

Safety in Numbers:
Creating and Contesting the Los Alamos Approach to Supercomputing, 1943 to 1980

A DISSERTATION
SUBMITTED TO THE FACULTY OF THE
UNIVERSITY OF MINNESOTA
BY

Nicholas Dale Lewis

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Thomas J. Misa, Advisor

August 2019

Copyright by Nicholas Dale Lewis

2019

Acknowledgements

This work would not have been possible without the financial support of the High-Performance Computing (HPC) Division, the Advanced Simulation and Computing (ASC) Program, and the Associate Directorate for Theory, Simulation, and Computation (ADTSC) at Los Alamos National Laboratory (LANL). I am indebted to those who created and provided key support for the HPC History of LANL Supercomputing Project, including, but not limited to, HPC Division leader Gary Grider, and HPC-Systems group leader Jeff Johnson. My boundless gratitude to Freddie Marshall for the herculean scanning and organizing effort that made the LANL archival resources underlying most of this dissertation readily accessible. Additional thanks to Thomas Herrera, Zachary Lewis, Zephram Tripp, and Pamela Smith for their support work on the document database. I am especially indebted for the crucial and tireless support of my LANL mentor, Carolyn Connor, whose vision for capturing the history of Los Alamos computing not only introduced me to Los Alamos and made this dissertation possible, but has been an unflagging source of encouragement and inspiration, no matter my undertaking on the HPC History Project.

Special thanks to LANL historian Alan Carr and the staff of the LANL Archives for going above and beyond in providing research access and assistance. Additional thanks to the current and retired LANL personnel, and the contributors outside LANL, who provided their time for the 60 interviews conducted for this project, and to those who also provided documentation, including William J. “Jack” Worlton, Ronald Christman, David Kratzer, Don Tolmie, Paul Iwanchuk, and Terri Morris.

I want to thank my committee members at the University of Minnesota for their invaluable assistance and comments that have made this dissertation a much stronger work of scholarship. I owe a great debt to my advisor, Thomas J. Misa, whose encouragement, advice, and guidance have made me a better scholar, and were essential for the success of this dissertation.

I also want to thank my greatest source of support, my family, including my parents, Dale and Linda Lewis, who have always encouraged my educational pursuits, my partner and favorite historian, Aimee Slaughter, and my frequent writing companion, Smaug the Cat.

Dedication

This dissertation is dedicated to my father, Dale Lewis. Although he did not get to see the final product, his unfailing support and dedication to his family are behind every page.

Abstract

Since its origins during World War II as the primary R&D site of the Manhattan Project, Los Alamos National Laboratory has pushed the limits of computational technologies and systems. In response to its heavy computing demands, and the infancy of the computing industry during the early Cold War, the Theoretical Physics (T) Division at Los Alamos, which provided computing resources for all divisions at the Lab, played an active role in researching and developing the computing technologies that its scientists required. The locally built MANIAC computer and the Monte Carlo method of simulation were among T Division's early developments. T Division's hardware and software R&D effort formed the basis of the distinctive Los Alamos approach to computing, which placed computing on an equal footing with other scientific research, and encouraged the formation of local, world-class expertise in computing technologies essential to the Lab's mission. As particular computing technologies matured, and vendors were able to provide commercial products adequate for the Lab's needs, T Division shifted its R&D focus to other, less mature lines of computing research. For example, following T Division's design collaboration with IBM on the "Stretch" supercomputer, and the maturation of the supercomputer industry with the Control Data 6600, Los Alamos transitioned its R&D efforts away from computer hardware, to supercomputer networking, data storage, operating systems, and multi-dimensional simulation codes. This approach provided Los Alamos computer users with technologies that were otherwise unavailable or unsuitable from outside sources, and ensured that T Division could rapidly adapt the Lab's computing operation to unpredictable changes in the technical and strategic demands of the Cold War. The Los Alamos approach to computing remained largely uncontested until its transfer in 1968 with T-Division personnel to the Lab's new Computing (C) Division, where the inherited approach became a point of contention for a growing number of weapons-program computer users who opposed the Lab's basic computer-science research and C Division's operational independence. Radical changes to the Lab's administrative structure over the 1970s upended the power dynamics between the supporters of the traditional Los Alamos approach to computing and those who sought to bring the Lab's computing resources under the direct control of the weapons program. A focal point for critics of C Division's R&D policies, the local development of a cutting-edge operating system for the iconic Cray-1 supercomputer became the final battle ground over the decades-old Los Alamos approach to computing.

Table of Contents

Acknowledgements	i
Dedication	ii
Abstract	iii
Table of Contents	iv
Preface	vi
Chapter One—Origins of Los Alamos Scientific Computing:	
Literature and Context	1
I. Introduction	1
II. Literature Review	3
III. The Origins of Los Alamos Computing	35
Chapter Two—MANIAC to Stretch:	
Making the Computer "Super" at Los Alamos	85
I. Chapter Introduction	85
II. The T-Division Approach to Computing	87
III. LASL and IBM	93
IV. The Stretch Project	122
V. The Home Stretch	149
Chapter Three—The Ties that Divide:	
Supercomputer Selection and the Formation of C Division	161
I. Chapter Introduction	161
II. "Super Stretch" and the Supercomputer Market	162
III. The Committees	181
IV. The Search	187
V. The Selection	192
VI. Changing Plans	202
VII. The Computing Sciences and Services Division	210
VIII. Chapter Conclusion	216

Chapter Four—Unshared Experiences:	
Time-Sharing and the	
First Intervention in Los Alamos Computing	219
I. Chapter Introduction	219
II. A Time to Share	222
III. "We can do it better."	240
IV. Reorganization	258
V. Intervening in C Division	262
VI. Chapter Conclusion	271
Chapter Five—The End of the Beginning:	
The Cray-1, DEMOS, and the	
Reorganization of C Division	275
I. Chapter Introduction	275
II. Control Data and the Case of the Wrong Computer	276
III. Falling STAR	292
IV. Suddenly Seymour	302
V. The Route Less Taken	333
VI. DEMOS and the End of the Beginning	343
VII. Chapter and Project Conclusion	369
Bibliography	377
Archival Collections	390
A Note on LANL Archival Collections	390
Cited Original Oral History Interviews	391

Preface

In 2013, the High-Performance Computing (HPC) Division, the division responsible for the supercomputing resources at Los Alamos National Laboratory (LANL), formed the HPC History of LANL Supercomputing Project to document and preserve the history of computing at Los Alamos. Recognizing the depth and importance of that history, the project entered a collaboration in 2014 with the Charles Babbage Institute (CBI) at the University of Minnesota in order to foster an academically driven program to capture and study the legacy of Los Alamos computing. The project focused on the collection and analysis of primary source materials, which included conducting oral-history interviews with current and retired LANL staffers, and negotiating an unprecedented lending agreement with the LANL archives. Project members borrowed computing-related collections from the LANL archives, and scanned them into an online database maintained by HPC Division. This dissertation is based upon the access and materials provided by the HPC History of LANL Supercomputing Project, and the cooperation and support of many current and retired LANL staff. The oral-history interviews produced as part of this project are permanently housed at LANL in the HPC Division digital archive.

-Chapter 1-

Origins of Los Alamos Scientific Computing: Literature and Context

I. Introduction

Between the early Cold War and 1980, the divisions in charge of computing at Los Alamos Scientific Laboratory (LASL) conducted extensive basic and applied computing R&D in order to supply Lab scientists with resources and technologies that computing vendors were either unable or unwilling to provide. This distinctive Los Alamos approach to computing, which elevated computing as an equal among the sciences at the largely physics-driven Lab, emerged out of the Theoretical Physics or "T" Division's unusually demanding computing requirements for its work in thermonuclear-weapons development during the early Cold War. T Division's leaders maintained this approach over subsequent decades to ensure that LASL's computing operation remained responsive and flexible in the face of changing Cold War pressures. When a particular field of computing matured, and the products of vendors were finally deemed adequate for the Lab's requirements, T Division transitioned its R&D focus to new computing technologies. The maturation of mainframe hardware in the mid-to-late 1950s, for example, allowed Los Alamos to purchase its production computers (the core of the Lab's extensive nuclear-physics and weapon simulations) rather than constructing them locally. As a consequence, T Division could focus more heavily on software R&D, an area where the products of computing vendors were still largely inadequate. This process of focusing and refocusing the Lab's computing R&D efforts repeated frequently at Los

Alamos with technologies such as supercomputer networking, fast high-volume data storage, operating-system development, and others.¹

The Los Alamos approach to computing, which successfully transferred with T-Division computing personnel in 1968 to the Lab's first dedicated Computing or "C" Division, remained largely uncontested while T Division controlled, and was the primary user of, LASL's computing resources. By the late 1960s and early 1970s, diminishing returns on extracting greater efficiencies out of nuclear-weapons designs, and enhanced administrative accountability pressures, resulted in larger numbers of weapons-program personnel using LASL's computer resources on a regular basis. Unlike early T-Division users, who were, by necessity, often experts in the computing technologies they utilized, many of the newer users from the weapons program had less experience or familiarity with the challenges of designing and maintaining LASL's computing operation. Growing friction between newer and more established users over C Division's R&D strategies and operational independence brought the Los Alamos approach to computing into question for the first time. Escalating tensions between these two groups in the 1970s came to a head when an administrative reorganization at the Lab brought weapons-program users far greater financial and political power over the computing operation, resulting in a sudden and definitive end to the decades-old Los Alamos approach to computing in 1980.

Based upon unprecedented access to largely untouched primary-source materials from the Los Alamos archives, and 60 original oral-history interviews with current and former Lab personnel, this dissertation provides an in-depth examination of the

¹ LA-UR-18-26793.

development and changing roles of computing at a US nuclear-weapons laboratory between World War II and the latter years of the Cold War. More than an institutional history, this dissertation situates the evolution of Los Alamos computing within key contextual frameworks, including the interplay between Lab computing experts and various sectors of the US computing industry, the changing Cold War political and strategic demands that strongly shaped LASL's mission objectives and computing requirements, and the important cultural distinctions between the labs of the US nuclear complex that resulted in unique approaches to computing, particularly evident between Los Alamos and its chief rival lab, Lawrence Livermore. Through extraordinary access to usually closed Los Alamos sources, this dissertation contributes to the historiographies of computing, the US nuclear-weapons complex, and, specifically, of Los Alamos during the Cold War, examining how the distinctive Los Alamos approach to computing reflected and shaped the Lab's unique culture and place within the rapidly changing Cold War environment.

II. Literature Review

The Laboratory at Los Alamos, New Mexico is best known as the birthplace of the atomic bomb during World War II. Most histories of Los Alamos focus on its role as the primary R&D site of the Manhattan Project, with comparatively few academic and popular accounts examining the continued existence and evolution of Los Alamos as it transitioned from a wartime crash effort, to a permanent nuclear-weapons-development

facility over the course of the Cold War.² A variety of factors have contributed to this focus on the Manhattan Project in the existing Los Alamos literature, and, correspondingly, the limited selection of academically rigorous histories of Cold War Los Alamos and its collection of Cold War research programs, including scientific computing. The contributing factors include a long-standing public interest in the wartime Laboratory, Cold War security issues, the Lab's own deliberate public-relations efforts, and even the Lab's internal presentations and perceptions of how it fits within the larger US nuclear-weapons complex. Those factors have done much to shape the research and presentation of Los Alamos history in ways sometimes unique even among the other US weapons laboratories.

Important academic histories of Cold War Los Alamos do exist, such as Anne Fitzpatrick's "Igniting of the Light Elements" and Peter Westwick's *The National Labs*. Although, as unusual as they are in their depth of research and analysis on Cold War Los Alamos, they are also circumscribed in the timeframes they address, and in the features of Los Alamos culture and organization they describe.³ Most academic accounts relating to

² Richard Rhodes' *Dark Sun: The Making of the Hydrogen Bomb* (New York: Simon and Schuster, 1995) is among the most popular and accessible accounts of post-World-War-II Los Alamos, discussing the protracted effort to develop the first thermonuclear weapon. Although, the qualitative differences between *Dark Sun* and Rhodes' earlier account of the Manhattan Project, *The Making of the Atomic Bomb* (New York: Simon and Schuster, 1987), highlight the comparative dearth of information available for the post-World-War-II period. *Dark Sun* suffered from a lack of source material on the title subject, and was heavily padded with unrelated, but more stimulating accounts of World War II espionage, and dramatized discussions of early Soviet atomic-bomb efforts.

³ Anne Fitzpatrick's "Igniting the Light Elements: The Los Alamos Thermonuclear Weapon Project, 1942-1952" (PhD diss., George Washington University, 1999) remains the most in-depth study of the early development of thermonuclear weapons at Los Alamos. Fitzpatrick, in her capacity as a Q-cleared consultant at LANL, was able to conduct extensive archival research into the early development of thermonuclear weapons at Los Alamos. Fitzpatrick's dissertation is a detailed, analytical look at the US nuclear-weapons complex as a technological system, which included one of the few scholarly examinations of Los Alamos computing during World War II and the early Cold War. Fitzpatrick's work has been one of the few sources on Cold War Los Alamos for historians of computing, and was an important indicator to

Cold War Los Alamos are even more limited in their scope, either providing high-level overviews that reveal little of the Laboratory's internal operations, or providing a granular focus on the biographies of individual or groups of scientists—predominantly of physicists. Accounts falling within this range are best characterized as fragmented glimpses of a much larger and more complex institution. Peter Galison and Barton Bernstein's "In Any Light: Scientists and the Decision to Build the Superbomb, 1952-1954" is one of the best of these narrowly focused academic treatments, as it includes a thoughtful analysis of the scientific, personal, and political choices that contributed to thermonuclear-weapon development. It remains one of the most insightful treatments on the topic, possibly only surpassed with the completion of Anne Fitzpatrick's "Igniting the Light Elements." Galison and Bernstein do acknowledge computing among several important factors contributing to the "Super" project, but the limited timeframe of the

the HPC History of LANL Supercomputing Project that a rich collection of primary-source materials were available at LANL and that scholarly analysis could be done on Los Alamos computing history. Although, the dissertation's topic meant a focus on a limited ten-year timeframe, which includes the entirety of the Manhattan Project at Los Alamos. More discussions of the importance of Fitzpatrick's dissertation will follow in this chapter. Peter J. Westwick's *The National Labs: Science in an American System, 1947-1974* (Cambridge, MA: Harvard University Press, 2003) is one of the few attempts at doing an analytical history of the US nuclear-weapons Laboratory complex. Westwick's account focuses on the peculiar circumstances following World War II that fostered the development of a *collection* of laboratories tied to the nuclear complex, including the relative weakness of the Atomic Energy Commission as an organizing agency, which allowed the various laboratories to grow, and make themselves increasingly difficult to streamline or shut down between the 1950s and 60s. Far from accidental, the labs coordinated their research programs in order to ensure that each lab was indispensable for particular lines of research. *The National Labs* is a high-level overview of the entire AEC Laboratory system, focusing primarily on the late 1940s through the early 1960s. The 1960s through the mid-1970s earn only a brief summary in the concluding pages, likely due to the availability of materials. Robert W. Seidel, who worked at Los Alamos National Laboratory as director of the Bradbury Science museum, produced a number of articles providing overviews of the histories of the national labs, and their place as purveyors of "big science" and technological R&D in the United States since World War II. These pieces, such as "The National Laboratories of the Atomic Energy Commission in the Early Cold War," *Historical Studies in the Physical and Biological Sciences* 32 (2001): 145-162, are important sources, but, as brief overviews, do not explore any one lab in great detail over a substantial length of time. Seidel's full-length institutional histories are focused on Lawrence Berkeley Laboratory, rather than Los Alamos, due to factors that will be addressed later in the chapter.

article reveals only small pockets of Laboratory culture and history. The account still focuses primarily on the experiences of physicists, despite the diverse range of influences on the "Super."⁴ The narrow focus of timeframe, subject, or cultural depth found in most academic treatments of Los Alamos of any time period leaves sizeable gaps in the broader understanding of the Laboratory's change over time. This dissertation follows a middle-ground institutional approach meant to ameliorate gaps in the existing literature by situating a key area of Los Alamos research and development within the broader Lab environment over a longer period of time.

Popular treatments of Cold War and post-Cold War Los Alamos tend to skirt the fringes of the Laboratory, and range from the repackaging of officially released technical or organizational information into timelines and photobooks, to public-interest pieces meant to engage with common (or uncommon) perceptions of the Laboratory and townsite. One common form of public-interest piece examines the consequences (or perceived consequences) of LANL's history and presence on the surrounding region. Jo Ann Shroyer's *Secret Mesa: Inside Los Alamos National Laboratory* digs into LANL's complex relationship with northern New Mexico through a melding of news reports, anecdotes, and local and regional gossip.⁵ Shroyer's journalistic approach portrays LANL largely from a disaffected outsider's perspective, hoping to reveal to an audience of other outsiders the alleged inner workings or intentions of the organization, which comes off as monolithic and unchanging, likely due to minimal inside access. Shroyer's

⁴ Peter Galison and Barton J. Bernstein, "In Any Light: Scientists and the Decision to Build the Superbomb, 1952-1954," *Historical Studies in the Physical and Biological Sciences* 19, no. 2 (1989): 267-347.

⁵ Jo Ann Shroyer, *Secret Mesa: Inside Los Alamos National Laboratory* (New York: Wiley & Sons, Inc., 1998).

work, and other works similar in approach and analysis, tend to mystify more than they reveal about the institution, particularly with respect to the involvement of an institution like Los Alamos in key areas of Cold War scientific and technological change.

Conversely, Historian Jon Hunner's *Inventing Los Alamos: The Growth of an Atomic Community* is a more academically rigorous presentation of the community that developed around the Laboratory during its first fifteen years. Rather than further mystifying the institution, using government documents and personal accounts from residents, Hunner describes the development of the Los Alamos townsite and how that community contributed to the early growth and evolution of the Laboratory.⁶

Joseph Masco's *The Nuclear Borderlands* also provides perspective and analysis of the relationships that Shroyer's book sought to investigate, but through an anthropological and culturally intensive lens. Masco examines how legacies of the Manhattan Project have interacted with a variety of cultural groups in the Los Alamos region, including local pueblos, LANL weapons designers, and anti-nuclear groups. Masco analyzes how the Cold War's framing of the debate on nuclear weapons contributed to post-Cold War interpretations of the resulting nuclear-weapons complex.⁷ Masco's book remains the best anthropological study of Los Alamos County and LANL.⁸

⁶ Jon Hunner, *Inventing Los Alamos: The Growth of an Atomic Community* (Norman, OK: University of Oklahoma Press, 2004).

⁷ Joseph Masco, *The Nuclear Borderlands: The Manhattan Project in the Post-Cold War New Mexico* (Princeton University Press, 2006).

⁸ More direct anthropological studies of nuclear-weapons laboratories are possible, as Hugh Gusterson demonstrated in *Nuclear Rites: A Weapons Laboratory at the End of the Cold War* (Berkeley and Los Angeles: University of California Press, 1996). Although Gusterson did not receive direct support or cooperation from Lawrence Livermore National Laboratory (LLNL), he was able to interview and engage with LLNL scientists and weapons designers in order to understand to a greater degree the cultural processes, norms, and rites that had developed around their particular group and occupation. Gusterson also delved into the culture of Livermore, California, itself, analyzing the interconnected cultural developments that support the mission and existence of a weapons Laboratory.

