy. This was the office responsible for the development and execution of the Accelerated Strategic Computing Initiative. Alex was involved with the program as it started in 1994 and participated in the installation of the world's first teraflop computer at Sandia National Laboratories. Prior to the start of the ASCI program, Larzelere was the manager of the DOE Defense Programs HPC technology transfer program. There he managed a portfolio of projects that involved Los Alamos, Lawrence Livermore, and Sandia National Laboratories working with industry.

Alex Larzelere II graduated from the U.S. Coast Guard Academy and spent 10 years as a commissioned officer in the U.S. Coast Guard. He was primarily stationed in Southeast Alaska where he commanded a patrol boat and was executive officer of a seagoing buoy tender.

---

[1]Since the draft of this manuscript was written, Alex Larzelere II has rejoined the Department of Energy. He currently (April 2009) serves as Director, Advanced Simulation and Modeling Office of Nuclear Energy, Fuel Cycle Management (NE-5). The programs he directs actively implement the central philosophy of ASCI: that computational simulation is essential to a new approach to science, where it stands as a peer to theory and experiment in the process of attaining scientific insight.

Glossary

AEC – Atomic Energy Commission, the original civilian agency charged with oversight of the U. S. nuclear weapons program and nuclear energy development.

ANL – Argonne National Laboratory, a DOE national laboratory focused on non-weapons science. Located in Argonne, Illinois.

Advanced Applications – An element of the DAM program area that provides physics and geometric fidelity for weapons simulations.

Advanced Architectures – An ASC program element that is focused on development of more effective architectures for high-end simulation and computing.

Alliances – A program element within the ASCI University Partnerships.

ASC – Advanced Simulation and Computing program. This program evolved from merging of the Accelerated Strategic Computing Initiative and the Stockpile Computing Program.

ASCI – Accelerated Strategic Computing Initiative.

ASCI Blue Mountain – 3.1 tera/FLOP/s SGI clustered SMP system installed at Los Alamos National Laboratory.

ASCI Blue Pacific – 3.8 tera/FLOP/s IBM clustered SMP system installed at Livermore National Laboratory.

ASCI Q – A Compaq, now Hewlett-Packard (HP), clustered SMP system located at Los Alamos National Laboratory.

ASCI Red – 1.8 tera/FLOP/s Intel MPP system installed at Sandia National Laboratories in Albuquerque; upgraded in 1999 to 3.2 teraFLOP/s.

ASCI White– 12.2 tera/FLOP/s IBM clustered SMP system installed at Livermore National Laboratory.

ASC Purple – 100 teraFLOP/s IBM system installed at Lawrence Livermore National Laboratory.

ASC Red Storm – 40 teraFLOP/s system installed at Sandia National Laboratories in Albuquerque.

ASIC – Application-specific integrated circuit.

AST – Alliances Strategy Team.

Attack of the Killer Micros – Lawrence Livermore report, edited by Eugene Brooks, for the Massively Parallel Computing Initiative. Report predicted in 1991 that commercial microprocessors would come to dominate high performance computing platforms.

Bertha – High resolution desktop display invented by IBM with partial support from ASCI PathForward.

C – General-purpose computer programming language often used for scientific applications.

C++ – Derivative of the C programming language with many important additions including the support for object-oriented programming.

Capability system – A system that can run the most demanding large single problems using the entire resources of the machine.

Capacity system – A system that that maximizes aggregate throughput for many simultaneous smaller (but still large and demanding) problems.

CAVE – Cave Automatic Virtual Environment, a walk-in theatre at Los Alamos National Laboratory, for displaying visualizations of simulation results.

Code Teams – Teams of 20 to 30 people at Laboratories working to develop simulation applications.

COTS – Commercial-off-the-shelf, referring to technology products available from vendor standard stock.

CRS – Congressional Research Service.

CTBT – Comprehensive Test Ban Treaty, banning the testing of nuclear weapons by full-scale nuclear detonation.

DAM – Defense Applications and Modeling, the program area that focuses on development of 3-D, physics-model-based codes that are formally verified and validated.

DARHT – The Dual Axis Radiographic Hydrodynamic Test Facility at LANL will examine implosions from two different axes.

DARPA – Defense Advanced Research Projects Agency, a major research funding entity.

DISCOM2 – Distance Computing and Communication, a program element within ASC focused on computing at a distant location and data communications between geographically distant locations.

DOE – U.S. Department of Energy.

DP – Defense Programs, one of the three major programmatic elements in NNSA.

DSW – Directed Stockpile Work, those SSP activities that directly support the day-to-day work associated with the refurbishment and certification of specific weapons in the nuclear stockpile.

Earth Simulator – Once the most powerful computer in the world. Built by Japanese Earth Simulator Research and Development Center and designed to use symmetric multi-processing

ENIAC – Electronic Numerical Integrator and Computer, one of the first large-scale, electronic, digital computers capable of being reprogrammed to solve a full range of computing problems

ER – Energy Research, an element of the Department of Energy focused on non-Defense Programs research. Later became the DOE Office of Science.

ERDA – Energy Research and Development Agency, successor to the Atomic Energy Commission and predecessor to the DOE.

Finite Element – A mathematical method for representing objects in a simulation, by dividing the object into many tiny "elements" and performing calculations on each element.

FLOP/s – floating-point operations (e.g., multiplication or division) per second. A computer's performance is often measured by the number of FLOP/s it can perform.

FORTRAN – Programming language used primarily for scientific computing.

FY – Fiscal Year. The U.S. Government's fiscal year runs from October 1 through September 30.

GAO – Government Accountability Office (formerly the Government Accounting Office).

GigaFLOP/s – One billion FLOP/s.

Hero Codes – Applications written at the Laboratories by small teams of three or four people, usually with one primary author.

HP – Hewlett-Packard.

HPC – High performance computing.

HPCMod – Department of Defense High Performance Computing Modernization Program.

HPCS – High Productivity Computer Systems – DARPA project in the early 2000s that investigated technology advances in high performance computing systems.

IBM – International Business Machines.

IBM 701 – Vacuum tube based machine that ran 12 times faster than the UNIVAC and was considered to be IBM's first commercially successful scientific computer.

ID – Initial Delivery system, an innovation of the ASCI Blue procurement in which an early system would be constructed to prove the concept and be employed for applications development while the final system was under development.

I/O – Input/output.

JASON – A group of university professors who study national security issues at the request of the federal government.

LANL – Los Alamos National Laboratory, a prime contractor for NNSA, located in Los Alamos, New Mexico.

LCD – Liquid crystal display monitor.