Los Alamos National Laboratory (LANL), the Atomic Energy Commission (AEC), and the Department of Energy (DOE)⁹ have themselves commissioned, produced, and supported a variety of historical treatments on Los Alamos history. Most of those treatments have, again, focused on the Manhattan Project.¹⁰ Discussions and summaries of Cold War and post-Cold War Los Alamos history developed with LANL's support were largely cursory, public-relations-oriented pieces, rather than in-depth institutional or programmatic histories. Part of the overwhelming focus on the wartime period is due to its reciprocal relationship with the public fascination with the "secret city on the hill," relating to its involvement with the atomic bombings of Japan and the beginnings of the atomic era, romanticized stories of espionage, and the post-war popularization of leading scientific figures from the Manhattan Project.¹¹

The comparative abundance of primary- and secondary-source materials covering wartime Los Alamos, including oral histories, biographies, novels, memoirs, primary documents (unclassified and declassified), photographs, film footage, and other materials,

⁹ The Atomic Energy Commission (AEC) was the civilian commission formed by the US Congress in 1947 to oversee the former sites of the Manhattan Project. The AEC was split up in 1974, due to conflicts of interest in its charter, which specified that the AEC both promote and regulate nuclear energy. The Energy Research and Development Administration temporarily replaced the AEC until the Department of Energy was formed in 1977.

¹⁰ For example, the DOE and LANL provided support and access in the production of Lillian Hoddeson, Paul W. Henricksen, Roger A. Meade, and Catherine Westfall's *Critical Assembly: A Technical History of Los Alamos during the Oppenheimer Years, 1943-1945* (Cambridge, MA: Cambridge University Press, 1993), and Robert W. Seidel's *Los Alamos and the Development of the Atomic Bomb* (Los Alamos: Otowi Press, 1995).

¹¹ Richard Rhodes' *The Making of the Atomic Bomb* has long been the standard in popular accounts of Manhattan Project Los Alamos, offering a compelling, if slightly dramatized account of the human strife and challenges involved with living and working in the secret city on the Los Alamos mesa. As with most accounts of wartime Los Alamos, Rhodes' story at Los Alamos effectively concludes with the successful Trinity test in July of 1945; It should be noted that the Los Alamos Historical Society has published a small selection of pieces (and conducted oral-history interviews) describing the Los Alamos community during the Cold War.

do not exist by accident.¹² Historians, journalists, Manhattan Project veterans, and others, have worked with and under a variety of sponsoring organizations to document various aspects of Los Alamos and the Manhattan Project in relatively great detail for public consumption. Sponsors of this research have included the Atomic Energy Commission (AEC), the Department of Energy (DOE), the Atomic Heritage Foundation, the American Institute of Physics, the University of California, LANL, and various academic and popular-press publishers. LANL, the AEC, and the DOE themselves found expedient uses for the resulting wartime-focused historical treatments, such as leveraging the narrative of an American scientific and wartime success story as public-relations and recruiting tools during the Cold War, even as the public perceptions of atomic energy, nuclear weapons, and the Cold War in general, grew more complicated and contentious in the 1960s through the 1980s. For example, LANL celebrates its milestone years, most recently the 70th in 2013 and the 75th in 2018, with a variety of talks and articles focused on its World War II past. The presentations chosen to represent the Laboratory's history read as historical artifacts themselves, due to the focus on American scientific exceptionalism, and often outmoded historical and cultural interpretations of the use of atomic weapons during World War II. Publishers, and even filmmakers, have likewise capitalized on, and contributed to, the public fascination with Los Alamos and the Manhattan Project, sometimes with fictionalized depictions of life on the Hill.¹³

¹² It should be noted that oral histories of Cold War Los Alamos are among the most important primary-source materials available describing that period (typically describing the early post-war years), but they remain rare compared to World War II oral histories. The work of historians like Janet Abbate, who interviewed several women who worked in early-Cold War Los Alamos computing, must be acknowledged for capturing pieces of Los Alamos history that otherwise would have been lost.

¹³ The 1952 film "The Atomic City" is perhaps the earliest and most glaring example of entertainers taking advantage of the American public's fascination with Los Alamos, with a fictionalized account of terrorists

The abundance of research and documentary materials produced in the decades since Los Alamos and the Manhattan Project were officially revealed to the American public in 1945 have provided a long-term resource for historians and authors. The level of historical research and access that occurred with Manhattan Project history was more difficult to achieve for Cold War era Los Alamos, when Los Alamos Scientific Laboratory (LASL)¹⁴ and the Atomic Energy Commission's associated research and weapons complex grew into an elaborate secrecy-shrouded technological system.¹⁵

From a security standpoint, the highly classified nature of Los Alamos' Cold War mission made in-depth access extremely difficult for outsiders hoping to study the Laboratory's history, and curtailed what could be produced internally and made available for public consumption. Unlike conventional military classification procedures, which require justification for marking materials as classified, AEC/DOE materials relating directly to the development of nuclear weapons are "born classified" (due to the Atomic

ransoming a Los Alamos scientist's son in exchange for H-bomb plans. More recently, the 2014-2015 WGN series "Manhattan" depicted the interpersonal drama behind the development of the first atomic bombs, using a realistic reproduction of the townsite and Laboratory, but with mostly fictional characters, with the exceptions of J. Robert and Kitty Oppenheimer. A variety of historical-fiction novels and books for children and young adults also take advantage of wartime Los Alamos as a setting for dramatic storytelling. An annual performance of Tchaikovsky's "The Nutcracker" at Los Alamos High School sets the action in World War II Los Alamos, complete with LANL's director assuming the role of his predecessor, J. Robert Oppenheimer. The villain of the piece is a spy who steals plans from Oppenheimer's pocket during a holiday dance at Fuller Lodge.

¹⁴ The Laboratory at Los Alamos earned the Los Alamos Scientific Laboratory (LASL) designation with the formation of the Atomic Energy Commission in 1947. The former sites of the Manhattan Project that remained in operation after World War II, including Los Alamos, fell under the AEC's purview, and formed the core of the US nuclear-weapons Laboratory system. LASL gave way to Los Alamos National Laboratory (LANL) in 1980, after sufficient access to Laboratory resources were allocated to outside researchers to earn Los Alamos its "National" Laboratory accreditation.

¹⁵ In "Igniting the Light Elements," Anne Fitzpatrick argues that the US nuclear weapons complex falls within the definition of a "technological system" as described in Thomas P. Hughes' *Networks of Power: Electrification in Western Society, 1880-1930* (Baltimore, MD: John Hopkins University Press, 1983). Fitzpatrick contends that the secrecy surrounding the weapons complex has largely obscured its systemic nature from historians of technology.

Energy Acts of 1946 and 1954), and remain so unless and until they are deliberately declassified.¹⁶ Under this stricture, many documents and other materials dating as far back as the Manhattan Project remain classified even today. Those materials from the Manhattan Project that are available to the public were either originally unclassified, were created for public consumption, or were purposely declassified.

As the US nuclear-weapons laboratories during the Cold War grew in size, funding, and research scope, the volume of classified materials they generated grew at an exceptional rate. As a result, the laboratories were procedurally and culturally organized to prevent the outward flow of information that could potentially aid in the unauthorized external development of nuclear weapons. Consequently, Laboratory personnel involved with weapons development were usually unable to discuss their work in significant detail outside their closed professional circles in the AEC/DOE weapons program, making them unable (or at least understandably reticent) to document publicly or be interviewed about most details involving the Laboratory's primary mission. That environment naturally inhibits the construction of in-depth histories for an organization as complex as a multi-billion-dollar national Laboratory. Of the history-related publications that Los Alamos personnel themselves have produced or contributed to during and after the Cold War, most were typically in the form of brief memoirs or special-interest pieces on specific projects or lines of research.¹⁷

¹⁶ Fitzpatrick, "Igniting the Light Elements," 2; Robert W. Seidel, "A Home for Big Science: The Atomic Energy Commission's Laboratory System," *Historical Studies in the Physical and Biological Sciences* 16, no. 1 (1986): 135-175.

¹⁷ It was not impossible to research institutional or programmatic histories of Los Alamos during the Cold War, as institutional histories of Sandia and the AEC (to be discussed later) were not impossible, but access and institutional support for history research activities were difficult to obtain for researchers without pre-existing clearance and justification for their work. The HPC History of LANL Supercomputing Project

Behind Tall Fences: Stories and Experiences About Los Alamos at Its Beginning, an edited volume from the Los Alamos Historical Society, is emblematic of these accounts. It consists of twenty-one short articles by eighteen authors sharing their own accounts of what it was like to live and work at Los Alamos, largely focusing on personal impressions of the Laboratory and town, and the impact on their lives.¹⁸ These accounts mostly provide snapshots of town and Laboratory culture, at least as perceived through personal retrospectives. *Behind Tall Fences* is also consistent with the dominance of Manhattan-Project Los Alamos on the historical narrative, with rare forays into the Cold War experience.¹⁹

remains unusual at Los Alamos for the level of access and institutional support it has received as a history-focused project geared for disseminating LANL's HPC history; C. Bruce Tarter's *The American Lab: An Insider's History of Lawrence Livermore National Laboratory* (Baltimore, MA: John Hopkins University Press, 2018) is a new addition to the insider literature about the history of the US nuclear labs. A former LLNL director, Tarter offers an overview of Livermore's history, highlighting what Tarter considers key scientific and technical achievements at the Lab. It also continues the tendency of historical works discussing the US weapons labs to come from insiders who focus on their own fields within the lab system (although, as a former director, Tarter's purview was broader than that of most insiders), rather than from historians.

¹⁸ John C. Allred, et al., *Behind Tall Fences: Stories and Experiences About Los Alamos at Its Beginning* (Los Alamos, New Mexico: Los Alamos Historical Society, 1996). *Behind Tall Fences* is probably one of the more informative pieces fitting within the Los Alamos memoir genre. Lawrence Badash, Joseph O. Hirschfelder, and Herbert P. Broida, eds., *Reminiscences of Los Alamos: 1943-45* (Boston: USA, London: England, 1980), for example, characterizes the bulk of the genre, which presents new and repackaged memoirs of wartime Los Alamos. Discussions of Cold War Los Alamos were usually created specifically for public consumption, and reveal little beyond high-level overviews.

¹⁹ A significant Cold War and post-Cold War era form of published outreach was the *Los Alamos Science* journal, in print from 1980 to 2006. *Los Alamos Science* was a quarterly publication produced at LANL for public distribution. It hosted brief descriptions of various researchers, projects, and results from work done at Los Alamos, or in conjunction with academic or commercial partners. The purpose of the journal was to demystify among the American public the work done at LANL (although the articles were carefully chosen from unclassified projects, or unclassified results connected with more sensitive projects.), in order to draw attention to Los Alamos, while also downplaying the Laboratory's weapons work. The journal occasionally featured snapshots of Los Alamos history, which were usually collected recollections from what were then current or retired Laboratory staff. These snapshots, while useful in the absence of more in-depth source material, were typically brief anecdotes meant for entertainment purposes, but with important exceptions. Mentions of Laboratory computing history in the journal were largely confined to Manhattan-Project recollections, with the most common exceptions being pieces about the original MANIAC computer, and the related development of the Monte Carlo method of simulation. Nicholas Metropolis co-authored several articles about the MANIAC and Monte Carlo, which appeared in both *Los Alamos Science* and

Despite security and secrecy creating sizable restrictions to access at Los Alamos, and in the production of historical treatments of the Cold War Laboratory, other AEC/DOE facilities have produced, or aided in the production of, their own institutional histories of various scales and depths of analysis. While academically driven institutional histories of national laboratories are difficult to assemble, in light of the size, complexity, and security issues of each facility, those impediments can be overcome. For example, Sandia National Laboratories (SNL), located in Albuquerque, New Mexico and Livermore, California, has sponsored two institutional histories, including Nicah Stewart Furman's *Sandia National Laboratories: The Postwar Decade*, and Leland Johnson's *Sandia National Laboratories: A History of Exceptional Service in the National Interest*.²⁰ Part of Sandia's interest in its own history relates to its relative obscurity within the DOE complex. Formed in 1947 out of Z Division from Los Alamos, the division that constructed the atomic weapons, Sandia has remained largely out of the American public's eye compared to the more attention-grabbing weapons-design labs of LANL and LLNL, or the reactor-centric labs of Argonne and Oak Ridge. Although a

outside journals. The importance of Nicholas Metropolis and his efforts to popularize key pieces of Los Alamos computing history will appear in greater detail later in this chapter.

²⁰ Stewart Furman's *Sandia National Laboratories: The Postwar Decade* (Albuquerque, NM: University of New Mexico Press, 1990) is focused on the decade following Sandia's official formation in 1947, when the facility evolved from a collection of temporary structures next to an airfield, into a permanent weapons-production operation. The book is heavily focused on the development and change over time of the Laboratory's unique culture within the AEC system, resulting from its divergence from Los Alamos, its almost "junior partner" status compared to Los Alamos and Livermore, and its greater focus on practical engineering over the more science and theory orientations at the two design labs; Leland Johnson's *Sandia National Laboratories: A History of Exceptional Service in the National Interest* (Albuquerque, NM: Sandia National Laboratories, 1997) has, as its title indicates, a noted bias in its treatment of Sandia's history, a history that has brought SNL into conflict, or at least contention, with Albuquerque and northern New Mexico on more than one occasion over the facility's existence. Regardless, the book offers a surprisingly detailed look at the variety of programs and devices SNL has developed since the 1940s, considering that SNL is often regarded as one of the more security-conscious (some might say zealous) facilities in the DOE complex.

multi-billion-dollar-per-year operation, Sandia's range of research and production work is narrower than most of the other DOE labs. This means that Sandia's institutional histories, although still significant, are less difficult to collect and analyze—assuming the support and initiative of a leadership interested in producing such histories—as a single piece than those of DOE facilities with broader research envelopes, while also serving as important tools to draw attention to Sandia's sometimes overlooked position and importance within the DOE complex.

Much like the Sandia-supported institutional histories, the Atomic Energy Commission and the Department of Energy directly supported the research and writing of a three-volume institutional history of the AEC between the creation of the Manhattan Project and the early 1960s. Published between 1962 and 1989, *The New World: A History of the Atomic Energy Commission, 1939-1946*; *Atomic Shield: A History of the Atomic Energy Commission, 1947-1952*; and *Atoms for Peace and War, 1953-1961: Eisenhower and the Atomic Energy Commission*, effectively document the interrelated political and technical drives that underpinned much of the AEC's first decades of operation.²¹ These histories benefited from access granted to Richard G. Hewlett and his three co-authors who contributed to the series.

The AEC series further indicates that academically oriented, relatively in-depth histories are feasible within the US nuclear complex, but only when adequate access,

²¹ Richard G. Hewlett and Oscar E. Anderson, Jr., *The New World: A History of the Atomic Energy Commission, Volume I, 1939-1946* (University Park: Pennsylvania State University Press, 1962); Richard G. Hewlett and Francis Duncan, *Atomic Shield: A History of the United States Atomic Energy Commission, Volume II, 1947-1952* (U.S. Atomic Energy Commission, 1972); Richard G. Hewlett and Jack M. Holl, *Atoms for Peace and War, 1953-1961: Eisenhower and the Atomic Energy Commission* (Los Angeles: University of California Press, 1989).

initiative, and support are available from the organizations being studied. Such initiative and support has largely been absent at the higher levels of LANL with regards to its Cold War and post-Cold War histories, meaning that individual divisions and groups at the Laboratory, such as the HPC Division, have been largely left to instigate and find support for the research and production of their own history projects. Such projects, as the HPC History of LANL Supercomputing Project demonstrates, are capable of operating and conducting meaningful research within the LANL environment, but require the support of LANL staff with sufficient interest and resources to overcome significant funding, security, and procedural hurdles. These obstacles have generally been sufficient to make the HPC History of LANL Supercomputing Project effectively unique at LANL, and highlight the difficulty of attempting such research at a national Laboratory as an outsider.

Other factors also contributed to the continued dominance of Manhattan Project Los Alamos in the public consciousness. One factor was the accessibility, attention, and use of the historical sites themselves. World War II landmarks in Los Alamos County have remained a popular destination for "Atomic Tourists" since the townsite was opened to the public in 1957.²² Although the use of Manhattan Project-era sites for promotional purposes was not new at Los Alamos in 1957, as the former Ranch School Ice House, an important fixture during the Manhattan Project, was turned into the first museum in Los

²² Los Alamos county, the smallest county in New Mexico by area, consists of the Los Alamos townsite, the properties containing LANL, and the bedroom community of White Rock. The Los Alamos townsite and White Rock are administered by the county, as there is no town government.

Alamos.²³ Despite decades of change, public interest in the project has not ended. In 2016, the US National Park Service established the Manhattan Project National Historical Park. One of the park's three sites (the others being in Hanford, Washington and Oak Ridge, Tennessee) is headquartered at Ashley Pond, the center of the original Los Alamos Laboratory technical area.

As of 2018, the Park Service is negotiating access to three Manhattan-Project sites that are within LANL's current properties, but the majority of Cold War and post-Cold War sites will remain off-limits to the public. LANL and Los Alamos County have long capitalized on their conjoined wartime history as part of their respective core identities. LANL utilizes its Manhattan-Project renown to distinguish itself from the other AEC/DOE facilities (particularly from its chief rival, Lawrence Livermore National Laboratory), and the county uses its holdings (such as Fuller Lodge and two houses on Bathtub Row) to sell itself to tourists as "The Atomic City." The Los Alamos Historical Society has, in recent years, expanded its focus to document and present the county's Cold War legacy, which includes Cold War exhibits in the local History Museum.

The greater attention to World War II-era Los Alamos history also applies to the history of Los Alamos computing. Although Los Alamos was among the earliest centers of electronic digital computing—and has remained a consistent presence in the scientific-computing industry and community—the current understanding of Los Alamos Cold War

²³ The Ice House was a rock structure built to store ice harvested from Ashley Pond before the US Army seized the Los Alamos plateau in 1943. The structure was used during World War II to store bomb components leading up to the Trinity test in July of 1945, and was subsequently used as a museum of Laboratory wartime artifacts for visiting dignitaries. Its use was partly due to the security of its locking metal door, and partly for the historical value of the structure itself, as it helped to underscore for visitors the wartime need to use whatever resources were available during the crash program on the Hill.

computing development remains fragmentary at best—even within LANL. Manhattan Project computing at Los Alamos has received the majority of historical attention, despite the Laboratory's four decades of Cold War computing experience. Even then, Manhattan Project *computing* rarely appears as a focus. Computing features for the sake of adding color or an interesting anecdote to the traditional narratives about physicists leading up to the Trinity test.

For example, the first hand-computing group at Los Alamos, formed in the Spring of 1943, is typically described as being made up of the "wives of scientists," with little more said about that group's work, or how the group's size and composition changed markedly in only a matter of weeks.²⁴ The IBM Punched-Card Accounting Machinery, or PCAM, used at the Laboratory, operated primarily by soldiers from the Army Special Engineering Detachment (SED) starting in 1944, typically receives more attention and credit than the human computers. When hand computing is mentioned, such as in Richard Feynman's "Los Alamos From Below," it is in a race that Feynman staged between the human computers and the PCAM equipment, which resulted in an apparent victory for the machines.²⁵ Manhattan Project computing, and particularly human

²⁴ The phrase "wives of scientists" captured a lot of the chauvinistic attitudes of the 1940s and later periods with regards to human computers, and specifically in their role in the Manhattan Project. It aided in diminishing the difficulty and importance of the hand-computing effort by castigating it as "women's work," using attitudes of the period to imply that it was less than skilled work, or less than a critical component of the Manhattan Project, which perpetuated the dominant focus on the work and role of physicists.