LDRD – Laboratory Directed Research and Development, a program at the Laboratories intended to provide internal funds for basic research.

LEP – Life Extension Program.

Linpack – One of many computer performance benchmarks; Linpack codes solve systems of linear equations. Used to rate computers on the TOP500 list of the world's most powerful platforms.

Linux Cluster – Parallel computing platforms built from COTS PC technology using the Linux operating system.

LLNL – Lawrence Livermore National Laboratory, a prime contractor for NNSA, located in Livermore, California.

MCR – Multiprogrammatic Capability Resource, an 11.2 teraFLOP/s Linux cluster installed at Livermore.

MegaFLOP/s – One million FLOP/s.

MIMD – Multiple Instruction Multiple Data, a programming model in which different processors in a machine working on the same application simultaneously perform different instruction sets on different data sets.

MPI – Message Passing Interface.

MPP – Massively Parallel Processor computing platform.

NERSC – National Energy Research Supercomputing Center at the Lawrence Berkeley National Laboratory.

NEWS – Numerical Environment for Weapons Simulation, an ASCI element designed to develop user environment technologies for scalable simulations. Later incorporated into VIEWS.

NIF – National Ignition Facility at Lawrence Livermore National Laboratory, the world's largest laser facility, designed to achieve fusion ignition by 2010.

NNSA – National Nuclear Security Administration, a semi-autonomous agency within DOE.

NSA – National Security Agency.

NSF – National Science Foundation.

NTS – Nevada Test Site.

NTS – Numerical Test Site.

Office of Science – An element of DOE that focuses on (non-nuclear weapons) energy and basic science research. This includes research into advanced simulation capabilities.

OMB – Office of Management and Budget.

OOP – Object-oriented programming, a paradigm for creating computer applications codes using "objects" (pieces of codes having characteristic properties) and their interactions.

ORNL – Oak Ridge National Laboratory, a DOE non-weapons science laboratory, located in Oak Ridge, Tennessee.

PC – Personal computer.

PI – Principal Investigator, the leading researcher on a project.

PathForward – An ASC program element that partners with industry to accelerate the development of critical technology leading to commercial products needed by ASC.

PetaFLOP/s **–** 1000 trillion floating-point operations per second.

POOMA – Parallel Object Oriented Methods and Applications.

Power Wall – Large wall-mounted display system using several tiled projectors.

PSE – Problem Solving Environment, an ASC program element focused on the development of an infrastructure that provides effective software development tools, production computing environments, and archival storage.

PVM – Parallel Virtual Machine, an early interprocessor-communication system for parallel processing.

R&D – Research and development.

RFP – Request for Proposals, a standard method whereby a funding agency solicits proposals from contractors to accomplish a given task.

RRW – Reliable Replacement Warhead, a proposed program to replace existing nuclear warheads with new designs that are safer, more reliable, cheaper, and easier to maintain.

SC – The International Conference for High Performance Computing, Networking, Storage, and Analysis held annually. Usually written with the year appended, as in SC|05. Generally referred to as (i.e.,) “SuperComputing 05.”

S&CS – Simulation and Computer Science, the program element that provided the infrastructure necessary to connect applications and platforms into integrated systems.

SBSS – Science Based Stockpile Stewardship - The effort to increase understanding of the basic phenomena associated with nuclear weapons, to provide better predictive understanding of the safety and reliability of weapons, and to ensure a strong scientific and technical basis for future U.S. nuclear weapons policy objectives. Currently known as the Stockpile Stewardship Program (SSP).

SDI – Strategic Defense Initiative, often called “Star Wars,” a program for developing defensive weapons that would operate in outer space, intended to intercept and destroy enemy ballistic missiles carrying nuclear warheads.

SFI – Significant Finding Investigation. An SFI results from the discovery of some apparent anomaly with the enduring stockpile. DSW Surveillance generally initiates an SFI. For complex SFIs, resolution comes from the Assessment & Certification element of DSW, often in partnership with ASC capabilities.

SIMD – Single instruction multiple data, a programming model in which different processors in a machine working on the same application simultaneously perform the same instruction sets on different data sets.

SGI – Silicon Graphics Inc.

SLEP – Stockpile Life Extension Program. SLEP is the DP element responsible for planning and execution of component and weapons refurbishments.

SMP – Shared Memory Processor or Symmetric Multi-Processor computing platforms. Shared Memory Processors employ commodity microprocessors using a special memory-access system that allow multiple processors to see the same bank of memory.

SNL – Sandia National Laboratories, a prime contractor for NNSA with locations primarily in Albuquerque, New Mexico, and Livermore, California.

SSP – Stockpile Stewardship Program, DP's program for ensuring the safety, performance, and reliability of the U.S. nuclear stockpile in the absence of nuclear testing.

SST – Sustained Stewardship TeraFLOP/s system, the final configuration called for in the ASCI Blue procurement.

STS – Stockpile-to-target sequence, a complete description of the electrical, mechanical, and thermal environment in which a weapon must operate, from storage through delivery to a target.

SuperLab – Tri-lab program that predated ASCI to explore the use of the Internet to allow the Labs to share computing resources.

TeraFLOP/s – One trillion floating-point operations per second.

Test-based – The traditional approach used for the development of nuclear weapons, based on full-scale nuclear tests.

TTI – Technology Transfer Initiative. A tri-lab program for technology transfer that ASCI used as a model of inter-Laboratory cooperation and governance.

TOP500 List – Rankings of computer power using the Linpack benchmark. Since 1993, a new list has been published biannually.

TR – Technology Refresh, an innovation in the ASCI Blue procurement in which the Initial Delivery (ID) system would be updated to keep it on the leading-edge of technology while the full Sustained Stewardship TeraFLOP/s system was being developed.

Tri-Lab – Refers to the three NNSA laboratories: LLNL, LANL, and SNL.

TST – Technology Support Team.

UC – University of California. UC operated the Los Alamos and Lawrence Livermore Laboratories under contract from DOE (previously ERDA and AEC) from the establishment of those Laboratories until 2006 and 2007, respectively. UC continues to operate the (non-weapons) Lawrence Berkeley National Laboratory.

UNIVAC – Commercial version of the ENIAC. Univac became a catch-all name for the American manufacturers of the lines of mainframe computers bearing the name. The actual manufacturer, as a result of mergers and acquisitions, underwent numerous name changes.

V&V – Verification and Validation. Verification is the process of confirming that a computer code correctly implements the algorithms that were intended. Validation is the process of confirming that the output of a code adequately represents physical phenomena.

VIEWS – Visual Interactive Environment for Weapons Simulation. VIEWS was the previous name for DVS, the ASC program element that provides the capability for scientists and engineers to “see and understand” the results of a simulation.

Bibliography

Attaway, Stephen A., et al. "Transient Solid Dynamics Simulations on the Sandia/Intel Teraflop Computer." In Proceedings of the ACM/IEEE SC97 Conference, November 15, 1997. Available online at http://www.cs.sandia.gov/pub/papers/bahendr/HTML/pronto/pronto.html (accessed January 7, 2009).

Barsis, Ed. Interview by author. Albuquerque, New Mexico, May 18, 2005.

Berggren, Steve. *Accelerated Strategic Computing Initiative: Strategic Planning Document*. United States: U.S. Dept. of Energy, 1994. Quoted in Marshall M. Sluyter, *Memorandum to Vic Reis,* July 8, 1994.

———*Accelerated Strategic Computing Initiative: Strategic Planning Document*. United States: U.S. Dept. of Energy. Defense Programs, July 8, 1994.

——— "Musings on the birthday of a bureaucracy." E-mail communication, May 29, 1996.

Bobrowicz, Frank. *Numerical Experiments: The Nuclear Weapons High-Performance Computing for the Twenty-First Century*. Nevada: Los Alamos National Laboratory, May 2005.

Brooks, Eugene D. and Karen H. Warrens, *The 1991 MPCI Yearly Report: The Attack of the Killer Micros*. Livermore, CA: Lawrence Livermore National Laboratory, March 1991.

Brown, Jeff et al. *ASCI Debugging Requirements*. Nevada: Los Alamos, Lawrence Livermore and Sandia National Laboratories, March 28, 1998.

Christensen, Randy. Interview by author. Lawrence Livermore National Laboratories, May 17, 2005.

Christensen, Randy. *The Numerical Test Site*. California: Lawrence Livermore National Laboratories, December 1992.

Christensen, Randy and Charlie McMillan. *Justification for the Need for 100 Tflop Computers to Solve Nuclear Weapon Simulation Problems*. Livermore, CA: Lawrence Livermore National Laboratory, 1994.

Chrzanowski, Paul and Pamela MacGregor, eds., "The Eighties." In *Serving the Nation for Fifty Years: 1952-2002 Lawrence Livermore National Laboratory Fifty Years of Accomplishments*, report no. UCRL-AR-148833. Livermore, CA: U.S. Dept. of Energy, Lawrence Livermore National Laboratory.

Crawford, Dona. Telephone interview by author, May 6, 2005.

D'Agostino, Tom. Interview by author. NNSA Headquarters, July 29, 2005.

Danneskiold, Jim. "Los Alamos to get Lightning computer system from Linux Network." *Los Alamos National Laboratory Daily News Bulletin,* August 14, 2003. http://www.lanl.gov/orgs/pa/newsbulletin/2003/08/15/text02.shtml.

Deri, Bob et al. *ASCI PathForward Project*. United States: U.S. Dept. of Energy, Defense Programs, Lawrence Livermore, Sandia, and Los Alamos National Laboratories.

Dongarra, Jack, Hans Meuer, and Erich Strohmaier. "Top 500 Supercomputer Sites." Universities of Mannheim and Tennessee, November 9, 1994.

Dongarra, Jack, Horst Simon, Hans Meuer, and Top500.org. *Lightning Details.* Universities of Mannheim and Tennessee, November 1, 2005. http://www.top500.org/system/6547 (January 7, 2009).

——— "Top 500 List for November 1994." Universities of Mannheim and Tennessee. November 1, 1994. http://www.top500.org/list/1994/11/ (accessed January 7, 2009).

——— "Top 500 List for November 1995." University of Tennessee and University of Mannheim, November 1995. http://www.top500.org/list/1995/11/ (accessed January 7, 2009).

——— "Top 500 List for November 1999." University of Tennessee and University of Mannheim, November 1999. Online: http://www.top500.org/list/1999/11/ (accessed January 7, 2009).

——— "Top 500 List for November 2005." Universities of Mannheim and Tennessee, November 1, 2005. http://www.top500.org/list/2005/11/ (accessed January 7, 2009).

Donofrio, Nicholas. Telephone interview by author, August 23, 2005.

Drell, Sid et al. *Science Based Stockpile Stewardship*, JASON report no. JSR-94-345. McClean, VA: MITRE Corporation, August 10, 1994.

Ewald, Bo. Telephone interview by author, July 28, 2005.

Feng, Wu-Chun. "Making a case for Efficient Supercomputing," *ACM Queue* 1, no. 7 (2003): 54-64. Available online at: http://acmqueue.com/modules.php?name=Content&pa=showpage&pid=80&page=1 (accessed January 7, 2009).

Germann, Timothy C., Kai Kadau, and Peter S. Lomdahl. "25 Tflop/s Multibillion-Atom Molecular Dynamics Simulations and Visualization/Analysis on BlueGene/L." Gordon Bell Award finalist, presented at Gordon Bell session at SC|05, Seattle, WA, November 17, 2005. Available online at http://sc05.supercomputing.org/schedule/pdf/pap122.pdf (accessed January 7, 2009).

Gibbons, Jack. "FY 1996 Science and Technology Budget: Press Briefing." United States: White House. February 6, 1995. Available online at http://clinton1.nara.gov/White_House/EOP/OSTP/html/gibbons96.html (accessed January 7, 2009).

Goldstone, Robin. *IBM Working Level Story*, e-mail communication to author. June 28, 2006.

Goodwin, Bruce. Interview by author. Building 132, Lawrence Livermore National Laboratories, June 30, 2005.

Hale, Art. E-mail communication to Alex R. Larzelere. April 4,1995.

——— E-mail communication to Gil Weigand and Alex R. Larzelere. October 17, 1996.

Hobson, David. "U.S. Nuclear Security in the 21st Century." Arms Control Association annual meeting luncheon address, Washington, DC, February 3, 2005.

HPCWire. "DOE Orders Teraflops Supercomputer for Nuclear Test Simulations." *HPCWire* 4, no. 36.

HPCWire. "IBM Demos ASC Purple Milestone Supercomputer," Vendor Spotlight, *HPC Wire* 14, no. 29, July 22, 2005.Available online at: http://news.taborcommunications.com/msgget.jsp?mid=436979 (accessed January 7, 2009).

IBM. "IBM and Department of Energy's NNSA Partner to Expand IBM's Blue Gene Research Project." IBM press release, November 9, 2001. Available online at http://www.research.ibm.com/bluegene/press_release.html (accessed January 7, 2009).