²⁵ Richard P. Feynman, "Los Alamos From Below: Reminiscences of 1943-1945," *Engineering and Science* 39, no. 2 (1976): 11-30. The tale, which is the most often repeated about early Los Alamos computing, was meant to be amusing, with the human computers falling behind the machines due to fatigue. This contributed to a growing narrative at Los Alamos and elsewhere that human computing was, at best, antiquated, and, at worst, an aberration that was not true computing. It aided in reinforcing among scientists and engineers the perception in the 1960s and later that the dominance of electronic digital computing was inevitable, just as it was "inevitable" that electromechanical computers replaced human computers in the Manhattan Project, despite human computing groups remaining in use at Los Alamos and

computing, is usually treated as a side note meant to highlight the ad-hoc nature of the early work at Los Alamos, and as an amusing "See how far we've come!" tidbit for audiences. Presentations such as those may, even inadvertently, have the effect of encouraging audiences to dismiss the difficulties and importance of the early computing work, and its consequences in shaping Los Alamos' culture and mission in the decades that followed.

While the heightened security concerns that impacted the development of Cold War histories of Los Alamos were also at play with the Laboratory's long and elaborate involvement with computing, other factors local to LANL aided in relegating computing to the sidelines. Despite most of Los Alamos' non-weapons-code-related computing activities being unclassified in nature, the dominance of the weapons program, and its associated physics-centric culture, has fostered within the LANL community a perception and presentation of computing as a background support service, rather than as a field of research in its own right, particularly one that was and remains critical to LANL's mission.²⁶ Although the Laboratory's diversity of research divisions and activities has grown markedly since the late 1940s to include various branches of chemistry, biology, climatology, and many other fields, nuclear physics was and remains the primary source

elsewhere into the Mid-1950s. This problem of "inevitability" is elaborated in James S. Small's *The Analogue Alternative: The Electronic Analogue Computer in Britain and the USA, 1930-1975*, where the science and engineering community of the 1940s through at least the Mid-1950s did not consider the dominance of electronic digital computing over the other, contemporaneous forms of computing to be inevitable. However, later recollections often insist that the dominance of one over the other was obvious from the start, even though the staunch proponents of electronic digital computing were, themselves, the aberrations in the 1940s and early 50s. For further discussions of the "obviousness" of the dominance of electronic digital computers, see also Larry Owens, "Where Are We Going, Phil Morse? Changing Agendas and the Rhetoric of Obviousness in the Transformation of Computing at MIT, 1939-1957," *IEEE Annals of the History of Computing* 18, no. 4 (1996), which will be addressed later in the chapter.

²⁶ Weapons "codes" are the algorithms that aid in the design of nuclear weapons and their components, and simulate the behaviors of those designs, either in conjunction with live nuclear testing, or independently.

of funding and user of computing resources at Los Alamos, and dominates the internal and external presentations of the Laboratory's history.

The perception within the general LANL population of computing as an ancillary service (a perception that was strongest during the Cold War), despite the early and consistent pattern of scientific-computing R&D the Laboratory supported internally and elsewhere, shaped the organizational support and drive for investigating and disseminating histories of Laboratory computing. As a consequence, the internal focus on nuclear physics as the prime mover at LANL, coupled with the decades-long focus on the Manhattan Project as the historical moment that defined the essence of LANL's post-war character and mission, have combined with the issues shaping academic and popular presentations of Los Alamos to produce an especially fragmentary history of Los Alamos computing.²⁷

²⁷ Oral-history interviews have been a revealing source of information about the prevailing, and changing attitudes at LANL with regards to computing. Although the majority of LANL's funding as of 2018 remains directed at its traditional weapons-related work, computing as a critical resource, and as a source of institutional pride, has garnered broader attention and acceptance within LANL, particularly compared to the Cold War era. The changes in attitudes reflect a variety of cultural changes, such as the aging of the weapons-design staff due to the smaller and smaller numbers of nuclear physicists graduating from universities and joining LANL, and the continued influx of computing and non-weapons-related staff members since the 1990s. Computing, and those who work in the field, have also become more culturally relevant and significant at LANL, just as they have in the US and internationally, with the rise of personal-computing and social-media platforms and services. Graduates in the fields of computer hardware and software engineering who are drawn to the information-technology sector, like social-media giants Google and Facebook, or other high-tech startups, have transferable skill sets to and from places like LANL. At the same time, careers in nuclear physics appear to incoming college and graduate school students as increasingly old-fashioned and lacking in mobility. Much as Los Alamos in the 1950s used its state-of-the-art computers to attract new staff members who wished to work with the latest technologies, LANL today capitalizes on its high-end cluster computers and related research problems to draw new computing expertise and talent. A consequence of this juxtaposition—the increasing draw and relevance of computing versus the declining importance of nuclear-weapons and physics—is a growing prominence of scientific computing in LANL's outward presentations of itself, and in the conceptualization of the institution of computing within LANL.

Despite these ongoing challenges, the prominence and visibility of computing has begun to change within the LANL community since the end of the Cold War, particularly with the end of full-scale nuclear testing. The shift to computer-based simulations as the primary means of stockpile verification since the mid-1990s has brought computing more prominently into the foreground within LANL and in LANL's outward presentations of itself to the world. That is, in part, because the development and use of world-class supercomputer systems, like LANL's Roadrunner and Trinity clusters, are attention-grabbing, unclassified (in hardware and some of their software), and less controversial fixtures of the Laboratory's mission that can be presented publicly. They also represent important tools for promoting recruitment and partnership, as Los Alamos becomes increasingly reliant upon computing personnel and R&D in the pursuit of its primary mission.

Collaborations with supercomputer vendors on the development of new systems and technologies are not new arrangements for LANL or the DOE complex, as Chapter 2 will demonstrate. Although, the scale, complexity, and price tag of high-end cluster computers has made collaborations with a diverse range of vendors and other partners essential. As a consequence, supercomputing contracts, products, and systems have become prominent areas of interface and promotion between LANL and the outside world. As its scientific computing efforts have moved closer to the forefront of LANL's internal and external image, a greater organizational interest has emerged for the study and presentation of Los Alamos computing history. Key figures have publicly promoted Los Alamos computing history ever since the 1970s. Their efforts to preserve the

Laboratory's involvement with early electronic digital computing began even as that history was taking place. Their public efforts produced most of the materials available to historians about Los Alamos computing history, while their document collection efforts formed the core of LANL's extensive computing archive.

The existing literature on Los Alamos computing history began with those who first shaped that history. Early practitioners of a professional or technical field are typically the first to produce histories of that field. As a consequence, many histories of and by "pioneers" are heavily internalist and technical in orientation, as the chroniclers are usually not historians themselves but have chosen to document the rise of their fields, and desire to see their own work recognized. The record of Los Alamos computing is no exception.

Los Alamos computing began during a period when the sciences and mathematics were encountering the limits in existing computational methodologies and systems. Computing between the 1930s and 1950s underwent something akin to a Kuhnian "revolution," in that a variety of technical and methodological solutions to a common problem entered the fray, with each offering relative strengths and weaknesses over the others. Manhattan Project-era Los Alamos was among the sites where multiple approaches to computing came into contact, and quite literally competed against one another. There was no clear "winner" among the available technologies as World War II came to an end and the Cold War began, and Los Alamos employed examples of each technological path (including human, analog, electromechanical, and electronic digital computing) as the various options matured. How Los Alamos contributed to, and

interacted with, the development of computing during this time of technological uncertainty occupies much of the early history of Los Alamos computing, and was most of the focus for the individuals who documented that history. Although, key individuals at Los Alamos who were most invested, personally and professionally, in the growth of one particular line of computing development, made a concerted effort to promote that technology over its alternatives in their recollections and assertions about the past.

Nicholas Metropolis, a name now synonymous with computing at LANL, was the first promoter of Los Alamos computing history, and an early promoter of electronic digital computing in general. Oppenheimer recruited Metropolis in 1942 from the University of Chicago, where he had worked with Edward Teller. It was Teller who persuaded Metropolis to transition from experimental to theoretical physics, and Oppenheimer offered Metropolis the opportunity to continue his collaboration with Teller at Los Alamos. Metropolis expended no small effort in documenting and disseminating the role he perceived himself as playing in computing development—and the importance he ascribed to electronic digital computing in the prosecution of Los Alamos' mission. In wartime T Division, Metropolis was first assigned to develop equations for materials in extreme states of temperature and pressure. That assignment brought him into close contact with T Division's computing operations. After the war, the new field of electronic digital computing offered Metropolis a place to follow his new interest in mathematics and computing, and to earn recognition. Metropolis conscientiously curated his image as a computing pioneer, even while the field of digital computing was still taking shape, and he took a keen interest in ensuring that the early history of computing,

particularly at physics-centric Los Alamos, gained wide recognition when his technical career began to wind down in the 1970s and 80s. As a consequence, Metropolis and his co-authors produced a large proportion of the pieces available on Los Alamos' computing history, which were largely focused on the technical developments of computing at Los Alamos during its first decade, the time when Metropolis had the most direct influence on the shaping of Laboratory computing.²⁸

As early as the 1940s, Metropolis began to publish technical pieces describing computing R&D done at Los Alamos. Among the first and most significant, "The Monte Carlo Method" described Metropolis and Stan Ulam's adaptation of the Monte Carlo method of modeling complex physical phenomena for use on digital computers, and first appeared in September 1949.²⁹ A follow-on 1953 piece, which Metropolis co-authored with four other seminal figures in Los Alamos science, elaborated on the utility of digital computers for Monte-Carlo-based simulations, based on Los Alamos' experiences with the newly completed MANIAC computer.³⁰ These early technical articles were tied to Metropolis' evangelism of electronic digital computing by encouraging the search and development of applications for the new technology, which he and other early figures in Los Alamos computing, such as Jim Richardson and Eldred Nelson, considered critical to the Laboratory's Cold War mission. Metropolis continued to promote Los Alamos within the small, but growing computing community as a leader in scientific-computing R&D in

²⁸ Herbert L. Anderson, "Metropolis, Monte Carlo, and the MANIAC," *Los Alamos Science* (Fall, 1986): 96-107.

²⁹ Nicholas Metropolis and Stanislaw Ulam, "The Monte Carlo Method," *Journal of the American Statistical Association* 44, no. 247 (September, 1949): 335-341.

³⁰ Nicholas Metropolis, Adrianna W. Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller, and Edward Teller, "Equations of State Calculations by Fast Computing Machines," *Journal of Chemical Physics* 21, no. 6 (1953): 1087-1092.

the US, and himself as an influential figurehead of the Los Alamos computing effort, over the 1950s and 60s. This effort continued even as Metropolis returned to the University of Chicago to build the MANIAC III in the early 1960s, then back to Los Alamos in a largely advisory capacity in the late 1960s. In the 1970s, Metropolis used his presence at Los Alamos and within the computing community to play an active role in the emerging effort to document the history of electronic computing, hosting at Los Alamos the International Research Conference on the History of Computing in 1976.

The conference was a foundational moment in establishing the history of computing as a field of study, with a large number of computing "pioneers" contributing talks about their work (the audio of which was recorded and preserved at Los Alamos). A truly international affair, attendees included Soviet computing pioneer Andrei Ershov, the controversial figure from Germany, Konrad Zuse, and well-known representatives from the US, Great Britain, Japan, the Netherlands, Czechoslovakia, and others. The talks they delivered, including three by current or former Los Alamos staff members, Metropolis, Stan Ulam, and Mark Wells, formed the basis of the conference publication, *A History of Computing in the Twentieth Century: A Collection of Essays*.³¹ Metropolis' contribution to the collected volume, "The MANIAC," outlined the development and use of Los Alamos' first electronic digital computer, and began a string of computing-history publications for Metropolis between 1980 and the early 1990s. Most of Metropolis' history articles focused on the early digital computing efforts and contributions of Los Alamos, primarily with the Monte Carlo method and the first MANIAC.

³¹ Nicholas Metropolis, J. Howlett, and Gian Carlo Rota, eds., *A History of Computing in the Twentieth Century: A Collection of Essays* (New York: Academic Press, 1980).

The primacy in Metropolis' accounts of electronic digital computing over traditional hand computing and electromechanical computing echo similar emphases on the "obviousness" of electronic systems becoming the dominant technology found in other, insider histories of computing.³² Metropolis' focus reflected, and contributed to, the overall emphasis in the history of computing on electronic digital systems, and in the same emphasis within Los Alamos, where the record of its extensive use of analog, human, and electromechanical computing after World War II has effectively gone unexplored. The major Metropolis articles, about six altogether, represented most of the available literature based on internal access about Los Alamos computing history until the late 1990s, with the publication of Anne Fitzpatrick's "Igniting the Light Elements" (discussed further later in the chapter), which provided excellent discussions of the various forms of computing used at Los Alamos between World War II and the early 1950s.

Metropolis' final history-related article, "The Age of Computing: A Personal Memoir (1992)," further discusses a comment Metropolis and Jack Worlton, a Los Alamos computing and procurement expert between the 1950s and 1980s, made in a 1972 conference talk about mistakes and misrepresentations in computing history.³³ In the 1992 article, Metropolis outlined what he considered to be significant detriments to the effective study of the history of computing, namely, the primacy of theoretical

³² See Owens, "Where Are We Going, Phil Morse?"

³³ Nicholas Metropolis, "The Age of Computing: A Personal Memoir," *Daedalus* 121, no. 1 (Winter, 1992): 119-130; N. Metropolis and J. Worlton, "A Trilogy on Errors in the History of Computing," in *First USA-Japan Computer Conference Proceedings: October 3-5, 1972, Tokyo, Japan* (American Federation of Information Processing Societies, 1972): 683-691.

physics in the history of science. He argued that the predominance of theoretical physics as the intellectual driving force in twentieth-century scientific and technological change served to obfuscate and even denigrate the history of engineering, particularly as it related to the development of computing, even among computer scientists.

In the context of his 1992 article, the body of Metropolis' work presenting Los Alamos computing history appears to put forward a two-part agenda. The first part was to establish and nurture the legitimacy of the history of computing as a field of study, just as he evangelized electronic digital computing in the 1940s and 50s as a preferred technological path. The second goal was to place Los Alamos at or near the center of what he considered the beginnings of scientific computing in the twentieth century. It was not unusual among early figures in computing history to draw attention to the history of their field, as the numerous "pioneers" taking part in the 1976 Los Alamos history of computing conference amply demonstrated. For Metropolis, emphasizing the legitimacy and importance of computing history was particularly pertinent in the physics-dominated culture within Los Alamos, as it was a key component for ensuring that his own legacy was preserved and recognized.³⁴

While Metropolis was not alone, his was the most significant computing-history effort and publication record at Los Alamos between the 1970s and early 1990s. Fragments of Cold War Los Alamos computing history appeared in various publications, such as the numerous histories of IBM and other computing corporations, and

³⁴ In 2002, LANL dedicated its new, 300,000-square-foot computing center as the Nicholas C. Metropolis Center for Modeling and Simulation, as testament to the presence Metropolis established in Laboratory computing history.

government- and military-sponsored computing development and projects.³⁵ Discussions of Los Alamos in these publications were largely based on materials and perspectives from outside the Laboratory, while also relying heavily on the brief historical treatments Metropolis and a few other Los Alamos computing figures, like Jack Worlton, Edward Voorhees, Mark Wells, and Roger Lazarus, had produced over the late-1970s and the 1980s.³⁶ It was clear to historians of computing that Los Alamos had a long-standing presence and involvement in scientific computing, but that history was mostly known only in reference to outside organizations and entities that had interacted with the Laboratory, with very little emerging from Los Alamos itself.

One of the first and most notable academic treatments on the history of Cold War computing at Los Alamos was "The Influence of the Los Alamos and Livermore National Laboratories on the Development of Supercomputing," a 1991 *IEEE Annals of the*

³⁵ For example, Los Alamos appears in a variety of computing-history-related works, from the largely technical Charles J. Bashe, Lyle R. Johnson, John H. Palmer, and Emerson W. Pugh, *IBM's Early Computers* (Cambridge, MA: MIT Press, 1986), the business-history-oriented Emerson W. Pugh, *Building IBM: Shaping an Industry and Its Technology* (Cambridge, MA: MIT Press, 1995), and the historiographically revolutionary Paul N. Edwards, *The Closed World: Computers and the Politics of Discourse in Cold War America* (Cambridge, MA: MIT Press, 1986). Los Alamos is typically a black box in most accounts in the history of computing, mostly as a place where machines disappear, rather than as a contributing participant in the scientific-computing field. As a consequence, the tales are usually one-sided, from the perspective of vendors or other computing organizations. This stands in contrast with the US Military, which was also a major source of R&D funding and contracts, like Los Alamos and the AEC/DOE, but the literature discussing military involvement in early digital computing development is quite extensive.

³⁶ For example, Roger B. Lazarus, Edward A. Voorhees, Mark B. Wells, and W. Jack Worlton's "Computing at LASL in the 1940s and 1950s," Los Alamos Scientific Laboratory, LA-6943-H (May 1978) was a collection of four technical overviews of Los Alamos computing, which were assembled for the 1977 National Computer Conference. The conference organizers featured LASL on the conference's "Pioneer Day," meant to honor early figures in digital computing. The four articles the LASL contingent presented were, again, largely discussions of wartime and early Cold War computing at the Laboratory, but with greater attention to software and Los Alamos' involvement with supercomputer development in the mid-1950s through the introduction of the Cray-1, but only in brief, cursory treatments. As basic as these articles are, they remain among the best insights into Los Alamos Cold War computing that have emerged from the Laboratory itself.

History of Computing article that sociologist and historian of technology Donald Mackenzie wrote, in part, with access to staff members of LANL and LLNL's computing operations.³⁷ Mackenzie's article raised basic questions about the influence of the US nuclear weapons design laboratories (LANL and LLNL) on the development of computing, questions that had long been raised and discussed regarding other computing R&D sponsors and centers, like the US military, IBM, and others, but were still almost entirely unexplored regarding the AEC/DOE laboratories. MacKenzie's technical information came largely from existing publications external to Los Alamos and Livermore, while internal information about the organizations, and about how machines, software, and other technologies were adapted for use at the laboratories, came almost entirely from interviews with computing personnel.

While historians of physics and computing appreciated that the computing operations at Los Alamos and Livermore were somehow connected with the broader missions of the laboratories, that is, the utility of computer-based simulations in nuclear-weapons development, relatively little was known about how nuclear-weapons development and scientific computing influenced one another, let alone how the interaction of the two changed over time. Those issues, and the various social, technical, and economic factors influencing them, were largely hidden within the secretive operations at either Laboratory. MacKenzie pursued those unknowns to make larger statements about the changing influence of the US weapons labs on computing between the 1940s and the 1990s. He concluded that, to a large degree, the influence of the

³⁷ Donald MacKenzie, "The Influence of the Los Alamos and Livermore National Laboratories on the Development of Supercomputing," *IEEE Annals of the History of Computing* 13 (1991): 179-201.

laboratories on supercomputing vendors had declined with the significant expansion of high-end scientific computing to a new range of customers in the 1970s and 80s, making the weapons laboratories less of a central focus for vendors than in the 1950s and 60s. MacKenzie also concluded that the inherent secrecy of Los Alamos and Livermore had constrained the overall technical influence of the laboratories on the broader supercomputing field, as their computing R&D efforts were largely known only to insiders. The weapons laboratories, MacKenzie continued, also focused on scientific computing, which had been the first prime mover in electronic digital computing development in the 1940s and 50s, but had since given way to data processing and consumer technologies between the 1970s and 90s as the foci of the computing market.