IBM, Lawrence Livermore National Laboratory, and the BlueGene/L Team, An Overview of the BlueGene/L Supercomputer (IBM and Lawrence Livermore National Laboratories; IEEE, 25, June, 1999).

IBM, Lawrence Livermore National Laboratory, and the BlueGene/L Team. *An Overview of the BlueGene/L Supercomputer*. Technical paper presented at SC02, Baltimore, MD, November 21, 2002. Available online at https://asc.llnl.gov/computing_resources/bluegenel/pdf/sc2002-pap207.pdf (accessed January 7, 2009).

JASON. *Simulation for Stewardship*, JASON report no. JSR 96-315. McLean, VA: MITRE Corporation, June 30, 1996.

Johnston, Don. "World's fastest supercomputers dedicated at Livermore Lab." Lawrence Livermore National Laboratory press release, October 27, 2005. Available online at http://www.llnl.gov/pao/news/news_releases/2005/NR-05-10-06.html (accessed January 7, 2009).

Kent, Gary. E-mail communication to Art Hale and Alex R. Larzelere. November 26, 1996.

Kissel, Lynn. Telephone interview by the author, May 12, 2005.

Larzelere, Alexander R. "The ASCI Team Report (1st Ed.)." Organizational newsletter distributed by e-mail. April 26, 1995.

——— E-mail communication to Art Hale. April 20, 1995.

——— E-mail communication to ASCI Execs. December 16, 1996.

——— E-mail communication to ASCI. July 1, 1995.

——— E-mail communication to ASCI members. March 27, 1995.

——— E-mail communication to Gil Weigand. March 2, 1995.

Levine, Michael et al. BlueLight *Project Review*. *BlueLight* Project Review Panel, April 23, 2001.

Louis, Steve. *Computational science at full-scale is the culmination of many years of work.* Lawrence Livermore National Laboratory, March 20, 2005.

Mailhiot, Christian. Telephone interview by author, May 3, 2005.

Mattson, Timothy G. and Greg Henry. "An Overview of the Intel TFLOPS Supercomputer," *Intel*

Technology Journal 2, no. 1 (1998): 1-12. ftp://download.intel.com/technology/itj/q11998/pdf/overview.pdf (accessed January 8, 2009).

.

Medalia, Jonathan. *Nuclear Weapons: The Reliable Replacement Warhead Program*. Washington, DC: Library of Congress, Congressional Research Service, June 3, 2005.

Merky, Phil. *Beowulf History*. Beowulf.org, 2004. http://beowulf.org/overview/history.html (accessed January 8, 2009).

Messina, Paul C. et al. *ASCI Technology Prospectus*, DOE/DP/ASC-ATP-001. United States: U.S. Dept of Energy, July 2001. Available online at http://www.sandia.gov/ASC/pubs_pres/pubs/prospectus.pdf (accessed January 8, 2009).

Moin, Parviz. *Center for Integrated Turbuelence Studies 2004 Annual Technical Report*. California: Stanford University, 1994.

Monroe, Laura, Steve Stringer and Jeff Brum. "Extreme Resolution Visualization Enabled New Discoveries." *Scientific Computing* 23, no. 3 (2006): 26-31.

Morrison, John. Telephone interview by author, June 23, 2005.

Moye, William T. *ENIAC- The Army Sponsored Revolution*. United States Army, January 1996. http://ftp.arl.mil/~mike/comphist/96summary/index.html (accessed January 8, 2009).

Munoz, Jose. Telephone interview by the author, July 23, 2005.

Narayanamurti, Venkatesh. *Blue Ribbon Panel Report.* Harvard University, March 2, 1999.

National Science Foundation Blue Ribbon Panel on Simulation-Based Engineering Science. *Simulation-Based Engineering Science: Revolutionizing Engineering Science through Simulation.* Washington, DC: National Science Foundation, February 1, 2006. Available online at http://www.nsf.gov/pubs/reports/sbes_final_report.pdf (accessed January 8, 2009).

Nielsen Jr., Dale E. *General Purpose Parallel Supercomputing* Technical Report UCRL-ID-108228. Lawrence Livermore National Laboratory, 1991. Quoted in Eugene D. Brooks, et al. *The 1992 MPCI Yearly Report: Harnessing the Killer Micros*. Livermore, CA: Lawrence Livermore National Laboratory, August 1992.

Nowak, David. Telephone interview by author, June 17, 2005.

Ortiz, Michael. "Multiscale Models of Materials: Linking Microstructure and Macroscopic Behavior." California: CalTech, Presentation to Colloquium at Los Alamos National Laboratory, April 21, 2005.

Parker, Ann. "Into the Wild Blue Yonder with BlueGene/L." *Science and Technology Review*, April 2005, 23-25. Available online at http://www.llnl.gov/str/April05/Seager.html (accessed January 8, 2009).

——— "Riding the Waves of Supercomputing." *Science and Technology Review*, June 2003, 4-13. Available online at http://www.llnl.gov/str/June03/McCoy.html (accessed January 8, 2009).

——— "Strategic Computing Comes of Age." *Science and Technology Review*, June 2004, 12-20. Available online at https://www.llnl.gov/str/June04/Christensen.html (accessed January 8, 2009).

Peery, James. Interview by author. Los Alamos National Laboratory, June 20, 2005.

Post, Douglass E. and Richard P.Kendall. *Software Project Management and Quality Engineering Practices for Complex, Coupled Multi-Physics, Massively Parallel Computational Simulations: Lessons Learned from ASCI*. New Mexico: Los Alamos National Laboratory, February 24, 2003 [revised March 16, 2004].

This paper has subsequently been published in the *International Journal of High Performance Computing Applications* 18 (Winter 2004): 399-416.

Post, Douglass E. and Lawrence G. Votta. "Computational Science Demands A New Paradigm." *Physics Today* 58, no. 1 (2005): 35-41. Available at: http://www.physicstoday.org/vol-58/iss-1/p35.shtml (requires subscription).

Reis, Vic. Interviews by author. Telephone and at the Forestall Building, April 21, and June 20, 2005.

Reis, Vic et al. *Summary of ASCI Workshop.* 28-29 September 1994.

Rhodes, Richard. *Dark Sun: The Making of the Hydrogen Bomb*. New York: Simon and Schuster, 1986.

Rowell, Jan. "Full ASCI Red Goes On-line at Sandia, Sets 1.34 TFlops Record." *HPCWire* 6, no. 24, (June 20, 1997). http://www.taborcommunications.com/archives/9151.html (accessed January 8, 2009).