What was most significant about MacKenzie's article for historians of computing was how it situated computing at LANL and LLNL within their long-term Cold War context, a perspective that had been largely absent in previous publicly available treatments. Despite this, MacKenzie's interest in situating both Los Alamos and Livermore, and his still-limited access to internal documentation (for example, he had access to only one internally produced LANL document, a Computing-Division Review), constrained the depth and scope of his discussion about any particular Laboratory or time period. MacKenzie also conducted his research at a time at the end of the Cold War when the future of the weapons laboratories was unclear, and while the scientific-computing industry was in a state of rapid flux. As a consequence, MacKenzie's presentation bears much of the perception of uncertainty and loss that was characteristic of the US weapons labs at the end of the Cold War, as they contended with the end of the

superpower conflict that had justified their existence. The computing operations at the laboratories were also contending with an uncertain technical path, due to the diminishing performance returns of conventional single- and multi-processor supercomputer systems, and the realization that the weapons labs were no longer the central focus of the supercomputing field.

MacKenzie's analysis of the weapons laboratories at that moment, despite the uncertainties of the laboratories and their computing efforts in the late-Cold War environment, and the limited access to internal documentation, offered a never-before-seen perspective on the changing place of computing at Los Alamos over the Cold War. Although, the article was a product of its time and of MacKenzie's limited access (despite his access being unusual for an outsider), offering a Cold War and computing perspective in the midst of both undergoing significant changes, and while attempting to understand the meanings and consequences of the Cold War as it concluded. Analyses conducted after the Cold War and with extensive internal access do exist of computing development at US AEC/DOE laboratories. Anne Fitzpatrick's dissertation "Igniting the Light Elements" was the first, and its importance for analyzing early Cold War Los Alamos has already been discussed earlier in this chapter. While not specifically about computing, "Igniting the Light Elements" marked the first extensive use of Los Alamos' computing archival collections for an in-depth academic history publication. Completed in 1999, Fitzpatrick's dissertation placed Los Alamos computing of the 1940s and early 50s into a rich Cold War context, partly due to the newly available range of post-Cold War

scholarship and perspectives, and, more importantly, due to the author's extensive access to LANL's archives as an internal researcher.

Because of its unparalleled, in-depth survey of the "Super" development effort at Los Alamos, and the interconnected computing operations, based on otherwise inaccessible archival materials, "Igniting the Light Elements" has remained an important resource for historians of computing since its release. For example, Thomas Haigh, Mark Priestley, and Crispin Rope cite Fitzpatrick extensively in their "Los Alamos Bets on ENIAC" and *ENIAC in Action*, one of a series of articles and the subsequent book that delve into the life and times of what is widely (although contestedly) considered the first general-purpose programmable electronic computer.³⁸ Los Alamos played an important role in the use and modification of the ENIAC over its service life, but most of the documentation describing the Los Alamos side of that relationship remained unavailable to these researchers. Fitzpatrick offered the first glimpse at what was available in the Los Alamos archives, informing historians of computing, who otherwise would not have access of any kind, and reinforcing for the HPC History of LANL Supercomputing Project collaborators the value of pursuing archival access previously unattempted at LANL.

While Fitzpatrick's dissertation revealed a wealth of archival materials available at Los Alamos, Charles N. Yood's *Hybrid Zone: Computers and Science at Argonne National Laboratory, 1946-1992* demonstrated to HPC History of LANL

³⁸ Thomas Haigh, Mark Priestley, Crispin Rope, "Los Alamos Bets on ENIAC: Nuclear Monte Carlo Simulations, 1947-1948," *IEEE Annals of the History of Computing* 35, no. 3 (July-Sept. 2014): 42-63; Thomas Haigh, Mark Priestley, Crispin Rope, *ENIAC in Action: Making and Remaking the Modern Computer* (Cambridge, MA: The MIT Press, 2016).

Supercomputing Project collaborators the potential for, and the value of, investigating the broader history of computing at a national Laboratory.³⁹ Working under a DOE-sponsored fellowship to explore the DOE's contributions to the history of computing, Yood chose Argonne National Laboratory (ANL) as his focus because of its long involvement with computing, the fact that many individuals involved with that history were still available for interview, and the fact that Argonne did not work on weapons development, meaning that most of the documentation at ANL would not be classified.

Unfortunately, most of ANL's computing-related archival collections from before 1968 had been destroyed, which, combined with difficulties in obtaining certain DOE materials after the events of September 11, 2001, added to the complexity of Yood's research effort, but the result was the first-ever in-depth, academic analysis of the development of computing at a national Laboratory over the Cold War. Yood's treatment is a thorough examination of how computer science and the computational sciences interacted to shape the perception and use of computing at Argonne during the Cold War, a treatment that echoes MacKenzie's interest in LANL and LLNL in the early 1990s, but with a far deeper understanding of the long-term meanings and consequences of the Cold War, a more robust literature to draw upon from the history of computing, and more extensive access into the social and political structures behind the shape of computing at a national Laboratory.

The publication of *Hybrid Zone* in 2013, and its revealing account of the social construction of ANL's computing history, indicated to the HPC History of LANL

³⁹ Charles N. Yood, *Hybrid Zone: Computers and Science at Argonne National Laboratory, 1946-1992* (Boston, MA: Docent Press, 2013).

Supercomputing Project collaborators the possibility and potential of pursuing a similar academic analysis of Los Alamos computing history, utilizing oral histories similar to Yood's at Argonne, and archival source materials that Anne Fitzpatrick's dissertation demonstrated existed at LANL. Although obtaining initial permissions for archival access was complex, once secured, the computing collections themselves were far more complete than those at Argonne, and revealed a much broader scope of computing activities at Los Alamos, and of extensive interactions between Los Alamos and the scientific-computing community. The broader scope of activities and interactions were due to the more extensive use of, and need for, computing at Los Alamos than at the other AEC/DOE labs, aside from Livermore, which ran a similar, but far from identical computing operation, as Donald MacKenzie had noted.

In summary, the established literature on Los Alamos computing history reflects much of what has defined Los Alamos in both the public's perception, and in the Laboratory's internal perception and presentation of itself, since the end of World War II. The reciprocal relationship between public fascination with the wartime lab, and the Cold War Laboratory's desire for both secrecy and public recognition, resulted in an overrepresentation of Los Alamos' formative years, focusing primarily on physics and physicists as the prime movers of that history. Despite computing's long presence at Los Alamos, and the existence of important figures championing awareness of that presence, layers of secrecy obscured much of the Laboratory's computing history from public view and study, and, along with the internal perception at Los Alamos of the Laboratory as a

physics-driven facility, aided in obscuring the extensive computing effort that underpinned that facility's mission.

A history of computing use and development as extensive as that at Los Alamos would, in facilities less sensitive to security and more focused on public presentation, ordinarily be a more actively cultivated source of accolade. Instead, computing has traditionally occupied the margins of Los Alamos' outward presentations of itself, only earning greater prominence in LANL's outreach and self-promotion since the end of the Cold War. As a consequence, the available literature surrounding Los Alamos computing history has remained limited, until the development of the HPC History of LANL Supercomputing Project, and its unprecedented access to computing-related primary-source materials at LANL.

The following chapters represent the first in-depth, historical analysis of the development of computing at Los Alamos during the middle-decades of the Cold War, from the mid-1950s to 1980. This period marks the emergence of "supercomputing" as a category among scientific-computing vendors and users, when a small cross-section of users, such as Los Alamos, exerted considerable influence in creating and shaping that category. While historians of computing are generally aware of Los Alamos as an early center of electronic digital computing, and as an important player in the formation of supercomputing, the internal processes through which Los Alamos and its computing operations shaped and were shaped by the development of supercomputing have been largely concealed from historical survey. Archival and oral history information reveal a variety of factors stemming from Los Alamos' place in wartime and early Cold War

computing, and the changing importance and perceived roles of computing at the Laboratory, that influenced how the Laboratory approached computing, and how that approach changed with the continuation of the Cold War and the maturation of the computing industry.

Los Alamos' approach to computing—particularly differing visions of the place of computing at Los Alamos, and how the Laboratory's approach could or should change based on those differing visions—was a contentious issue at the Laboratory that culminated in a near-total restructuring of its computing operations. The remainder of this chapter establishes the existing scholarship and timeline of early Los Alamos computing, in the period between the founding of the Laboratory and the inception of "supercomputing." More importantly, the remainder of the chapter also provides the groundwork for what would come to characterize the Los Alamos approach to computing for nearly four decades. That approach became a controversial point of division between the Laboratory's computing specialists and its most influential (weapons designer) users, with escalating consequences for Los Alamos computing in the late 1970s.

III. The Origins of Los Alamos Computing

The origins of the Manhattan Project and Los Alamos are, as discussed at the beginning of this chapter, well-studied in both popular and academic writing. What follows in this chapter focuses on the development of computing specific to wartime and early Cold War Los Alamos. While this chapter offers sufficient context for the variety of influences on Los Alamos computing into the early 1950s, an in-depth discussion of

Manhattan Project and Los Alamos history is beyond the scope or need of this dissertation. For more extensive background information beyond computing, the reader is encouraged to consult the myriad of available published works. *The Making of the Atomic Bomb* by Richard Rhodes gives a compelling presentation of Manhattan Project and wartime Los Alamos history, and Peter Westwick's *The National Labs* or Robert Seidel's "A Home for Big Science" provides a more complete picture of how Los Alamos developed out of a secret government scientific project and evolved during the Cold War. Marjorie Bell Chambers' "Technically Sweet Los Alamos" similarly offers a detailed examination of how the townsite itself changed over its early decades.⁴⁰

Computing at Los Alamos began in the wartime spring of 1943, as the Laboratory itself grew atop a remote volcanic mesa in northern New Mexico. The Laboratory's lands were seized from local homesteaders and the Los Alamos Ranch School by the US War Department starting in 1942, with construction on the lab facilities beginning that December. Scientists and support staff began to arrive in late March 1943. Theoretical physicist J. Robert Oppenheimer, the Laboratory's wartime director, and Lieutenant General Leslie Groves, the US Army officer in charge of the Manhattan Project, selected the Pajarito Plateau as the project's primary R&D site for its safety and security benefits, and for Oppenheimer's aesthetic appreciation of the region. The idea behind creating a new, isolated Laboratory for the development of an atomic weapon, rather than expanding on an existing Manhattan Project research site, like the MetLab at the

⁴⁰ Marjorie Bell Chambers, "Technically Sweet Los Alamos: The Development of a Federally Sponsored Scientific Community," (PhD diss., University of New Mexico, 1974); Seidel, "A Home for Big Science: The Atomic Energy Commission's Laboratory System," 135-175.

University of Chicago, was to move the potentially dangerous work away from population centers, and to curtail potential leaking of top-secret information by isolating the personnel working on the project.⁴¹ While the remote location ultimately did not prevent the leaking of atom-bomb plans to the Soviets during World War II, it did play an important role in shaping wartime and post-war Los Alamos and its computing efforts.

Among the early arrivals at the Laboratory amid the flurry of construction on the Hill in March of 1943 were Stanley Frankel and Eldred Nelson, both physicists from the University of California, Berkeley. Oppenheimer, their former professor, recruited Frankel and Nelson to Los Alamos to continue Manhattan Project computing work they had begun at Berkeley. Working for the Theoretical or "T" Division and its leader, Hans Bethe, Frankel and Nelson assembled the first computing resources at Los Alamos. Because of its computationally demanding work in modeling proposed weapon designs and in estimating needed quantities of fissile materials for those weapons, T Division directed most of the computing capacity at Los Alamos during the Manhattan Project.⁴² T Division's most challenging computational wartime tasks included the neutron-diffusion problem of the gun-type weapon, what would become the so-called "Little Boy," the bomb dropped on Hiroshima, and shaping the explosive lenses of the implosion device, which would become the "Gadget" detonated near Alamogordo, New Mexico, the proof-of-concept for the finished "Fat Man" weapon, dropped on Nagasaki. T Division also undertook most of the computation work supporting next-generation thermonuclear-

⁴¹ Westwick, *The National Labs*, 30-32; Chambers, "Technically Sweet Los Alamos," 81-82.

⁴² Rhodes, *The Making of the Atomic Bomb*, 460-461; Anne Fitzpatrick, "Igniting the Light Elements," 51, 53.

weapons development between World War II and the early Cold War. By contrast, as of the spring and summer of 1943, T Division's computational capabilities were as basic as the new lab itself.

Frankel and Nelson had created a hand-computing group in support of Manhattan Project neutron-diffusion research at Berkeley in 1942, and ordered a contingent of mechanical desktop calculators to expedite T Division's continuation of that effort at Los Alamos. The purpose of that effort was to calculate, or at least estimate as accurately as was feasible given the project's emphasis on expedience, the critical mass of a uranium or plutonium core for a gun-type weapon, which was the primary research focus at Los Alamos in 1943. The wrong mass could detonate too early, reducing the weapon's power and wasting expensive fissile material, or not detonate at all. The series of lectures that underpinned the famous *Los Alamos Primer* began in April at Los Alamos, and aimed at bringing new arrivals up to speed on the multitude of challenges facing the design effort. Those challenges highlighted the need for the calculation effort to move as quickly as possible. The project's time constraints (the Manhattan Project's primary goal was to develop an atomic weapon before Nazi Germany), secrecy, and the cost of fissile materials, precluded extensive live testing of designs, making numerical or computed solutions essential.⁴³

Labor-saving computing technologies were critical at Los Alamos from the outset, due to the shortness of time and the complexity of the Laboratory's central mission, but the range of technologies and methodologies available to Los Alamos as it

⁴³ Rhodes, *The Making of the Atomic Bomb*, 461-463; Anne Fitzpatrick, "Igniting the Light Elements," 56.

began operations were limited. Even though the mechanical calculators Frankel and Nelson ordered upon arriving at Los Alamos were in high demand, the Manhattan Project's AAA-priority rating (the highest acquisitions priority during wartime) and the talent of the Los Alamos procurement officer, Dana Mitchell, secured the Lab a quantity of 8- and 10-digit desktop calculators that were distributed among project scientists. The 8-digit calculators were quickly abandoned, due to their slow speed and small digit capacity. The 10-digit Marchant and Friden machines were of greater utility, but the critical-mass calculations were stretching the limits of traditional hand computation, while still not providing the accuracy needed for the scientific and engineering staff to make design choices with the level of confidence they desired. Committing to design decisions without adequate numerical data to support those decisions could have led to technical dead ends down the road.⁴⁴

The computing effort escalated with the weapons research and design effort at Los Alamos. The population on the Hill (a common nickname for wartime Los Alamos, as those who knew of its existence were prohibited from calling it by its actual name) quickly ballooned beyond the 300 research and support personnel that Oppenheimer had initially estimated. Before the war's end, a total population of over 7,000 were living and working at Los Alamos. In 1942, Oppenheimer and other lab planners considered the fission-weapon problem to be "solved" compared to the thermonuclear-weapon concept

⁴⁴ N. Metropolis and E. C. Nelson, "Early Computing at Los Alamos," *Annals of the History of Computing* 4, no. 4 (October 1982): 348-349; Jon Hunner, *Inventing Los Alamos*, 19; Anne Fitzpatrick, "Igniting the Light Elements," 60-61; Lillian Hoddeson, et al. *Critical Assembly: A Technical History of Los Alamos during the Oppenheimer Years, 1943-1945* (Cambridge [England]: Cambridge University Press, 1993), 179-180; Francis H. Harlow and N. Metropolis, "Computing & Computers: Weapons Simulation Leads to the Computer Era," *Los Alamos Science* (Winter/Spring, 1983): 132-133.

that Edward Teller had proposed. As a consequence, they underestimated the complexity of the R&D undertaking, particularly in their need to push the existing boundaries of material physics and numerical methods. Due to personnel shortages and the volumes of calculations involved with early gun-type-weapon design work, Mary Frankel, Josephine Elliott, Mici Teller, and other spouses of T Division personnel formed the first organized computing group at Los Alamos, the T-5 Computations Group. Under its leader, the New York University mathematician Don "Moll" Flanders, T-5 grew to about 20 personnel (mostly women) by late summer of 1943 with the addition of human computers from the Women's Army Corps (WACs). Members of the US Army Special Engineering Detachment (SED) later joined the computing effort. The SED was created for the Manhattan Project, and consisted of graduate students in the sciences and engineering who were capable of assisting with research at Los Alamos.⁴⁵

By the end of 1943, the hand-calculation group and its desktop calculators could not keep pace with the time-sensitive neutron-diffusion calculations needed for the gun-type weapon. One impediment was that the calculators themselves broke down frequently due to their heavy use and the dusty environment on the Hill. Located on the top floor of a hastily-constructed two-story building between Ashley Pond⁴⁶ and the technical area's main thoroughfare (named "Trinity Drive" after the war), the T-5 group had little choice but to keep the windows open in the unairconditioned building during

⁴⁵ Metropolis and Nelson, "Early Computing at Los Alamos," 349; Personal communication with Los Alamos Historical Society Museum Educator, 2017; *Inventing Los Alamos*, 29; Harlow and Metropolis, "Computing & Computers," 132.

⁴⁶ The original, main technical area of wartime Los Alamos centered around Ashley Pond, named after the founder of the Los Alamos Ranch School, Ashley Pond. Current Los Alamos residents still sometimes refer to the water feature, jokingly, by its full name, "Ashley Pond Pond."

the warmer months. Frequent truck traffic introduced dust into the complex calculator mechanisms, accelerating wear and tear. T-Division theoretical physicists Richard Feynman and Nicholas Metropolis began an impromptu calculator repair service, due to the time wasted sending the machines off the Hill to repair shops in Santa Fe or Albuquerque. The lab administration, after initial qualms with Feynman and Metropolis' unsanctioned repair service, turned a blind eye to their effort as the backlog of broken machines began to grow.⁴⁷

Feynman often told and retold his account of repairing calculators as an example of his mirthful adventures at Los Alamos. Metropolis followed suit, using the repair effort as an amusing prologue when writing at the end of his career about his experiences with early Los Alamos computing. Despite the common presentation, the calculator repair effort and the strained limits of existing computing methodology under the extreme circumstances at Los Alamos were emblematic of T Division's emerging approach to computing. It was an approach that did not end with the war, but evolved into the division's normal operating procedure, and its overall interpretation of computing as a necessarily hands-on activity at Los Alamos. T Division's approach, which carried forward into the division that inherited its computing operations in the late 1960s, emphasized the need to maintain local computing expertise and the capacity for self-reliance, in response to the unusual computational pressures and security requirements facing lab scientists and engineers.

⁴⁷ Metropolis and Nelson, "Early Computing at Los Alamos," 349.

A related aspect of that emerging approach was the need for flexibility in the selection and modification of technologies and methodologies to suit Los Alamos' unusual situation. An early example of T Division's flexibility of methodologies was Feynman's division of computational tasks onto index cards, and assigning specific tasks in a problem, such as an addition, multiplication, or division to individual human computers. The reorganization, according to Feynman, resulted in a considerable speed-up of the computing operation. The concept of breaking down mathematical problems into individual steps for human computers was by no means new when Feynman developed his strategy at Los Alamos, but it emphasized the need and desire within T Division for flexibility in the handling of its computational tasks. This was due to the wartime and early Cold War urgency to mitigate the limitations of traditional computational methods, and the period of uncertainty surrounding the concurrent existence of several lines of computing development, new and long-established.⁴⁸

Although computing was often represented and described as simply a tool in service of scientific ends among many early figures in computing, including John von Neumann, for many of the T-Division personnel who directly shaped the Lab's computing efforts during and after the war, working on computing meant working on the front lines—the front lines of Los Alamos' mission, and the front lines of technologies capable of fundamentally altering how and where science could peer into the natural world.⁴⁹ T Division's approach, and how it evolved in relation to differing perceptions of

⁴⁸ David Alan Grier, *When Computers Were Human* (Princeton, NJ: Princeton University Press, 2005), 275.

⁴⁹ For further discussion of the connections between computing and nuclear-weapons development at Los Alamos during the early Cold War, see Peter Galison, "Computer Simulations and the Trading Zone," *The*

where computing fit within Los Alamos' mission and image, will appear frequently in later chapters, but the approach found its footing during the lab's Manhattan Project and early postwar years.

The attempt to adapt mechanical computational aids to Los Alamos' burgeoning wartime demands has been one of the most oft-cited examples among retirees and chroniclers of lab computing history of T Division's development into a center of computing expertise. Although, in many accounts, the mechanization of Los Alamos computing has typically been presented as a prelude to electronic digital computing, particularly to the creation of the MANIAC computer (1952), rather than as a formative moment for T Division's distinctive approach to computing. Nicholas Metropolis was among those who cast IBM's punched-card accounting machinery (PCAM) as a gateway technology to electronic digital machines. He also portrayed the initial adaptation of the IBM PCAM equipment to Los Alamos' demanding and unusual simulation calculations as a key moment that demonstrated to T Division the potential of new computing technologies. The IBM PCAM systems also served as an alert to Metropolis and others in T Division that, when operating at the fringes of nuclear and material physics, finding and adapting relevant computational technologies had to be an active process for Los Alamos.

In late 1943, the technical and physical limitations of traditional hand-computing methods spurred the lab's first significant computing innovation. Prior to his work in procurement at Los Alamos, Dana Mitchell had experienced the scientific-calculation

Disunity of Science: Context, Boundaries, Power, P. Galison and D. Stump, eds. (Stanford: Stanford University Press, 1996), 121-122.

capabilities of IBM PCAM equipment during his time at Columbia University. The astronomer Wallace J. Eckert and his staff at Columbia were only the second group to use accounting equipment for scientific calculation, using it to compute astronomical tables. The first group was the London National Almanac Office, which used the equipment in the 1920s to calculate moon-phase tables. In the 1930s, Eckert urged IBM to expand the capabilities of its product line for scientific uses, which it did with the support of its CEO, Thomas Watson, Sr. Mitchell suggested that T Division order IBM PCAM equipment to offload some of the calculation effort for the gun-type weapon from the hand-computation group. The punched-card machinery could perform the simpler but laborious neutron-diffusion calculations, freeing the human computers for more complex problems.⁵⁰ Ordered in the fall of 1943, the punched-card equipment arrived the next spring and consisted of three IBM 601 multipliers, one 402 tabulator, one reproducer-summary punch, one verifier, one keypunch, one sorter, and one collator. Due to lab secrecy, no one from IBM was allowed to accompany the machines to install them. While the army tracked down and cleared IBM maintenance specialist John Johnston to setup the machinery, Feynman, Stanley Frankel, and Eldred Nelson decided to speed up the process by assembling the equipment themselves, despite IBM's warnings to the contrary. With some snags, they had the machinery mostly functional before Johnston arrived.⁵¹

⁵⁰ Francis H. Harlow and N. Metropolis, "Computers and Computing: Weapons Simulation Leads to the Computer Era," *Los Alamos Science* 7 (1983): 134; Fitzpatrick, "Igniting the Light Elements," 63-65, 78-79, 107-108; Feynman, "Los Alamos from Below: Reminiscences of 1943-1945," 25-26; Charles J. Bashe, Lyle R. Johnson, John H. Palmer, and Emerson W. Pugh, *IBM's Early Computers* (Cambridge: The MIT Press, 1986), 23-25

⁵¹ Metropolis and Nelson, "Early Computing at Los Alamos," 350.

Feynman, Metropolis, and Nelson wrote of the assembly and set-up of the accounting machinery as an introduction for the involved T-Division personnel to the new computing paradigm. Feynman joked that Stanley Frankel became enamored with the equipment to the detriment of doing useful work. The machines were intended for the gun-type weapon's neutron-diffusion and critical-mass calculations, but T Division's computational priorities had begun to change during the months the PCAM equipment was on order. Hedging its bets on two designs with two different nuclear-material requirements, Los Alamos had begun development of a new, plutonium-based implosion-weapon concept, but needed to understand better the behavior of compressed plutonium to prevent spontaneous detonation. Oppenheimer requested in the fall of 1943 that John von Neumann, the Hungarian mathematician and polymath, work as a Los Alamos consultant because of his experience with hydrodynamics problems. Von Neumann confirmed the belief at Los Alamos that using compression to reach critical mass would increase the efficiency of a fission device, but the volume of calculations required for development and design of the concept was too great for desk-calculator-aided hand computation. The accounting machinery offered a promising alternative.⁵²

Even before the IBM equipment arrived, T Division began to assemble personnel for the equipment's operation, hiring the mathematician and IBM PCAM expert Naomi Livesay as the machine-crew supervisor and testing simulation procedures for the machines via hand calculation. While the machinery was operational within a week of

⁵² Robert W. Seidel, "From Mars to Minerva: The Origins of Scientific Computing in the AEC Labs," *Physics Today* 49 (October, 1996): 33; Metropolis and Nelson, "Early Computing at Los Alamos," 350; Feynman, "Los Alamos from Below," 25-26; William Aspray, *John von Neumann and The Origins of Modern Computing* (Cambridge, MA: MIT Press, 1990), 28-29.

arrival, training personnel and optimizing operating procedures took longer, as the scale, speed, and complexity of the computing operation at Los Alamos was greater than had ever been attempted with electromechanical accounting machines. Although IBM had tailored some of its product lines to suit scientific interests, like those of Wallace J. Eckert's astronomy work, the machines still had limited data storage, processing speeds, and numerical precision, requiring the scientific staff and equipment operators to be creative in how they simplified and approximated the problems to be run, how the machines were configured, and how the crews fed jobs into them.⁵³

Members of the Special Engineering Detachment were trained to operate the equipment, with the number of operators soon expanding to run more machines for longer hours. The first hydrodynamics problem run on the IBM equipment took three months to complete, a process that required inputting a punched-card deck representing each time-step in the integration equation through the machines in sequence until the simulation was finished. The SED crews—working in three shifts by early spring—and Richard Feynman reorganized the implosion calculations by color-coded cards (a system first tested with the T-5 hand-computing group), and rewired the machines to allow some operations to occur in parallel, reducing the calculating time for a hydrodynamics problem to three weeks.⁵⁴

The machines were not appreciably faster than human computers during a direct competition, but the machines did not tire as the humans did. Although, the dusty

⁵³ Metropolis and Nelson, "Early Computing at Los Alamos," 350; Fitzpatrick, "Igniting the Light Elements," 64-65.

⁵⁴ Metropolis and Nelson, "Early Computing at Los Alamos," 350; Feynman, "Los Alamos from Below," 25-26; Fitzpatrick, "Igniting the Light Elements," 64-65; Seidel, "From Mars to Minerva," 33.

environment impacted the reliability of the equipment, as the same dust that caused problems for desk calculators stuck to the relay contacts of the accounting machinery. With hundreds of relays in each machine, jobs sometimes needed to be run repeatedly. Fortunately for Los Alamos, many of the jobs were approximations that were usually forgiving of scattered errors, mitigating some of the reliability problem. The need for approximations, due to the limitations of the equipment, time constraints, and the limited understanding at Los Alamos of the complex material physics involved, was another issue that impacted the overall utility of the calculating effort, and will be addressed later in the chapter. In the spring of 1944, the implosion calculations were still not proceeding quickly enough for the wartime pace. Eldred Nelson negotiated with IBM to construct more elaborate PCAM models, as the 601 multipliers in use at Los Alamos could not multiply three products simultaneously, nor did they have a built-in divide function, requiring additional steps and operator interventions to perform these more complex tasks. IBM agreed, but the more-capable machines would not arrive until late that year.⁵⁵

Von Neumann, seeking additional computing capacity for Los Alamos, arranged for an unclassified test problem to be run on the IBM-Harvard Mark I electromechanical computer, which had been completed in March of 1944. Although the Mark I with 23 decimal digits offered greater precision than the PCAM equipment running the same problem at Los Alamos, it still did not offer an increase in calculating speed and lacked the flexibility T Division was seeking. The Mark I's limitations were a moot point, as its time was fully booked for other military problems. Von Neumann, in his capacity as a

⁵⁵ Metropolis and Nelson, "Early Computing at Los Alamos," 351; Aspray, *John von Neumann*, 27-28, 30-33.

consultant for Los Alamos and other military projects, continued his search for new and available sources of computing capability. One option he encountered was a proposed relay computer from Bell Labs. Bell Labs had had experience building sophisticated calculating machines using telephone relays since the 1930s, but the new machine would not have been available soon enough to aid Los Alamos, even if its performance had proven sufficient to the Laboratory's wartime needs.⁵⁶

Von Neumann's search in 1944, although not immediately fruitful for Los Alamos, had important consequences for both him and the Laboratory. It gave von Neumann an introduction to, and a growing interest in, the changing array of computing technologies that were in use or in development in the United States during World War II. As William Aspray and Nicholas Metropolis noted, von Neumann's experience with parallel processing on the PCAM equipment at Los Alamos contributed to his preference for simpler serial processing (processing a single datum point at a time on a central-processing device, rather than multiple data points being processed simultaneously across several devices) with later electronic digital system design. Von Neumann and T Division's conjoined search for computing capacity also shaped their impressions of either on the capabilities and possible applications of the various technologies and methods available during the most intense phase of Los Alamos' computing effort. Hand computing and electromechanical systems remained in use at Los Alamos well into the 1950s, but von Neumann and key members of T Division, such as Nicholas Metropolis and Edward Teller, were generally dissatisfied with the capabilities of the existing

⁵⁶ Ibid; Kenneth Flamm, *Creating the Computer: Government, Industry, and High Technology* (Washington, D.C.: The Bookings Institution, 1988), 30.

computing technologies during World War II. Although, his search for new computing capabilities for Los Alamos intrigued von Neumann with the possibilities of large electronic computing systems in automating computing work in physics applications. As von Neumann learned more about the various computing projects under development in the US between 1944 and 1945, he conducted a series of lectures for T Division about their possible applications for Los Alamos, spreading his excitement for those technologies to T Division. Based on von Neumann and T Division's growing, mutual understanding of the technologies then available, conventional computing technologies and methodologies seemed increasingly less suitable for modeling the complex physical processes that comprised Los Alamos' workload.⁵⁷

Von Neumann's chance encounter in 1944 with Herman Goldstine, the Army Ordnance liaison officer with the ENIAC project, and the resulting introductions of von Neumann and Los Alamos to the ENIAC are well-documented, and need not be analyzed in detail here. Instead, an important point needs to be made about ENIAC and the illustrious mathematician. The noteworthy point is how and why von Neumann and T Division were intrigued with the general-purpose electronic digital computing concept, and were amenable to pushing forward in testing it (which will be discussed later in this chapter) at that moment—when the technology was in its extreme infancy and with few guarantees as to its feasibility. Historians of computing such as James S. Small in *The Analogue Alternative* and Larry Owens in "Where are We Going, Phil Morse?" have

⁵⁷ Nicholas C. Metropolis, OH 135, oral history interview by William Aspray (29 May, 1987), Los Alamos, New Mexico. Charles Babbage Institute, University of Minnesota, Minneapolis, <http://purl.umn.edu/107493>; Fitzpatrick, "Igniting the Light Elements," 75.

convincingly noted that the "obviousness" that general-purpose electronic digital computing would become dominant over the variety of other forms of computing, a viewpoint often depicted in earlier histories of computing, was not the case among most scientific, engineering, and business professionals connected with computing in the 1940s and early 1950s. Early evangelists of *electronic digital* computing, like Herman Goldstine, Nicholas Metropolis, and others, who were instrumental in the early promotion and search for uses of the technology, often later presented the predominance of electronic digital computing as a natural conclusion to a "Darwinian" contest. Historians of computing in the 1970s and 1980s, who were often industry insiders, tended to maintain this view of "obviousness," which became a common narrative until the 1990s, even among some participants who, in the 1940s and 1950s, had not been so convinced of electronic digital computing's alleged "natural" superiority.⁵⁸

Stripping away the "obviousness" narrative, which was usually presented as justification after the fact for pursuing early electronic computing systems, helps to illuminate the reasoning behind von Neumann and T Division's growing interest in the potential of new computing technologies for mathematics and the sciences. Although couched in immediate (and not-so-immediate, in the case of the hydrogen bomb) wartime concerns, von Neumann and T Division's early pursuit of electronic digital computing

⁵⁸ Aspray, *John von Neumann*, 34-35; As cited previously, Thomas Haigh, Mark Priestley, and Crispin Rope, *ENIAC in Action: Making and Remaking the Modern Computer* (Cambridge, MA: The MIT Press, 2016) is the most detailed discussion of ENIAC history, including Los Alamos' use of the computer, currently available, and elaborates on the authors' article, "Los Alamos Bets on ENIAC: Nuclear Monte Carlo Simulations, 1947-1948," that focused specifically on how Los Alamos used and modified the ENIAC; Aspray interview with Metropolis; James S. Small, *The Analogue Alternative: The Electronic Analogue Computer in Britain and the USA, 1930-1975* (London and New York: Routledge, 2001), 2-8, 15-16, 271-276; Owens, "Where Are We Going, Phil Morse?," 35-36.

came to represent a larger suite of ambitions that became driving forces in and of themselves. The immediate interest in new computing technologies, particularly at Los Alamos over the course of 1944 and 1945, helped to inform and fuel the aspirational interests in electronic digital computing that came to characterize von Neumann and T Division's later evangelism of the technology, and their deliberate efforts to promote its development and use. T Division's aspirational interests in new computing technologies for mission-critical activities at Los Alamos continued to shape the Lab's approach to computer storage, networking, software security, operating systems, cluster computing, and other technologies for decades to come.

The immediate and formative experiences that contributed to the interest in unconventional computing at Los Alamos were intimately tied to the most salient features of the Laboratory's wartime work. The extreme cost of fissile materials and the risky nature of conducting tests with those materials for the weapons development program had shifted much of the material-physics research burden onto numerical methods. Existing numerical methods and technologies were already reaching their limits in a variety of fields in mathematics and the physical sciences well before the lab at Los Alamos was created. For example, it was John Mauchly's interest in applying electronic calculation to statistical problems in meteorology that initially spurred the development of the ENIAC. Meteorology had been, along with astronomy and ballistics, a significant user of hand computing, with its more ambitious efforts in weather forecasting and modeling having been stymied by the quantity of calculations required, particularly for time-sensitive work. The development of numerous technical solutions, such as the electromechanical

IBM-Harvard Mark I, George Stibitz's relay computers at Bell Labs, Vannevar Bush's analog Differential Analyzer, and the ENIAC, represented a flowering of approaches to similar impasses in existing numerical methods for science and engineering.⁵⁹

Specific to Los Alamos, the partial differential equations associated with atomic-weapons hydrodynamic simulations contributed to a growing appreciation within T Division of the enormity of Los Alamos' task, and the likely complexity of nuclear-physics research after the war. That appreciation stood in contrast with the earlier perception on the part of Oppenheimer and other early Manhattan Project planners who considered the fission-weapon question effectively "solved" in 1942. That appreciation expanded with Edward Teller's wartime research on the feasibility of thermonuclear weapons, the "technically sweet" problem that would come to occupy much of Los Alamos' computing effort after the war. T Division scientists such as Teller had grown increasingly impatient with the limitations of PCAM technology over the course of 1944. PCAM equipment helped offset the labor shortage at Los Alamos, a function it would continue to perform after the war, but the machines were not capable of processing problems of the complexity or at the level of precision that human computers were capable of performing. Anne Fitzpatrick noted that wartime T Division's experience with its limited computing capability and limited understanding of the physical processes being modeled was to rely on simplified estimates in its calculations, hoping to build-in

⁵⁹ Haigh, et al, *ENIAC in Action*, 17; See Grier, *When Computers Were Human*, for a more detailed discussion of how and when various fields in science, mathematics, and engineering began to encounter computational impasses that fostered interest in hand and mechanical computing between the eighteenth and twentieth centuries. In particular, the attempts at developing meteorological prediction centers based on hand computing offer important insights into the draw of automated computing to particular individuals and fields, without falling into the trap of electronic, digital computing being retroactively deemed the "obvious" choice among those interested in automating the tedious computing effort.

acceptable margins of error into Los Alamos' weapon designs to make up for the uncertainties facing the design processes.⁶⁰

The updated IBM PCAM equipment that Eldred Nelson had ordered in the spring of 1944 arrived at Los Alamos late that year, and was a technical improvement over the older 601 multipliers. Despite the improvement, the new machines provided neither the orders-of-magnitude performance increases, nor the flexibility that von Neumann, Metropolis, Teller, and others were seeking for simulation work. Metropolis and others grew increasingly convinced that mechanical systems, while still the most readily available source of automation anywhere, were not capable of the performance T Division sought as World War II approached its final year.⁶¹

In the mid-1940s, and into the mid-1950s, electronic digital computing was one possible technical solution, which had numerous "reverse salients" that left many professionals in the sciences, mathematics, and engineering skeptical about its feasibility, utility, and cost-effectiveness. Von Neumann, Metropolis, and T Division's interest in the ENIAC and early electronic digital computing was a gamble, just as the ENIAC project had been a gamble for the US Army, on a technical approach that had, on paper, the potential to deliver the kinds of computational performance and functionality that hand computing and electro-mechanical systems could not. Experiencing the limitations of conventional approaches in the face of the escalating computing demands from Los Alamos' wartime work, and personal interests in pushing the limits of technology and

⁶⁰ Aspray, *John von Neumann*, 30; Aspray interview with Metropolis; Harlow and Metropolis, "Computing & Computers," 134; Rhodes, *Dark Sun*, 476; Fitzpatrick, "Igniting the Light Elements," 100.