Schwitters, Roy et al. *Requirements for ASCI,* JASON report no. JSR-03-330. McLean, VA: MITRE Corporation, October 2003.

Seager, Mark. *LLNL Platform Strategy*. California: Lawrence Livermore National Laboratory, December 2005.

——— Telephone interview by author, June 21, 2005.

Shay, Hank. E-mail communication. October 4, 1994.

Smyth, Norris Parker. "Midsummer's Night's Nightmare: Teraflops REX." *HPCWire* 4, no. 30 (July 28, 1995).

Smith, Paul H. and John van Rosendale. *Data and Visualization Corridors*, report on the 1998 DVC workshop series. Technical report no. CACR164. California: California Institute of Technology, September 1998.

Spence, Floyd D. *The Clinton Administration and Nuclear Stockpile Stewardship: Erosion by Design.* Washington, DC: House National Security Committee, October 1996.

Sterling, Thomas. *Facing the Multi-Core Fear Factor.* Presentation at the High Performance Computing and Communications conference, Newport, RI, March 28-30, 2006. Available at:

http://www.hpcwire.com/features/17883159.html (accessed January 8, 2009)

Streitz, Frederick H., James N. Glosli, Mehul V. Patel, Bor Chan, Robert K. Yates, Bronis R. de Supinski, James Sexton, and John A. Gunnels. "100+ TFlop Solidification Simulations on BlueGene/L." Gordon Bell Award finalist, presented at Gordon Bell session at SC|05, Seattle, WA, November 17, 2005. Available online at http://sc05.supercomputing.org/schedule/pdf/pap307.pdf . (accessed January 8, 2009).

Swan, Elliot. E-mail communication to Gil Weigand. August 11, 1995.

Tarter, Bruce. Interview by author. Building 111, Lawrence Livermore National Laboratory, August 3, 2005.

Tollefson, Jeff . "Part 1 of 3: Stockpile Stalemate." *Santa Fe New Mexican*, August 1, 2004.

U.S. Dept. of Energy. *Biography of Gil Weigand*.

U.S. Dept. of Energy. Defense Programs. *Above Ground Testing Capabilities Needed in the Absence of Underground Nuclear Testing*. U.S. Dept. of Energy, October 1994.

———— *ASCI Academic Strategic Alliances Program*. U.S. Dept. of Energy, September 2000.

———— *ASCI PathForward Description*. Lawrence Livermore, Sandia and Los Alamos National Laboratories, December 27, 1996.

———— *ASCI Program Plan 1996*. U.S. Dept. of Energy, September 1, 1996.

———— *ASCI Program Plan 2001*. U.S. Dept. of Energy, October 1, 2001.

————. *ASCI Program Plan 2002-2003*. U.S. Dept. of Energy, October 1, 2002.

———— *ASCI Program Plan FY04*. U.S. Dept. of Energy, August 2003.

———— *ASC Program Plan FY05*. U.S. Dept. of Energy, September 1, 2005.

———— *Early Draft of First ASCI Program Plan*. United States: Department of Energy, October 1, 1995.

———— Lawrence Livermore National Laboratory. "IBM to Build World's Fastest Supercomputers for U.S.Department of Energy— Approaching a Half a Quadrillion Calculations Per Second, the Machines Are More Powerful Than the World's 500 Fastest Supercomputers Combined." LLNL press release, November 19, 2001. Available online at http://www.llnl.gov/pao/news/news_releases/2002/NR-02-11-08.html .

———— Lawrence Livermore National Laboratory. *System Attributes: ASCI Blue-Pacific System*. California: LLNL, October 17, 2003. Available online at http://www.llnl.gov/asci/platforms/bluepac/blue.table.html (accessed January 8, 2009).

———— Los Alamos National Laboratory. *The ASCI Blue Mountain 3-TOps System: On the Road to 00 TeraOps*. ASCI and Los Alamos National Laboratory, November 1, 1998.

———— Los Alamos National Laboratory. "Department of Energy, national labs select five universities to bolster large-scale computer simulation effort." *Los Alamos News and Public Affairs*, July 31, 1997. http://www.lanl.gov/news/index.php?fuseaction=home.story&story_id=1630 (accessed January 8, 2009).

——— *Nevada Operations Office*. *United States Nuclear Tests July 1945 through September 1992, DOE/NV*-209-REV 15. Las Vegas: U.S. Dept. of Energy, Nevada Operations Office, December 2000.

——— *Nevada Test Site: Post Cold War Initiatives and Facilities*. November 10, 1999.

——— Sandia National Laboratories. *ASC Highlights 2003*, 2004-0281P. Sandia National Laboratories, March 2004. Available online at http://www.sandia.gov/ASC/pubs_pres/pubs/ASC%20HighLts03w.pdf. (accessed January 8, 2009).

——— Sandia National Laboratories. "Red Storm to be Assembled in New Mexico." *Sandia National Laboratories Newsroom*, July 27, 2004. http://www.sandia.gov/news-center/news-releases/2004/comp-soft-math/redstormrising.html (accessed January 8, 2009).

——— "Statement of Work." In *ASCI Blue RFP*. U.S. Dept. of Energy, February 12, 1996.

——— National Nuclear Security Administration. Office of Management, Budget, and Evaluation/CFO. *FY2005 Congressional Budget Request*. Vol. 1. Washington, DC: U.S. Dept. of Energy, NNSA, February 2004.

———. National Nuclear Security Administration. Public Affairs Office. "NNSA Announces a New Phase of its Academic Computational Science Partnership Program." April 25, 2004.

U.S. General Accounting Office. *Information Technology: Department of Energy Does Not Effectively Manage Its Supercomputers*, RCED-98-208. United States: General Accounting Office, July 1, 1998. Available online at http://www.gao.gov/archive/1998/rc98208.pdf. (accessed January 8, 2009).

Ward, Gary. "Sample Subcontract." In *Option Purple RFP*. United States: U.S. Dept. of Energy, Lawrence Livermore National Laboratory, March 18, 2002.

——— "Statement of Work." In *Option Purple RFP*. United States: U.S. Dept. of Energy, Lawrence Livermore National Laboratory, February 15, 2002.

Watson, Dick. Telephone interview by author, July 5, 2005.

Weigand, Gil. Telephone interview by author, December 19, 2005.

——— "Accelerated Strategic Computing Initiative: Alliance Pre-Proposal Conference." Unpublished presentation, Dallas, TX, December 5, 1996.

White House. Office of the Press Secretary. "Statement by the President." Washington, DC: The White House, August 11, 1995. Available online at http://clinton6.nara.gov/1995/08/1995-08-11-president-statement-on-nuclear-test-ban.html. (accessed January 8, 2009).