⁶¹ Aspray, *John von Neumann*, 30; Aspray interview with Metropolis; Harlow and Metropolis, "Computing & Computers," 134.

technique in the physical sciences, left T Division and von Neumann open to the unconventional when introduced to the ENIAC in 1944.⁶²

In late summer of 1944, von Neumann became a persistent figure at the University of Pennsylvania's Moore School, where the ENIAC was being completed alongside planning for a more capable follow-on machine, the EDVAC. Von Neumann's frequent presence at the Moore School and rapid (and later controversial) involvement with the EDVAC design, despite his grueling wartime schedule as a consultant at Los Alamos and elsewhere, is an indicator of the speed with which he became familiar and deeply interested in computing machinery. During von Neumann's search over 1944 for new computing solutions, and his involvement with the ENIAC project between late 1944 and early 1945, T Division had undergone significant organizational changes. One impetus to change was the shaky theoretical footing of the Super. Teller's pursuit of thermonuclear weapons, often at the exclusion of his assigned work on the implosion problem, had invited conflicts between himself and Oppenheimer and with Hans Bethe. A second spur to change was the growing computational demands of the implosion-weapon design. Enrico Fermi had joined Los Alamos in the summer of 1944, and made Teller group leader of F-1. Fermi created the new F Division upon his arrival, and placed F-1 specifically in charge of the Super.⁶³

A major consequence of the more focused work on the Super after the summer of 1944 was the realization, as with the implosion effort continuing in T Division, of the

⁶² Aspray, *John von Neumann*, 30; Aspray interview with Metropolis; Harlow and Metropolis, "Computing & Computers," 134.

⁶³ Aspray, *John von Neumann*, 34-39; Fitzpatrick, "Igniting the Light Elements," 108-112.

complexity of the physics and associated calculations needed for developing a fusion weapon. T Division's hand- and mechanical-computing efforts continued at a rapid pace in the months leading up to the Trinity Test in July of 1945, with the successful atomic explosion bearing out the numerical simulations conducted for the implosion weapon. Oppenheimer and his staff were confident enough with the gun-type design that it did not undergo full-scale testing, avoiding the waste of the US's limited U-235 supply. While the ENIAC likely would have been helpful with the simulations for the implosion weapon, had it been used for that application, the machine was still nearly a year from completion when von Neumann informed Los Alamos of the ENIAC project. Teller's group had begun to realize the scope of the computing effort that lay ahead for the Super, particularly when exploring the feasibility of the rudimentary design concepts under consideration. Not knowing if or when the ENIAC might be available, and with Los Alamos' IBM PCAM equipment occupied with implosion calculations, Teller began to seek the aid of Wallace Eckert, an advanced user of IBM PCAM equipment at Columbia University. This was ultimately unnecessary, because, in the summer 1945, von Neumann arranged for Los Alamos to be the first to use the ENIAC for Teller's feasibility calculations, arguing with ENIAC Project members that the Los Alamos Problem would be a more complete test of the machine's capabilities than any other available, including the Army's artillery-trajectory calculations for which the computer was originally intended.⁶⁴

⁶⁴ Fitzpatrick, "Igniting the Light Elements," 113-119; Harlow and Metropolis, "Computing and Computers," 134; Aspray interview with Metropolis.

F Division dissolved in October 1945, and Teller's F-1 group transferred back to T Division as T-7. T-7 group members Nicholas Metropolis and Stanley Frankel had toured the ENIAC facility at the Moore School that summer to begin developing the test program in conjunction with members of the ENIAC Project. Filling one-million punched cards, the first test program, run between December 1945 and March 1946, was significant for a multitude of reasons, not the least of which being that it was the first large-scale program run on a general-purpose electronic digital computer. In *ENIAC in Action*, Haigh, Priestley, and Rope address in detail the significance of the test program for the ENIAC, but for Los Alamos, the ENIAC test was most significant for what it revealed to T Division about the limitations of early electronic digital computing. The feasibility calculations had to be limited to one-dimension in space and excluded a significant array of physical variables, due to the limitations of the ENIAC's 1,000-word, tube-based main memory. Occupying about 95% of the ENIAC's control capacity, the Los Alamos test was, as von Neumann promised, far more strenuous on ENIAC's hardware than a typical artillery-trajectory calculation, which would have occupied closer to 25% of the system's resources. For the ENIAC team, the test program was a success; but, according to Nicholas Metropolis and Stanley Frankel, it produced largely inconclusive results about the feasibility of the Super for Los Alamos. Teller would, in April 1946, proclaim that the ENIAC series bore results in favor of the Super's feasibility, but numerical results generally considered favorable to Teller's project would not appear for another four years.⁶⁵

⁶⁵ Fitzpatrick, "Igniting the Light Elements," 118-121; Haigh, et al, *ENIAC in Action*, 81-84.

Teller's debatable conclusions aside, the ENIAC test provided Los Alamos with a demonstration, even within the ENIAC's hardware constraints, of the computational speed and flexibility of electronic digital computers compared to electro-mechanical systems. Metropolis recalled directly comparing the number-crunching speeds of the ENIAC with the IBM-Harvard Mark I, requiring seconds rather than minutes for the same calculation. The test also provided T-Division computing personnel with access to the ENIAC and the computer's designers at a point when the machine itself was still classified to the outside world. As much as electronic digital computing intrigued von Neumann, Metropolis, Frankel, and others connected to the Los Alamos computing effort, and as compelling a demonstration as the ENIAC provided for the potential of the technology for the computational impasses Los Alamos had encountered during and just after World War II, the ENIAC also provided a first-hand impression of the shortcomings of general-purpose electronic digital computing as it existed in the mid-1940s. Von Neumann quickly involved himself in the EDVAC design process, hoping to correct shortcomings of the ENIAC design. He likewise threw himself into searching for applications for the new technology, and putting his considerable name and reputation behind the support of its development, an effort that he maintained until his death in 1957.⁶⁶

Like von Neumann, key T Division personnel involved with computing became increasingly invested in the electronic digital computing model after World War II. For physicists like Edward Teller, and mathematicians like Stan Ulam, electronic-digital

⁶⁶ Aspray interview with Metropolis; Aspray, *John von Neumann*, 41-53.

systems like the ENIAC were the first technical solutions with the apparent means of addressing computational problems with precision and flexibility that had challenged Manhattan Project weapons research from the start, and, more importantly for Teller and Ulam, also appeared to open up avenues of fusion research that were entirely beyond the practical capabilities of traditional computing systems and methods. For Metropolis, Frankel, and others involved with the development and maintenance of Los Alamos' hand- and mechanized-computing facilities, electronic digital systems also appeared to have the performance potential to mitigate the Laboratory's severe post-war staff shortage. Approximately two-thirds of Los Alamos' wartime personnel departed between the end of the war and December 1945. Uncertainties about the lab's future spurred much of its scientific staff, including Oppenheimer, to return to industry and academia.⁶⁷

Norris Bradbury, the director of the Los Alamos lab after Oppenheimer departed in November 1945, was primarily occupied with keeping the Laboratory open in the immediate post-war years. The end of the war in August 1945, and the official end of the Manhattan Project in December 1946, signalled to many at Los Alamos and in the US government that the no-longer-secret city was also a no-longer-needed city. Those who departed reasoned that the status quo of limited government support for science would likely return, particularly because Leslie Groves and Vannevar Bush of the OSRD, among other planners in the wartime government scientific-research effort, refused to commit to research programs that extended beyond the end of the war. Oppenheimer, himself, was inconsistent in his views of whether Los Alamos should continue to exist

⁶⁷ Aspray, *John von Neumann*, 47; Chambers, "Technically Sweet Los Alamos," 159; Westwick, *The National Labs*, 32, 35.

after the war. Despite this, Groves and Oppenheimer, plus Norris Bradbury, became important forces for the continuation of the Manhattan Project lab system. Groves and Oppenheimer argued before the US Congress to maintain operational funding for the Manhattan Project labs until a permanent funding and management system could be arranged. Bradbury, following Groves' lead immediately after the war, committed to improving the spartan living conditions on the Los Alamos mesa, one of the major driving forces in the post-war departure of staff. Bradbury also began recruiting new staff, while attempting to retain those who had remained.⁶⁸

Meanwhile, the problems at Los Alamos were mounting. The Laboratory's staff levels were still much-reduced, after about two-thirds of the wartime staff had departed by late 1945. T Division, alone, lost twenty-seven senior staff members, leaving only eight in 1946. Those losses brought most research projects to a near standstill. In addition, many of the lab's hastily constructed temporary facilities were falling apart, requiring an infusion of funds, a new research direction, and a governing organization to provide them. Just as critical, Los Alamos had lost its drive toward a common goal, a feature that had characterized and unified the lab's work during the war. T Division's hand-computing and PCAM groups continued operating largely in support of refining atomic-weapon designs, the primary goal Bradbury chose for Los Alamos during the early Cold War years, when the lab was otherwise largely directionless. According to Bradbury, the two quickly developed wartime weapon designs were not practical for production, necessitating significant design improvements before atomic weapons could

⁶⁸ Chambers, "Technically Sweet Los Alamos," 160; Westwick, *The National Labs*, 36; Fitzpatrick, "Igniting the Light Elements," 84-85.

be produced in quantity. Los Alamos' electromechanical computing systems played an even more important role in offsetting the lab-wide loss of personnel. The Super development effort, as Anne Fitzpatrick noted, was, at the time, a low-priority at Los Alamos, compared to the task of keeping the Laboratory open. The limited computing resources at Los Alamos during this transitional post-war period made T Division, and members of Teller's group in particular, motivated to find more and better computation capacity, wherever and whatever that might be.⁶⁹

For Los Alamos, the years between 1946 and 1949 were largely about reconstructing the wartime Laboratory into a permanent research and development site, while lacking a strong unifying goal behind that effort. The formation of the Atomic Energy Commission (AEC) in 1947 by the US Congress provided a permanent funding source and organizing body for Los Alamos and the other assets of the Manhattan Engineer District (MED), the Army trusteeship of the Manhattan Project. The University of California, which had administered Los Alamos during World War II, entered a contract to administer Los Alamos Scientific Laboratory, the Laboratory's designation under the AEC system. Delegating facility administration to various contractors, including universities and corporations, had been the norm for many Manhattan Project sites during World War II. The logic behind delegating site administration was to utilize pre-existing, mission-relevant management structures for sites of complex scientific or industrial production.⁷⁰

⁶⁹ Chambers, "Technically Sweet Los Alamos," 159-163; Westwick, *The National Labs*, 32, 36-38; Fitzpatrick, "Igniting the Light Elements," 87, 96, 125-126.

⁷⁰ Westwick, *The National Labs*, 46-48, 53-55.

The distribution of management contracts for the former Manhattan Project labs among differing academic and industrial organizations—for example, the University of California for Los Alamos, the University of Chicago for Argonne, and Union Carbide for Oak Ridge (formerly Clinton Lab)—contributed to each of the AEC labs developing unique characteristics that stemmed from the mission focus of each lab, the backgrounds and professional compositions of their workforces, and the management styles of their contractor. While the general perception of the University of California's (UC) managerial effectiveness was not always favorable at Los Alamos during the war, and not all (including members of the university leadership) were in favor of continuing its management of Los Alamos after the war, UC was considered preferable to military management at Los Alamos. Los Alamos, under its postwar UC contract, developed a collegiate management style and atmosphere, where LASL staff members were designated as UC faculty. Some staff members at UC-managed Los Alamos recalled the managerial atmosphere as affording more flexibility in personal and group research approaches than an industry job or at a typical government-managed facility, while offering better pay than a regular academic position.⁷¹

While the Laboratory adapted to the post-war climate and UC's permanent management, and Los Alamos developed its own culture and characteristics that would persist for years, T Division's post-war character and approaches to computing began to take shape. Of the many technical influences on Los Alamos computing in the period between 1946 and 1949, the period that ended with the detonation of the first Soviet

⁷¹ Ibid.

atomic weapon, a few influences decisively shaped T Division's emerging computing strategy. One influence was the extensive use of the ENIAC. The ENIAC, despite its hardware limitations, provided Los Alamos with a singular test platform for increasingly sophisticated forays into new numerical simulation techniques before Los Alamos itself had any sort of comparable computer in-house. The most important of the early computational developments was the adaptation of Monte Carlo simulations for electronic digital computers, which Los Alamos first employed on the ENIAC in 1948.⁷²

The Monte Carlo approach encompasses a variety of techniques for utilizing probability to resolve otherwise deterministic problems. The central concept dates back to the eighteenth century, and Enrico Fermi experimented with similar hand-calculated techniques in the 1930s before they were given the name. Stan Ulam, inspired through games of solitaire he played during an illness in 1946, first suggested using what became known as "Monte Carlo" to contend with Los Alamos' mathematical difficulties with calculating the behavior of subatomic particles. These difficulties were a continuation of the neutron diffusion problems that first inspired the use of PCAM equipment at Los Alamos during the Manhattan Project. Large volumes of Monte Carlo calculations, which would be necessary in a detailed simulation, were impractical to perform by hand (although Enrico Fermi's analog FERMIAC trolley of the early 1950s could provide useful data one neutron at a time), but high-speed electronic computers were, according to Ulam and von Neumann, ideally suited for the task.⁷³

⁷² Aspray, *John von Neumann*, 111; Fitzpatrick, "Igniting the Light Elements," 129.

⁷³ Haigh, et al, "Los Alamos Bets on ENIAC," 44-46; H. L. Anderson, "Scientific Uses of the MANIAC," *Journal of Statistical Physics* 43, nos. 5/6 (1986): 732-733.

Ulam (who had returned to Los Alamos after a brief stint at the University of Southern California), Metropolis, Robert Richtmyer (the physicist and post-war T-Division leader until 1947), von Neumann, and others, adapted the Monte Carlo concept, which Metropolis named after the casino in Monaco, to be employed on computing machinery. The simulation of subatomic particle behavior through statistical sampling (in which a statistical model of probable particle motion is utilized to approximate how particles would behave within a given set of conditions) was a key development in the application of electronic digital computing for exploring systems too complex to model using other methods. The continued development and importance of Monte Carlo, and other numerical simulation methods that would follow, to Los Alamos computing will recur throughout this dissertation. It is sufficient to note here that the initial adaptation of Monte Carlo methods for electronic computing systems was intimately connected to Los Alamos' unusually early use of and familiarity with electronic digital computing in the immediate post-war period, and the desire of T Division and von Neumann to find and map their most pressing applications to the known capabilities, and the anticipated capabilities, of the new machines.⁷⁴

Connected with the development of increasingly sophisticated new simulation methods at Los Alamos in the immediate post-war period was the continued search for additional computing resources at and beyond the Laboratory. Hand computing still provided most general computational support at Los Alamos, and would continue to do so well into the 1950s, as human computers offered the most capable and versatile

⁷⁴ Roger Eckhardt, "Stan Ulam, John von Neumann, and the Monte Carlo Method," *Los Alamos Science*, Special Issue (1987): 131, 133, 137.

computing services. The AEC was unwilling to fund any significant expansion of the hand-computing effort at Los Alamos (let alone the "hundreds" of human computers that von Neumann indicated would be necessary for the Super's feasibility assessment), pushing Los Alamos further toward labor-saving computer machinery. New punched-card machines were added to T Division's existing roster, and were primarily used for implosion calculations in the refinement of fission weapons. They also found some use on the Super's design and feasibility assessment, largely because they were the only computing machines at Los Alamos, and the hand-computing group was unable to expand under AEC funding. Electromechanical calculators of the late 1940s and early 1950s, such as IBM's 603 and 604 Calculating Punch, were more capable than the 1930s-era models used during the Manhattan Project. The punched-card equipment at Los Alamos, despite its recent improvements, was still too limited in capacity, flexibility, and precision to push very far into the mathematical requirements Teller's group encountered with the Super. Although, electromechanical systems were still being explored at Los Alamos in the late 1940s for the significant computing challenges facing T Division.⁷⁵

Robert Richtmyer, on von Neumann's suggestion in 1948, created a Monte Carlo simulation relating to both the fission and fusion aspects of the Super for IBM's electromechanical Selective Sequence Electronic Calculator (SSEC). The SSEC was a unique machine intended as both a technical successor to the IBM-Harvard Mark I and as an exercise for IBM's Watson Scientific Computation Laboratory. Von Neumann suggested Richtmyer employ the SSEC, as it had a larger data store and a larger

⁷⁵ Metropolis and Nelson, "Early Computing at Los Alamos," 354; Fitzpatrick, "Igniting the Light Elements," 135-136, 140, 290.

instruction set than ENIAC, despite being a much slower system overall. Coding the simulation for the SSEC proceeded slowly, and Richtmyer's job did not run until 1950. This long-delayed use of the SSEC took place after T Division had better solidified its commitment to general-purpose electronic computing, largely due to its first-hand experience exploring the capabilities and addressing some of the difficiencies of the new computing paradigm.⁷⁶

In 1948, the same year the SSEC first entered service, John von Neumann, his wife, Klari von Neumann, and Nicholas Metropolis aided in modifying the ENIAC to provide the machine with a limited stored-program capability and an expanded instruction set. The ENIAC modifications, derived in part from the EDVAC planning effort and von Neumann's research-computer project at Princeton, significantly reduced the setup time of a problem on the ENIAC (from days to hours), while improving its programming feature set. The Los Alamos Monte Carlo runs on the ENIAC were all conducted after these modifications. Like T Division's initial use of the ENIAC in its original configuration in 1945 and 1946, the ENIAC modifications and Monte Carlo runs of 1948 left a cautiously optimistic impression in T Division about the potential of the new computing paradigm, but also a recognition of how much work lay ahead to adapt the still immature general-purpose computing technologies to the problems facing Los Alamos.⁷⁷

⁷⁶ Fitzpatrick, "Igniting the Light Elements," 130-136.

⁷⁷ Martin Campbell-Kelly, William Aspray, Nathan Ensmenger, and Jeffrey R. Yost, *Computer: A History of the Information Machine*, Third Edition (Boulder, CO: Westview Press, 2014), 104-105; See the original technical description for the IAS computer for a detailed outlining of the stored-program concept that influenced the modification of the ENIAC, and set forth the architecture underlying the IAS-based systems, including the first MANIAC: A. W. Burks, H. H. Goldstine, and J. von Neumann, *Preliminary Discussion*

The modifications that made the ENIAC easier to program reflected important system-architecture and programming concepts that were spreading through numerous computer development efforts worldwide, and were making fully electronic computers more capable and attractive as tools (at least among certain high-end scientific users like Los Alamos). Uncertainties remained among scientists and engineers in the late 1940s over whether or not electronic digital computers would be capable of tackling T Division's fission- and fusion-related research problems, but interest in the technology grew steadily at Los Alamos as some of the major questions about usability were being addressed. T Division's leadership and its technical experts in computing (who were mostly the same individuals, as T Division only consisted of fourteen permanent staff members in 1948, up from eight in 1946) built upon the already active role Los Alamos had played in scientific-computing development since World War II with plans to bring Los Alamos to the front lines of computing R&D.⁷⁸

As early as 1946, Bradbury, Richtmyer, Teller, and others began to suggest that Los Alamos build or purchase its own general-purpose, electronic digital computer; one that would surpass the limitations of the ENIAC, while reducing the time-consuming task of borrowing time on distant machines. The slow rate at which electronic digital computers were being constructed across the US tempered T Division's hopes to purchase a computer in the near future. For example, the EDVAC design effort encountered

of the Logical Design of an Electronic Computing Instrument (Princeton, NJ: Institute for Advanced Study, Princeton, 1946).

⁷⁸ Fitzpatrick, "Igniting the Light Elements," 136; J. Carson Mark, "A Short Account of Los Alamos Theoretical Work on Thermonuclear Weapons, 1946-1950," Los Alamos Scientific Laboratory, LA-5647-MS (July 1974), 3.

intellectual-property disputes in 1946 that delayed its design and construction for three years. J. Presper Eckert and John Mauchly's Eckert and Mauchly Computer Corporation (EMCC) was incorporated in 1947, and there was little indication of when that company's first large-scale machine (what ultimately became the UNIVAC in 1951) might become available. Established business-machine companies, like Remington Rand and IBM, were primarily busy producing punched-card accounting machinery and typewriters following the post-war demand, and were reluctant to move too quickly into the fledgling general-purpose computing concept and its unproven market. Los Alamos' options were obviously limited, and T Division initially sought a collaborative effort with a trusted ally to save time and reduce the risk of developing a completely independent system.⁷⁹

John von Neumann's Princeton Institute for Advanced Study (IAS) computer entered development in late 1945. While its construction encountered significant delays into the early 1950s, largely due to technical difficulties with memory and storage device hardware, von Neumann lined up significant support for the machine from Princeton University, industry, and the military. Von Neumann and his allies successfully promoted the construction of IAS-based systems at a variety of research centers. The computer, according to von Neumann and his fellow advocates, was a powerful scientific tool, similar to cyclotrons or other prestigious pieces of hardware, which self-respecting centers of scientific research dare not be without. Three AEC facilities were ultimately interested in building IAS-based systems: Argonne, Oak Ridge, and Los Alamos.⁸⁰

⁷⁹ Cuthbert C. Hurd, "Early IBM Computers: Edited Testimony," in Jeffrey R. Yost, ed., *The IBM Century: Creating the IT Revolution* (IEEE Computer Society: CS Press, 2011): 71-73; Aspray, *John von Neumann*, 51, 59-60, 85.

⁸⁰ Fitzpatrick, "Igniting the Light Elements," 129-132; Aspray, *John von Neumann*, 91.

By 1948, with its growing desire to possess a local general-purpose computer, and the realization that such systems were not entering the market anytime soon, T Division at Los Alamos settled upon building its machine locally. The decision to build a machine was primarily in response to Monte Carlo calculations that T Division anticipated to be an integral part of determining the needed fuel composition for a thermonuclear weapon, among other numerical-simulation problems. Carson Mark, a Canadian-born mathematician who joined Los Alamos in 1945, became T Division's leader after Richtmyer stepped down in 1947. Mark, as was the norm for most scientist-managers at Los Alamos during the early Cold War, was directly involved with determining the research direction of his division, while also working as a research scientist himself. As a consequence, Mark was intimately familiar with the computational problems relating to T Division's work in refining fission weapon designs and exploring the thermonuclear problem, just as he was directly familiar with the state of the computational technologies needed to address those problems. It was Mark who initiated the development of a local computer at Los Alamos, and placed Nicholas Metropolis as the project lead. That is, after Mark enticed Metropolis back to Los Alamos from the University of Chicago with the promise that he could construct a computer at the Laboratory.⁸¹

Metropolis and Stanley Frankel had joined the Chicago faculty with the understanding that the university would begin a computer-development effort under their direction. When the university president's plans changed, and indefinitely delayed the computer project, Carson Mark's proposal drew Metropolis back to Los Alamos. After

⁸¹ Ibid; Mark, "A Short Account of Los Alamos Theoretical Work," 5; Frank Harlow, "J. Carson Mark," Los Alamos National Laboratory Daily News Bulletin (March 4, 1997).

his return, Metropolis conferred with von Neumann's group about producing a copy of the IAS computer, and organized the engineering staff at Los Alamos. Numbering six staff members, including chief engineer, Jim Richardson, the engineers quickly duplicated the IAS arithmetic processor, the number-crunching central unit of the machine. The new T-7 Computer Research and Development Group, which Metropolis formed in February of 1949, undertook the computer's construction effort.⁸²

T-7, as its full title suggests, was about more than building, maintaining, or apportioning user access to a single computer. The new group reflected T Division's growing, dual-prong approach to computing as both a lab-wide service and as a field of research. It was a response to the significant and escalating computational hurdles Los Alamos had faced since the founding of the Laboratory. It was also a reaction to the limited range of technical possibilities (and outside experts in those technologies) that appeared conducive to tackling LASL's numerical challenges. Not to be ignored, T-7's research focus also represented the professional and personal interests that many staff members in the still modestly sized T Division shared in the largely unexplored possibilities with machine-based computing. This dual approach to computing was not unusual among early scientific-computing centers in the 1940s and 1950s, at a time when computing technologies and the means of using them were in a state of rapid change and

⁸² Mark B. Wells, "MANIAC," in "Computing at LASL in the 1940s and 1950s," Los Alamos Scientific Laboratory, LA-6943-H (May 1978), 16-20; Aspray interview with Metropolis; Anderson, "Scientific Uses of the MANIAC," 737; N. Metropolis, "The MANIAC," in Nicholas Metropolis, J. Howlett, and Gian Carlo Rota, eds., *A History of Computing in the Twentieth Century: A Collection of Essays* (New York: Academic Press, 1980), 460-461; Aspray, *John von Neumann*, 91; Carson Mark, OH 22909, oral history interview by Kenneth W. Ford (24 February, 1995), Los Alamos, New Mexico. American Institute of Physics, College Park, MD, <http://www.aip.org/history-programs/niels-bohr-library/oral-histories/22909> (accessed 10 March 2018).

exploration. T Division considered itself to be at an unusually advanced frontline in the need and development of computing technologies, due to its direct connection to atomic-weapons development and the US position in the new Cold War with the Soviet Union.

T Division valued, and prided itself in, developing a measured independence in its computing capabilities. Possessing local, world-class expertise in mission-critical technologies became a priority in T Division. The fear was that overdependence on outside experts and industry vendors could impair the Los Alamos mission, should technical requirements or Cold War strategies change, and Los Alamos find itself without the local expertise on tap to change with them. Los Alamos' position as an advanced center of computing, a reputation earned starting in World War II, was also a point of pride within T Division. This form of pride was a common sentiment among the staff at Los Alamos, particularly relating to the Laboratory's brief but prominent history in atomic and material physics.⁸³

T Division, as it transitioned out of the mid-to-late 1940s, embraced computing as both a service and as a field of local, mission-critical research. In the view of Carson Mark, the Laboratory's unusual computational requirements had only grown more unusual after World War II had ended, and keeping pace with technology and hiring qualified staff was essential. Underlying that viewpoint was the use of computing equipment to offset the loss of wartime personnel, and the directives to improve the efficiency (requiring more accurate modeling) of post-war fission-weapon designs and to

⁸³ Ford interview with Carson Mark; Aspray interview with Metropolis; John Morrison, oral history interview by Nicholas Lewis, Los Alamos, NM (July 1, 2014); Westwick, *The National Labs*, 154.

examine the feasibility of the "Super," an effort fraught with problems in material physics that were unanticipated during wartime.

T Division still maintained the majority of the Laboratory's computing resources, while also being their primary user. As T Division became more deeply aware of the mathematical complexity that its work with numerical modeling entailed, Carson Mark, Metropolis, Teller, and other leading figures in the division became convinced, as Metropolis phrased it, that "it was inevitable that the Laboratory would suffer total immersion in computing." This perception fueled another important drive behind T Division developing its own computing R&D capacity: attracting new staff members to Los Alamos by offering the latest in scientific and technical equipment for their use. The argument throughout the AEC lab system during the early Cold War period was that funding basic research and providing world-class equipment was essential to staffing the laboratories with top-quality personnel. The promise of building a computer had been, after all, what drew Nicholas Metropolis back to Los Alamos. As Cold War pressures escalated in late 1949, T Division, like most science divisions at Los Alamos, focused heavily on recruitment, and its computing operations and research groups became important tools in that recruiting effort.⁸⁴

Full-scale nuclear tests were the primary means of verifying weapon designs from the "Gadget" in 1945 up until the end of live testing in the US in the 1990s, but numerical methods were indispensable for conserving the US's limited supply of expensive fissile materials, and for exploring design iteration and concept feasibility

⁸⁴ Metropolis, "The MANIAC," 457; Harlow, "J. Carson Mark."; Mark, "A Short Account of Los Alamos Theoretical Work," 6-8; Westwick, *The National Labs*, 154-155; Lewis interview with Morrison.

where real-world testing would not be practical or desirable. The complexity of the Super, which required generating physical forces sufficient to fuse lighter elements into heavier elements, made accurate and detailed computer modeling essential. The behavior of materials, including the determination of which elements and in what concentrations could undergo fusion, were major stumbling blocks, not only in the development of thermonuclear weapons, but also in determining whether such a device was even technically possible. Following World War II, there was little political push for pursuing thermonuclear weapons, partly due to the US nuclear hegemony, and partly due to Oppenheimer, Bradbury, and the AEC restricting discussion of the topic outside very closed circles at Los Alamos and its immediate consultants, like John von Neumann, Hans Bethe, and Enrico Fermi. There was also little military interest in a weapon that was orders-of-magnitude more powerful than an atom bomb, and, even if there were interest within the military, the US Air Force had no systems capable of delivering one. It was in this context that T Division worked in the late 1940s on determining the feasibility of a fusion device, in what was largely a technical exercise waiting for sufficient computational capacity to perform the needed, more comprehensive calculations than were possible with existing equipment.⁸⁵

On August 29, 1949, the Soviet Union detonated its first atomic weapon at Semipalatinsk, in what is now Kazakhstan. That the Soviets developed an atomic bomb was not a surprise to the Western Powers, but it was completed far more quickly than anticipated. As a consequence of the US having lost its nuclear hegemony, US President

⁸⁵ Metropolis, "The MANIAC," 457; Norris E. Bradbury, interview by Arthur Lawrence Norberg, Los Alamos Scientific Laboratory, Los Alamos, NM (1980); Fitzpatrick, "Igniting the Light Elements," 29-30.

Harry S. Truman ordered the initiation of a "crash" program on thermonuclear-weapon development. The AEC and the General Advisory Committee (GAC), the committee of scientific and technical advisors to the AEC, opposed the decision for largely moral and strategic grounds, just as the GAC had opposed a thermonuclear program after World War II. The AEC and the GAC argued that the crash program would waste limited resources, and would divert staff needed elsewhere. The US public's and leadership's reaction to the Soviet atomic test, and the beginning of the Korean conflict only five months after Truman's announcement, cemented the emergency response of accelerating thermonuclear development in the US. Los Alamos and the other AEC labs moved forward with the new directive, with Los Alamos adopting a six-day work week, and spurring a flurry of research activity not seen at the Lab since the Manhattan Project.⁸⁶

Like the rest of Los Alamos, the T-Division computer project felt the strain of the new program. Coping with the demands of the crash program, T Division jumped in size from fourteen staff members in 1948 to thirty-five staff and three full-time consultants in 1950. The thermonuclear-development group was in desperate need of computing resources, but the Los Alamos computer, which Nicholas Metropolis dubbed the "MANIAC" in a self-professed attempt to end the convention of giving computers peculiar-sounding names ending in "IAC," and von Neumann's IAS computer were not completed until March and June of 1952, respectively. Delays in these computers' construction spurred Los Alamos scientists to continue using outside systems, including the ENIAC, IBM's SSEC, and the newly available, fully programmable electronic SEAC

⁸⁶ Norberg interview with Bradbury; Westwick, *The National Labs*, 120-121.

at the National Bureau of Standards. The SEAC was, like the MANIAC, based on the IAS design, but was scaled-down to ensure that it could be completed more quickly (completed in 1950) than the other IAS-based systems being constructed at the time.⁸⁷

T Division also added the latest in punched-card machinery to address some of the local demand, installing the sophisticated IBM Card-Programmed Electronic Calculator (CPC) starting in late 1949, with six soon in operation at Los Alamos until 1956. The CPC, a hybrid electronic and electromechanical calculator with 1,400 vacuum tubes and 2,000 relays, provided the bulk of T Division's local computing capacity before the completion of the MANIAC, during the time in which numerically testing the feasibility of a thermonuclear device became most frenetic. The CPC began IBM's effort to collect and develop the personnel needed for creating a dedicated scientific-computing development group, just as high-end scientific customers like Los Alamos began to clamor for fully electronic machinery. It was between the use of onsite CPC machines at Los Alamos, and off-site machines like the SSEC, New York University's UNIVAC I, the ENIAC, and particularly simulations on the SEAC, that Stan Ulam and Edward Teller settled upon a solution to the thermonuclear problem that largely ended the uncertainty over feasibility among physicists in January 1951.⁸⁸

The IAS computer's construction began in 1945 with an exploration of the technical options to complete the computer. The Princeton effort hit numerous delays

⁸⁷ Aspray, *John von Neumann*, 79; Mark, "A Short Account of Los Alamos Theoretical Work," 3; Metropolis, *The MANIAC*, 458; Wells, "MANIAC," 16; "Monthly Record of SEAC Computation, June 1952-July 25, 1953," Margaret R. Fox Papers, CBI 45, box 4, folder 7, CBI.

⁸⁸ Emerson W. Pugh, *Memories that Shaped an Industry: Decisions Leading to the IBM System/360*, (Cambridge, MA: MIT Press, 1984), 28; Bashe, et al., *IBM's Early Computers*, 68-72, 131; W. Jack Worlton, "Hardware," in "Computing at LASL in the 1940s and 1950s," Los Alamos Scientific Laboratory, LA-6943-H (May 1978), 1-6; Galison, "Computer Simulations and the Trading Zone," 136-137.

due to problems with finding adequate storage and memory devices, and then adapting the machine's design to work with the available technical solutions. Storage and memory technologies (although there was considerable overlap between the two) were significant reverse salients in early electronic computing. Finding technologies that were fast, precise, and reliable enough for the unusually high demands of electronic computing was no easy task in the 1940s and early 1950s, and resulted in delays and compromises for many computer-construction projects. Truman's order to begin a crash thermonuclear development effort in the US impacted the IAS-computer group, just as it had Los Alamos. The IAS computer was intended to be of use to the AEC labs, but the Truman directive, after several years' worth of project delays, spurred the Princeton group to cancel summer vacations and work long hours to complete the machine to aid in hydrogen-bomb calculations.⁸⁹

At Los Alamos, T-7 was able to avoid many of the pitfalls of the IAS machine by either finding different suppliers, or developing some solutions locally. For example, Jim Richardson, the MANIAC's chief engineer, had experience with the electrostatic, CRT-based memory technology, or "Williams tubes," which had caused delays for the IAS team, due to vendors failing to produce functional units as promised. Richardson found a more readily available two-inch cathode-ray tube for the MANIAC instead of the troublesome five-inch tubes used on the IAS system. T-7 also opted to avoid problems the IAS team had experienced with its memory-control system by developing a completely custom system at Los Alamos, saving considerable construction time in the

⁸⁹ Aspray, *John von Neumann*, 79, 82-85.

process. The MANIAC's I/O (input-output) systems were also quite different from the IAS computer's, with the T-7 engineers opting to use electronic over electromechanical I/O technologies whenever possible. This increased the speed and reliability of moving data onto and off the computer, a lesson learned through many years of experience with electromechanical computing at Los Alamos. Metropolis was fond of repeating:

"Whenever electromechanics can be replaced by electronics, do so!"⁹⁰

After their completion, the MANIAC and the IAS computer differed enough that the two were not code compatible, meaning that applications written for one would not execute on the other. T-7 had modified parts of the original design to suit Los Alamos' purposes, which included faster components and alterations to the system architecture. As a consequence, Los Alamos scientists, including Richtmyer and Foster Evans, ran several problems on the IAS system rather than the MANIAC, as the problems had been developed for the original IAS design. Neumann chided the T-7 group for diverging from the IAS design to such an extensive degree. Although, by letting the IAS team go through the development process first, T-7 was able to make adjustments to avoid most of the known issues with both the procurement of components from vendors, and the known problems with the IAS computer's architecture, saving several years of technical-exploration and construction time over the IAS system.⁹¹

As a consequence of being non-code compatible with the IAS system, the MANIAC was better suited to the work at Los Alamos than the IAS computer,

⁹⁰ Mark, "A Short Account of Los Alamos Theoretical Work," 3; Anderson, "Scientific Uses of the MANIAC," 737; Metropolis, "The MANIAC," 461.

⁹¹ Campbell-Kelly, et al., *Computer*, 168; Aspray interview with Metropolis; Aspray, *John von Neumann*, 92.

supporting a more refined electrostatic memory technology, and I/O systems that were faster and easier to manage. Although faster and more reliable than the mercury delay-line memory used in several early electronic computers of the early 1950s, and far larger than the expensive vacuum-tube memory of the original ENIAC design, the electrostatic memory tubes on the MANIAC, and all the IAS-based systems, were temperamental, and prone to wearing out over time. Growing out of radar-system research, the tubes stored data by firing electrons at a phosphorescent layer at the front of the tube, which would momentarily retain a charge. A fine wire mesh mounted on the front of the tube would let the computer read the charge differential of the individual locations on the tube that represented a particular sector of the machine's main memory. The data had to be continually refreshed to maintain the difference in charge. William (Bill) Spack, who joined the T Division computing effort as a technician immediately after World War II because of his training with radio equipment, recalled that Jim Richardson was among the few engineers who could keep the MANIAC's memory tubes properly calibrated.⁹²

The MANIAC also used a magnetic drum (similar in concept to later hard disk drives, but with a rotating drum instead of platters) sourced from Engineering Research Associates (ERA) in St. Paul, Minnesota, as a secondary memory system. MANIAC's secondary drum memory was similar in function to "virtual" memory in later personal computers, where a PC's operating system would use hard-disk space to simulate the presence of more memory than was actually installed on the computer, but at a performance penalty. Slower than the electrostatic main memory, the rotating drum had

⁹² Campbell-Kelly, *Computer*, 70-71; William Spack, oral history interview by Nicholas Lewis and Robert Cunningham, Los Alamos, NM (June 25, 2014).

a much larger capacity, at 10,000 40-bit words compared to the 1,024 words of main memory, and was a less expensive means of adding a large amount of additional memory to the system. Each facility that produced an IAS-based computer tended to follow its own path with regard to its adherence to the IAS design. Argonne, for example, made nearly an exact copy with its AVIDAC, while Oak Ridge's ORACLE, like the MANIAC, used improved components over the IAS computer.⁹³

The formation of T-7 also began the development of software (although the term itself would not be coined until the late 1950s) needed for the MANIAC upon its completion. Under the physicist John Jackson, a team comprised largely of human computers who had transferred from the T-5 Hand Computation Group to T-7 began research on assembly languages to produce one for the MANIAC. An assembly language was a low-level programming language that was closely coupled to the architecture of a specific computer. Not as easy to use or as portable as a high-level programming language (the earliest of which were entering development in the early 1950s), an assembly language allowed for an exact (or near exact) symbolic representation of the computer's architectural instructions. Because it was tailored to a specific architecture, and mapped directly (or almost directly) to that architecture's built-in instruction set, an assembly language was more efficient than a high-level language, which was beneficial among early, resource-limited computers. As computer resources improved in the early 1950s, and computer memories became larger, the need for extremely parsimonious coding became less of a concern, and made high-level languages more practical. The

⁹³ Aspray, *John von Neumann*, 83-85, 92; W. Jack Worlton, "Hardware," 1-6.

MANIAC assembler (which converted code in assembly language into the system's machine language) was completed in time for the computer's testing phase, beginning in late 1951.