Wolfe, Matt, Laurie McGavran, and Neil Pundit. "Scalable Programs Debugger (Etnus): Project Overview." National Nuclear Security Administration, ASC Program.

Zosel, Mary. Interview by author. Lawrence Livermore National Laboratory, California, May 17, 2005.

Figure Sources

Figure #	Source
1-1	http://www.lanl.gov/orgs/pa/News/ReisVic99.jpeg
1-2	http://ftp.arl.army.mil/ftp/historic-computers/png/eniac1.png
2-1	http://www.sandia.gov/media/periodic/STech/ST199910.pdf
2-2	U.S. Dept. of Energy, ASCI Programming Plan, 1996, 23.
3-1	http://science.nasa.gov/headlines/y2004/06may_lunarranging.htm
3-2	http://www.sandia.gov/ASC/library-pages/library-image-mesh.html
3-3	http://www.sandia.gov/ASC/library-pages/library-image-vnv-snl.html
3-4	http://www.sandia.gov/ASC/library-pages/library-image-vnv-lanl.html
3-5	http://www.sandia.gov/ASC/library-pages/library-image-teravoxel.html
3-6	ASCI Program Plan
3-7	http://sc05.supercomputing.org/schedule/pdf/pap122.pdf#search=%22spasm%20copper%20simulation%22
4-1	http://www.sandia.gov/ASC/library-pages/library-image-platforms-red.html
4-2	http://www.sandia.gov/ASC/library-pages/library-image-red-storm-fisheye.html
4-3	1996 ASCI Program Plan
4-4	http://www.sandia.gov/ASC/library-pages/library-image-platforms-bluemtn.html
4-5	http://www.llnl.gov/asci/platforms/bluepac/sstgallery/series4-2.html
4-6	SAIC Advanced Technology Group
4-7	http://www.llnl.gov/linux/mcr/build/images/P0002717_large.jpg
4-8	http://www.sandia.gov/LabNews/LN08-11-00/images/cplant_pix.jpg
4-9	http://www.es.jamstec.go.jp/esc/eng/ES/hardware.html
4-10	http://www.llnl.gov/asci/platforms/bluegene/papers/2hardware_overview.pdf
4-11	http://www.sandia.gov/ASC/library-pages/library-image-bluegene-l.html
4-12	http://www.llnl.gov/str/JulAug06/Streitz.html
4-13	http://www.llnl.gov/asci/platforms/purple/
5-1	http://www.etnus.com/TotalView/index.html
5-2	http://www.llnl.gov/icc/sdd/img/highres.shtml
5-3	http://www.sandia.gov/ASC/library-pages/library-image-DVC.html
5-4	http://www.sandia.gov/ASC/library-pages/library-image-VIEWS-029255.html
5-5	http://www.sandia.gov/ASC/library-pages/library-image-rave.html
6-1	http://www.sandia.gov/ASC/library-pages/library-image-alliances-illinois.html
6-2	http://www.sandia.gov/ASC/library-pages/library-image-alliances-utah.html
6-3	http://www.sandia.gov/ASC/library-pages/library-image-alliances-caltech.html
6-4	http://flash.uchicago.edu/website/information/AnnualReport02.pdf
6-5	http://www.sandia.gov/ASC/library-pages/library-image-alliances-stanford.html
6-6	http://www.llnl.gov/accomplishments/fifties.pdf
6-7	Mark Seager, LLNL

Index

Abraham, Spencer, 98
Adams, Tom, 45
Advanced Architectures, 95, 96, 97
Advanced Simulation and Computing (ASC), 171
Alliance Strategy Team, 141, 144
Alliances Strategy, 27, 139, 140, 141, 142, 147
Amdahl’s Law, 57
Anastasio, Michael, 102
Ang, Jim, 159
Applications
 DYNA3D, 44, 45
 FLASH, 146
 Miranda, 102
 ParaDiS, 102
 PRONTO, 61, 77
 SPaSM, 53, 102
Argonne National Laboratory, 59, 137
ASCI Executive Committee, 23, 24, 128, 144, 158, 159
ASCI Implementation Plan (IP), 159
ASCI Principal Investigator Meetings, 160
ASCI Program Plan, 23, 49, 51, 80, 118, 140, 159
ASCI Team Report, 27, 157, 162
Atomic Energy Commission (AEC), 136
Attack of the Killer Micros, 13, 55, 56, 57, 61, 150
Barsis, Ed, 58, 67, 72, 74, 77, 81
Becker, Donald, 90
Bell Telephone, 135
Berggren, Steve, 14, 15, 22, 26, 119, 158, 159
Bliley, Tom, 161
Blue Ribbon Panel, 161
Bobrowicz, Frank, 13, 58, 67
Borchers, Bob, 97
Brooks, Linton, 13, 56, 57, 62, 101, 105, 109, 150
Brown, Jeff, 120
Bush, President George H. W., 8, 41
California Institute of Technology, 51, 90, 96, 142, 145, 146, 159
Capability Computing, 88, 89
Capacity Computing, 88, 89, 162
CHAMPS, 97
Chandler, Elaine, 54
Christ, Norman, 97
Christensen, Randy, 13, 14, 42, 67, 68, 69, 88, 105
Clinton, President William J., 8, 24, 27, 41, 142
Cold War, 8, 9, 10, 11, 16, 74, 82, 135, 158
Comprehensive Test Ban Treaty (CTBT), 8, 9, 10, 41, 42, 72, 152, 157
Computational Engineering and Science Conference, 144
Computer Companies
 California Digital, 93
 Cambridge Parallel Processing, 26
 Compaq, 93, 95, 96
 Computational Engineering International, 129
 Convx Computer Corp., 26
 Cray Computer, 26
 Cray Inc., 61, 77, 121
 Cray Research, 26, 56, 67, 72, 83, 84, 86, 94, 150, 151
 Dolphin, 120, 121
 Etnus, 121
 IBM, 26, 61, 62, 67, 83, 84, 85, 92, 95, 97, 98, 99, 100, 101, 102, 103, 107, 108, 121, 130, 149, 150, 151, 152, 153, 154, 155, 156
 Intel, 26, 38, 58, 74, 75, 76, 77, 79, 85, 90, 92, 93, 150
 Kendall Square, 67, 150, 152
 Linux NetworX, 92, 93, 94
 MasPar, 67
 Meiko Scientific Inc., 25, 26, 38
 nCUBE, 26
 NEC, 98
 Remington Rand, 149
 Silicon Graphics Inc. (SGI), 60, 67, 83, 84, 121, 128, 129, 130, 151
 Sun Microsystems, 46, 61, 121
 Thinking Machines Corp., 26, 58, 67, 150, 152
Computing Platforms
 ASC Purple, 61, 91, 98, 101, 103, 105, 106, 107, 108, 109, 118, 123, 152, 153, 155, 166
 ASCI Blue Mountain, 84, 85, 86, 88, 95, 103, 105
 ASCI Blue Pacific, 79, 85, 86, 88, 95, 101, 103, 105, 107, 151, 152, 153, 154, 155
 ASCI Linux Cluster, 92, 108, 152
 ASCI Q, 91, 95, 96
 ASCI Red, 71, 72, 73, 74, 75, 76, 77, 78, 79, 81, 82, 86, 87, 88, 90, 93, 95, 96, 101, 118, 155
 ASCI Red Storm, 77, 101, 166
 ASCI White, 91, 95, 105, 107, 118, 129, 152, 153, 155
 BBN Butterfly, 25
 Beowulf Clusters, 90, 92
 BlueGene/L, 52, 53, 61, 95, 96, 97, 98, 99, 100, 101, 102, 105, 107, 108, 109, 152, 155, 166, 173
 CPlant, 93
 ENIAC, 11, 12, 56, 149
 IBM 701, 149, 150
 IBM SP1, 151
 IBM SP2, 151
 Japanese Earth Simulator, 98, 99, 101
 Lightning, 93
 Linux Clusters, 88, 91
 Multi-programmatic Cluster Resource (MCR), 92, 93, 94, 107
 Origin2000, 84
 Thinking Machine CM-2, 58
 Thinking Machine CM-5, 38, 58
 Thunder, 93
 UNIVAC, 149
Cooper, David, 151
Cooperative Research and Development Agreement (CRADA), 121

Crawford, Dona, 141
Cray, Seymour, 56
CTBT. *See* Comprehensive Test Ban Treaty
Culhane, Candace, 97
D'Agostino, Tom, 158, 173
Dally, William, 97
DARPA. *See* Defense Advanced Research Project Agency
Dayem, Hassen, 72, 81
Defense Advanced Research Project Agency (DARPA), 21, 22, 67, 82, 160, 174
Defense Programs, 9, 10, 11, 14, 15, 16, 17, 20, 22, 23, 26, 27, 30, 31, 41, 42, 52, 65, 69, 73, 74, 75, 79, 83, 86, 91, 96, 97, 118, 119, 121, 126, 128, 130, 135, 136, 138, 148, 158, 161, 162, 163, 167, 168, 169, 171, 173
Department of Commerce, 27, 28
Department of Defense, 27, 89, 135, 137
Department of Energy, 10, 16, 18, 27, 28, 31, 59, 67, 82, 89, 90, 96, 117, 119, 120, 126, 135, 141, 147, 161, 162, 163, 167, 174
Deri, Bob, 91
Directed Stockpile Work (DSW), 172
Distributed Computing and Distributed Communications, 29
DOE Energy Research, 27
DOE Office of Science, 89, 135, 141, 161, 174
Dolphin, 120
Dongarra, Jack, 67, 149
Donofrio, Nick, 62, 103, 152, 156
Dual Axis Radiographic Hydrodynamic Test Facility, 41
Empirical Data, 50
Energy Research and Development Agency, 136
Ewald, Bo, 94
Fat Man, 8
File Systems
  Lustre, 91, 92
Frameworks
  Parallel Object Oriented Methods and Applications (POOMA), 59
  Sierra, 59
Galileo, 34, 35, 36, 50
General Motors, 21
Germann, Timothy, 53
Gerstner, Lou, 26, 152
Gibbons, John, 27
Goldstone, Robin, 154
Goodwin, Bruce, 48
Gordon Bell Award, 76, 102
Government Accountability Office, 161, 162, 170
Government Accounting Office, 161
Hale, Art, 74, 91
Hammer, Derrol, 141
Hamming, R. W., 132
Hayes, Ann, 141
Headrick, Sean, 158
Hecker, Sig, 24
High Performance Storage System, 115, 158
High Productivity Computing Systems, 61
High-Performance Computing Modernization Program (DoD), 137
Hiroshima, 8
Hoang, Thuc, 141, 158
Hobson, David, 131
Hooper, Brooks, 158
Hopson, John, 25
HPCWire, 57, 75, 76
IBM Thomas J. Watson Research Center, 97, 98
Immele, John, 9
Interconnection Networks
  Ethernet, 90, 100
  HIPPI, 84, 85
  IBM Federated Switch, 108
  Myricom, 93
  Quadrics, 91, 92, 93
  SeaStar, 77
Intergovernmental Personnel Agreements (IPA), 159
JASON, 26, 89, 161, 162
Job Manager
  LoadLeveler, 108
Kadau, Kai, 53, 62
Kendall, Richard, 48, 67, 150, 152
Kent, Gary, 91
Kissel, Lynn, 98
Kusnezov, Dimitri, 135, 171
Laboratory Directed Research & Development, 55, 89
LANL. *See* Los Alamos National Laboratory
Lawrence Berkeley National Laboratory, 108, 137, 149
Lawrence Livermore National Laboratory, 8, 9, 13, 14, 20, 25, 42, 44, 45, 48, 52, 54, 61, 72, 81, 82, 83, 84, 86, 91, 92, 93, 94, 97, 98, 100, 101, 102, 103, 107, 108, 116, 118, 119, 120, 121, 129, 131, 135, 141, 149, 150, 151, 153, 154, 155, 156, 159, 173
Levine, Michael, 97
Life Extension Projects (LEP), 172
Linnebur, Eldon, 58
Little Boy, 8
LLNL. *See* Lawrence Livermore National Laboratory
Lockheed Martin, 135
Lomdahl, Peter, 53, 62
Los Alamos National Laboratory, 8, 9, 13, 15, 18, 20, 24, 25, 44, 46, 48, 51, 53, 58, 59, 72, 81, 82, 83, 84, 85, 86, 91, 93, 95, 96, 116, 120, 129, 131, 135, 141, 150, 159, 160
Lucas, Robert F., 97
Mailhiot, Christian, 52, 159
Manhattan Project, 7, 135, 139
Martin Marietta, 135
Massively Parallel Computing Initiative, 25, 55, 61, 62, 150
McCoy, Mike, 44
Meisner, Bob, 162
Message Passing Libraries
  Message Passing Interface, 59, 61, 122
  MPI Chameleon (MPICH), 59
  Open MPI, 59, 122
  Parallel Virtual Machine, 59

Messina, Paul, 96, 135, 159, 171
Meuer, Hans, 149
Michaelis, Mike, 158
Milestones, 45, 159, 160
MITRE, 161
Moin, Parviz, 147
Moniz, Ernie, 141, 142
Moore, Gordon, 79
Morrison, John, 81, 93
Morse, Norm, 91
MPCI. *See* Massively Parallel Computing Initiative
Munoz, Jose, 96
Nagasaki, 8
Narayanamurti, Venkatesh, 161
NASA. *See* National Aeronautics and Space Administration
National Aeronautics and Space Administration (NASA), 27, 88, 90, 101, 126, 174
National Energy Supercomputer Center, 137
National Ignition Facility (NIF), 41
National Institute of Science and Technology (NIST), 174
national Nuclear Security Administration (NNSA), 173
National Nuclear Security Administration (NNSA), 7, 31, 101, 103, 105, 136, 158, 163
National Oceanographic and Atmospheric Administration (NOAA), 174
National Science Foundation (NSF), 54, 126, 137, 174
National Security Agency, 137
NNSA. *See* National Nuclear Security Administration
Not-invented-here syndrome, 136
Nowak, David, 23, 24, 52, 96, 97
Numerical Environment for Weapons Simulation (NEWS), 128
Numerical Test Site, 13, 14
Oak Ridge National Laboratory, 59
One Program – Three Laboratories Strategy, 23, 29, 136, 157, 168, 169
Operating Systems
  AIX, 108
  Linux, 87, 88, 90, 91, 92, 93, 94, 107, 152, 166, 169
Oppenheimer, Robert, 8
Ortiz, Michael, 51, 52
PathForward, 91, 92, 93, 94, 109, 121, 152
Patrick, Merrell, 141
PC video games, 129
Peery, James, 44, 160
Peña, Federico, 142
Post, Doug, 46, 47, 48, 112
Power wall. *See* Visualization Displays
Predictive Science Academic Alliance Program, 148
Principal Investigator Meetings, 160
Problem Solving Environment (PSE), 30, 119, 129
Processors
  AMD Opteron, 77, 93
  DEC Alpha, 93
  Intel Itanium2, 93
  Intel Pentium, 75
  Intel Xeon, 92
  Power5, 108
  PowerPC, 99
Programming Languages
  C, 24, 52, 58, 60, 121, 144, 159
  C++, 60, 121
  FORTRAN, 60
  High Performance Fortan, 60
  OpenMP, 121, 122
  Unified Parallel C, 60
Programming Tools
  PAPI, 122
  Tau, 122
  TotalView, 120, 121, 122
  Vampire, 122
PSE. *See* Problem Solving Environment
Purdue University, 21
Quantification of Margins and Uncertainties (QMU), 172, 173
R&D 100 award, 129, 130
Ratter, Justin, 76
Reagan, President Ronald W., 10
Reed, Bill, 135
Reis, Vic, 10, 14, 15, 16, 17, 20, 21, 22, 23, 26, 30, 38, 62, 67, 72, 73, 78, 82, 84, 119, 135, 152, 158, 160, 161, 163, 167, 168, 173
Reliable Replacement Warhead, 173
Request for Proposals, 74, 83, 92, 98, 106, 107, 143, 155
Reynders, John, 59
Robinson, C. Paul, 24
Rosendale, John van, 126, 128, 132
Sandia National Laboratories, 8, 9, 14, 20, 21, 24, 44, 58, 59, 61, 71, 72, 73, 74, 75, 76, 77, 79, 81, 82, 85, 86, 91, 93, 102, 109, 116, 120, 131, 135, 141, 159
SBSS. *See* Science Based Stockpile Stewardship
Science Based Stockpile Stewardship, 11, 15, 26, 42, 46, 57, 65, 72, 147, 152, 157, 165, 167, 168, 170, 172, 173, 174
Seager, Mark, 61, 91, 94, 97, 107, 108, 120, 153, 154
Shared Memory Processors, 60, 81, 82, 84, 87, 91, 105, 118, 143, 151
Significant Finding Investigations (SFI), 172
Simon, Horst, 149
Simulation & Computer Science (S&CS), 119
Site Y (Los Alamos Laboratory), 139
Smith, Paul, 91, 126, 128, 132
Smyth, Norris Parker, 75, 78
SNL. *See* Sandia National Laboratories
Software Engineering Institute, 39
Soviet Union, 8
Spence, Floyd, 24
Stanford University, 130, 142, 144, 145, 147, 148
Sterling, Tom, 57, 90
Stockpile Stewardship Program (SSP). *See* Science Based Stockpile Stewardship
Stovepiping, 144
Strategic Defense Initiative, 10
Streitz, Frederick, 102

Strohmaier, Erich, 149
SuperComputing Conference
  SC|02, 98
  SC|05, 53, 101
  SC|97, 76
SuperLab, 119, 158
Tarter, Bruce, 25, 151
Technology Support Team, 144
Technology Transfer Initiative, 28, 158
Terascale Simulation Facility, 108
Theory of Relativity, 7
Tomkins, Jim, 74
TOP500 List, 93, 101, 110, 149, 150, 152, 155
Turpin, Steve, 159
U.S. Congress, 21, 22, 23, 29, 31, 161, 168
University of California, 135
University of Chicago, 102, 142, 145, 146, 148
University of Illinois at Urbana/Champaign, 143, 144, 145, 148
University of Utah, 142, 144, 145, 148
University of Utah/Salt Lake, 145
University of Virginia, 130
Verification and Validation, 29
Visual Interactive Environment for Weapon Simulations, 131
Visual Interactive Environment for Weapons Simulation (VIEWS), 29, 128, 130, 131, 132
Visualization Displays
  Bertha, 130, 152
  Cave Automatic Virtual Environment (CAVE), 128, 131
  Power Walls, 131
Visualization Software
  Chromium, 130
  Ensight Gold, 129, 130
  Kitware, 129
  ParaView, 129, 130
  VisIt, 129, 130
  Visual Tool Kit, 129
Voigt, Bob, 141
Watson, Dick, 91, 97, 98, 141
Weigand, Gil, 21, 22, 23, 25, 26, 30, 38, 43, 52, 62, 67, 72, 73, 74, 76, 78, 81, 82, 84, 96, 132, 135, 139, 141, 147, 148, 150, 152, 160, 162, 163, 168
Westinghouse, 21
Wheat, Stephen, 74
William and Mary, College of, 141
Williams, Alan, 120
World War II, 7, 136
Zosel, Mary, 118, 120, 123
Zwakenberg, Rick, 120