The T-7 coding team also produced a continually growing subroutine library to allow coders (the term "programmer," much like "software," did not yet exist) to save time by using existing subroutines for their coding work, rather than write new code each time it was needed. With the completion of the MANIAC, the mathematician Mark B. Wells also launched T-7's research into high-level programming languages, which would develop into the long-running MADCAP programming language R&D program, an effort that would continue into the 1970s.

The code-development work in T-7 was not simply an addition to the hardware development of MANIAC. According to Mark Wells and Carson Mark, the two were tightly coupled from the beginning, with Jackson's coding team working with Metropolis and Richardson's hardware team to incorporate desirable features and improvements into the MANIAC, based on the problems Los Alamos scientists were designing to run on the machine. One consequence was that the MANIAC deviated from the IAS computer, but was easier to use, and better suited to the needs of Los Alamos. Another consequence of the tightly coupled development effort was the formalization of an interconnected hardware and software R&D program at Los Alamos. Computing, in its various forms, had existed as a service at Los Alamos since 1943, but T-7 institutionalized computing R&D as a sanctioned, mission-critical research focus. It was, by no coincidence, located

in the division with the greatest dependence on numerical methods in the H-bomb's development.⁹⁴

The MANIAC was not fully completed and available for use after testing until March 1952, and the IAS computer in June 1952, but the MANIAC went into immediate service to perform implosion calculations. The MANIAC was the first computer available to Los Alamos that was capable of simulating the proposed "Teller-Ulam" thermonuclear device. For the device to work, the computer was needed to determine the correct fuel mixture and geometry, problems that could not be solved empirically. Over spring and summer in 1952, the MANIAC ran calculations needed to complete LASL's proof-of-concept thermonuclear device, a massive assembly hauled to the South Pacific that fall. Producing a 10.4 megaton explosion, the Ivy Mike test of November 1952 obliterated the Enewetak Atoll island of Elugelab where the test was conducted, conclusively demonstrating the Teller-Ulam concept's feasibility and dramatically verifying the MANIAC's simulation results. The drawn-out feasibility testing of the H-bomb demonstrated to von Neumann and Los Alamos the necessity of powerful computing machinery. Hand computing, even conducted with the large numbers of human computers von Neumann had indicated would be needed in the 1940s, would have required years longer. It also indicated that, as Carson Mark had stated in 1949, earlier general-purpose computers, like the ENIAC, lacked the resources to produce conclusive

⁹⁴ Edward A. Voorhees, "Software and Operations," in "Computing at LASL in the 1940s and 1950s," Los Alamos Scientific Laboratory, LA-6943-H (May 1978), 7; Campbell-Kelly, et al., *Computer*, 170-175; Howard B. Demuth, John B. Jackson, Edmund Klein, N. Metropolis, Walter Orvedahl, James H. Richardson, "MANIAC," Los Alamos Scientific Laboratory, LA-1725 (1954), i-ii, 1; Harlow and Metropolis, "Computing & Computers," 136; Mark B. Wells, "MANIAC," in "Computing at LASL in the 1940s and 1950s," Los Alamos Scientific Laboratory, LA-6943-H (May 1978), 16-18; Fitzpatrick, "Igniting the Light Elements," 155.

results for T Division's thermonuclear work. Machines of the MANIAC's capability, in T Division's estimation, would be the bare minimum required for adequately detailed simulation work at Los Alamos, and for moving forward from the Teller-Ulam proof-of-concept device, to more compact, practical explosive devices the US military was capable of delivering.⁹⁵

The new MANIAC quickly proved its value to T Division in the rush to develop a thermonuclear device before the Soviet Union, after Los Alamos had been stymied for years by inadequate numerical methods and resources. This demonstration did much to settle any lingering doubts about the capability of general-purpose computers to handle large-scale simulations that previously appeared insurmountable, even on earlier computing machines. It also stressed for T Division the likely future importance of high-end general-purpose computers for the Laboratory's mission, which aided in solidifying the computing R&D approach T Division established during the MANIAC's construction. As T Division began to explore the wider capabilities of the MANIAC (from hydrodynamics codes, to genomic research, to chess-playing programs), and to consider Los Alamos' next steps in computing, Los Alamos' place in the AEC weapons complex began to change.⁹⁶

Between Truman's call for an accelerated H-bomb development effort in 1950 and the successful test of the Teller-Ulam device in the Marshall Islands, Teller's personal incompatibility with Los Alamos and its management came to a head. Teller argued that

⁹⁵ Fitzpatrick, "Igniting the Light Elements," 166-167, 289-290.

⁹⁶ Harlow and Metropolis, "Computing and Computers," 136; Wells, "MANIAC," 19; Metropolis and Nelson, "Early Computing at Los Alamos," 355-356.

Los Alamos was hampering the thermonuclear development program, and that the Laboratory lacked the expertise and drive to produce a thermonuclear device. Much of the delay in the feasibility testing was owed to the time required for the IAS-based computers to become available. Teller himself had suggested in 1947 that Los Alamos deliberately delay its thermonuclear development program to wait for better computer resources. In 1952, Teller argued that creating a second weapons Laboratory would foster competition and ease some of the burden in developing thermonuclear weapons.

Teller's certainty about the H-bomb's feasibility predated the empirical justification for that certainty, which fueled many of his altercations with other Los Alamos staff. Indeed, it was largely Teller's misrepresentation of the accuracy of his Super calculations of the 1940s to Norris Bradbury that had spurred Bradbury to deny Teller control over Los Alamos' thermonuclear program, placing another scientist at its lead instead. Teller's dishonesty with Bradbury, and Bradbury's response, cemented each man's views of the other, and spurred Teller's departure (after threatening to leave on many occasions) from Los Alamos. The AEC's General Advisory Committee initially rejected Teller's pleas to form a second weapons lab, but Teller's connections in the Congressional Joint Committee on Atomic Energy, or JCAE, (the powerful committee that oversaw the AEC and all civilian and military matters in the US concerning atomic energy and weapons) pushed the matter forward, resulting in the formation of the Livermore Radiation Laboratory in 1952.⁹⁷

⁹⁷ Westwick, *The National Labs*, 123-126; Ford interview with Carson Mark; Norberg interview with Norris Bradbury; Galison, "Computer Simulations and the Trading Zone," 135; Fitzpatrick, "Igniting the Light Elements," 164-165.

The immediate implication with the formation of Livermore was an implied slight against Los Alamos, one which the JCAE attempted to smooth over, but Teller's continued conduct further escalated the issue in the coming years. A longer-term implication was a competition, friendly or otherwise, between the two laboratories. The benefits of such a competition, cited by advocates of forming Livermore when approaching both the US Congress and the Department of Defense, were subject to fierce debate later in the 1950s, by which time the AEC was too late to overcome the institutional inertia of the new arrangement. With the two weapons labs firmly in place, their differing research and management philosophies resulted in distinct computing approaches and operations, due to the intimate ties between computing and weapons development. As a consequence, computing became an important arena where the labs differentiated from one another, intentionally and otherwise.⁹⁸

T Division's approach to computing developed out of complex and unusually intensive experiences with the limitations of traditional computational methods and technologies between World War II and the escalation of the US thermonuclear program in the early 1950s. In addition to fostering an interest and, by 1952, a commitment to general-purpose electronic computing at a point when the utility and practicality of early electronic systems were still unsettled, T Division's experiences with and impressions of electronic computers in relation to the division's responsibilities to Los Alamos' weapons program contributed to an approach that would remain the operational norm at Los Alamos for decades to come. That approach melded the computing services at the

⁹⁸ Westwick, *The National Labs*, 128.

Laboratory with R&D work meant to develop or refine computing technologies and methods for the potential benefit of Los Alamos. Furthermore, it was the lens through which the purpose and place of computing at the Laboratory would be viewed as the Cold War arms race intensified. The direct correlation between the speed and success of the thermonuclear program at Los Alamos with the power and availability of computing resources, specifically general-purpose electronic computing resources, left an indelible impression on T Division in the early 1950s, just as Los Alamos' place within the AEC lab system began to change amid heightened Cold War pressures and scrutiny.

The following chapter delves into the concurrent development of Los Alamos and its computing effort during the post-thermonuclear period, when the AEC weapons laboratories settled into their Cold War missions, and the general-purpose electronic computing field and industry began to mature, largely in conjunction with Cold War pressures and funding. Chapter Two contends that the interaction of these pressures, including Los Alamos negotiating its place within the AEC lab system, directly contributed to the development, and long-term social construction of the supercomputing category. That is, scientific computers designed to be orders of magnitude more powerful than the majority of other, contemporaneous computer systems, an arena where the AEC labs would dominate in funding and influence for decades.

-Chapter 2-

MANIAC to Stretch: Making the Computer "Super" at Los Alamos

I. Chapter Introduction

Between 1955 and 1961, the Theoretical Physics Division at Los Alamos Scientific Laboratory collaborated with IBM on the development of the IBM 7030, better known as "Stretch." The Stretch collaboration did more than produce a large scientific computer often retroactively labeled as a "supercomputer."⁹⁹ Stretch culminated LASL's decade-long transition from a mixed-computing environment of human, electronic and electromechanical, and homegrown stored-program computing methods and technologies, to an environment predominantly of IBM-sourced general-purpose systems. The transition reflected Los Alamos' increasing reliance on general-purpose stored-program computers as its primary computing technology, after a long, concurrent use of multiple technologies and methods. Neither simple nor "obvious," the transition reflected a multitude of technical and social factors that shaped Los Alamos computing for decades to come. Those factors included the maturation of the computing industry in the early 1950s, the goal within the Theoretical Physics Division of making general-purpose computing more widespread and accepted among sometimes reluctant users at LASL,

⁹⁹ Donald MacKenzie in "The Influence of the Los Alamos and Livermore Laboratories on the Development of Supercomputing" pointed out in 1991 the difficulty of pinning down when the supercomputing category began, due to the vagueness of what would qualify as a supercomputer, particularly before the late 1960s. The term "supercomputer" was not in widespread use until the 1970s, although the concept of a special grouping of extremely powerful scientific computers that stood apart from the majority of other, contemporaneous systems emerged in the 1960s in the US, UK, France, and the Soviet Union.

and changing research strategies at the weapons labs in the mid-1950s, born out of high-level concern for the labs' long-term futures.

Upon its completion, the IBM 7030 "Stretch" elicited dramatically different reactions from IBM and Los Alamos, with IBM considering the machine a failure and a setback, and Los Alamos considering it the near-ideal embodiment of a scientific computer. Utilizing extensive primary-source materials from the LANL archives and from numerous oral history interviews, this chapter examines how IBM's desire to dominate the scientific-computing market, and Los Alamos' distinct approach to computing, resulted in a multifaceted machine that shaped what would become "supercomputing."¹⁰⁰

¹⁰⁰ Bruce Boghosian provided in his "Computational Physics on the Connection Machine" article in *Computers in Physics* 4, no. 14 (1990) what might have inadvertently been the most broadly accurate definition of a supercomputer in any given time period when he stated: "A supercomputer is a machine that turns a compute-bound problem into an I/O-bound problem." Originally an old industry joke that Boghosian was repeating in his article, it accurately summarizes the major distinguishing features and enduring limitations of very large scientific computers, as supercomputers tend to exist at the limit of contemporaneous storage and memory I/O technologies, even more so than with processing technologies. The scientific-computing field and industry expends significant R&D resources on both hardware and software optimization to reduce the time a supercomputer's CPU (and, later, CPUs) spends waiting for its input/output devices to read and write data during normal use cycles. The operational time of a supercomputer tends to be extremely expensive (in terms of operator salaries, electrical power, cooling, floorspace, etc.), and time spent waiting on I/O can undermine the justification for purchasing a supercomputer in the first place. A diverse range of computers, from the Princeton IAS system to the Cray-1, have been labeled the first of the supercomputer category, usually depending on the strictness of the writer's definition, and the weight of personal interests over which machine received the accolade. Sources from IBM and Los Alamos have claimed Stretch to be the first supercomputer, despite similar claims for machines preceding and following Stretch. The May 4, 1961 issue of *LASL News* does refer to Stretch as a "super computer," with a space separating the two words. The article used the term "super computer" to describe Stretch as a computer, rather than place the machine into a category. The debate over which machine was the first supercomputer is less constructive for this chapter than recognizing the intent behind the Stretch development project. That recognition does more to elaborate on the new conceptions about high-end scientific computing and its potential applications that Stretch helped to foster within IBM, Los Alamos, and other advanced computing centers of the 1950s and 60s.

II. The T Division Approach to Computing

Los Alamos was among the earliest and most extensive users of general-purpose stored-program computers, and remained one of the world's most intensive centers of scientific computing between the 1950s and the 1970s. The Theoretical Physics or "T" Division operated most of the computing resources at LASL (Los Alamos Scientific Laboratory), and was largely responsible for establishing LASL's approach to computing as both a service to users, and as a field of research and development actively pursued at Los Alamos. That approach emerged from the informal practice within T Division beginning in World War II, and became formalized in 1949 when the division leader J. Carson Mark and Lab Director Norris Bradbury approved the formation of the T-7 Computer Research and Development Group. T-7's staff, at first consisting of seven members dedicated to hardware and ten to software (half of those ten were women), first developed and then operated the MANIAC computer. The staff members were divided into two subgroups; the first operated and maintained the computer, and developed tools needed for immediate use, while the other conducted more generalized projects, such as programming-language research with MADCAP. Nicholas Metropolis stated that T Division's computing research effort was modeled after von Neumann's group at the Princeton Institute for Advanced Study, just as the IAS computer served as a model for the MANIAC's design. The T-7 group's R&D efforts reflected a broader approach to research that was taking shape in the AEC labs. T Division's support of basic and applied, mission-focused research on general-purpose computing technologies developed

concurrently with the expansion of basic scientific-research efforts across the board in T Division, at Los Alamos, and in the AEC lab system.¹⁰¹

In the late 1940s and early 1950s, the AEC lab system overcame AEC leadership and GAC members' resistance to the inclusion of basic research. The AEC and its General Advisory Committee had feared mission-creep and a ballooning of budgets at the labs if they were allowed to add and expand on research programs not directly tied to programmatic work. The laboratories' struggles with remaining fully staffed, and particularly the crisis surrounding the Soviet atomic-bomb test in 1949, loosened the reins on the labs expanding their areas of research. Some of the AEC labs presented basic research as important for staff morale, while other labs, like Los Alamos, argued that basic research was an important tool for attracting and retaining top-quality staff. Those staff members, the Los Alamos argument continued, offered a "reservoir" of expert personnel who could be transferred from the basic research projects that had enticed them to the labs, to mission-critical programmatic work, if needed.¹⁰²

Norris Bradbury pushed Los Alamos to diversify its research efforts, as he hoped to minimize LASL's dependence on weaponry as the major reason for its existence. Bradbury's major challenge was to contend with the variety of staffing problems that faced early-Cold War Los Alamos. In addition to industry, LASL was in direct competition with universities for qualified staff (despite technically being part of the University of California system), and the Lab had the liability of being less flexible in the

¹⁰¹ Westwick, *The National Labs*, 154; N. Metropolis, "The Los Alamos Experience, 1943-1954," in N. Metropolis, *A History of Scientific Computing* (New York: ACM, 1990), 242-243; R. B. Lazarus, "MANIAC II," Edward A. Voorhees Papers, box 11, folder 4, series 17, LANL Archives, 3.

¹⁰² Westwick, *The National Labs*, 154.

research programs it offered, and of being a new kind of institution that offered few guarantees of long-term employment. The competition for staff naturally grew more heated with the formation of the Livermore Lab in 1952, despite AEC stipulations that the two were not to poach each others' staff. The limited availability of experts in fields like nuclear physics and electronic engineering made conflicts between Los Alamos and Livermore inevitable. The staffing conflicts compounded tensions stemming from the perceived insult at Los Alamos of Livermore's formation, and the general antipathy at the New Mexico Lab toward Edward Teller, both of which remained relevant at LASL in the coming decades.¹⁰³

¹⁰³ Westwick, *The National Labs*, 154, 156; Norberg interview with Bradbury; Edward F. Hammel, "Los Alamos Scientific Laboratory Energy-Related History, Research, Managerial Reorganization Proposals, Actions Taken, and Results, 1945-1979," Los Alamos National Laboratory, LA-13072-H (1997), 2-3, 59; LASL's overarching system of management during the Bradbury years (1945-1970) remained largely a continuation from the Oppenheimer years. Bradbury's most extensive changes to the Lab's administrative structure occurred in 1948, with the addition of three new Assistant Director positions for Classification and Security, Production, and Weapon Development. He also oversaw the formation of three new divisions: J Division, the Field Testing Division; GMX, the Explosives Division; W Division, the Weapon Nuclear Engineering Division. While J Division was a new entity, created to conduct the Lab's newly formalized live-testing program, GMX and W were formed mostly out of the Manhattan Project-era M and X Divisions. W and GMX Divisions quickly became the second- and third-largest users of LASL's computing resources, after T Division. Even after Bradbury's minor reorganization in 1948, remarkably little administration existed at Los Alamos for a laboratory of its size and complexity. Twenty-one division leaders still reported directly to Bradbury, with very little delegation to others in upper management, just as wartime division leaders reported directly to Oppenheimer. Oppenheimer asserted that his managerial strategy made it possible to maneuver the complex research requirements and capabilities of the bomb project quickly, as he did with several major reorganizations during World War II. Bradbury continued with Oppenheimer's managerial strategy for similar reasons during the Cold War (but without conducting major reorganizations), contending that twenty-one divisions reporting to him (there were a few technical divisions that reported to an assistant director, but they were the exception) were the maximum of which he could actively keep track. With this direct connection between the director and the majority of LASL's division leaders, Bradbury's instructions went directly to those most responsible for determining the Lab's research emphasis. As Westwick notes in *The National Labs*, it was largely division leaders who set the research agenda within the AEC, as they determined which projects to include in their budget requests to the Joint Committee on Atomic Energy. Individual research projects were usually lumped together into single, all-inclusive requests, making smaller projects, in particular, difficult to see or remove from the higher-level lab budgets. Although the AEC and the GAC had hoped to contain the budget allocations for basic research, the lumping of projects into division-level requests made curtailing research diversification nearly futile. The first major organizational restructuring at Los Alamos would not occur until 1970, when Harold Agnew replaced Bradbury as director. This provided a remarkably stable quarter-century of administration under a single director, under the same contract with the University of California, and under

T Division's dependence since World War II on leading figures in theoretical physics and mathematics—such as Hans Bethe, Enrico Fermi, Stanislaw Ulam, Edward Teller, and John von Neumann—made it among the LASL divisions most driven to attract top-quality experts, particularly after the mass postwar exodus from Los Alamos. An important part of T Division's recruiting effort was its willingness to support a wide variety of non-mission-related research programs, including projects of the staff members' own choosing. The most prominent example of T-Division-sponsored basic research was of physicists Frederick Reines and Clyde Cowan's Nobel-prize-winning co-discovery of the neutrino, an effort that began as a side project for Reines during the thermonuclear-development crisis in 1951. Allocating Reines programmatic funding from the Military Application Division, Norris Bradbury and J. Carson Mark allowed the scientist to choose a less-stressful side project as a diversion from the intense pressure of his regular work in the thermonuclear program. Management's goal was to prevent staff burnout during a critical development effort. T Division (and other science divisions) soon allowed staff members to devote up to 20% of their paid time to projects of their choosing.¹⁰⁴

From a staffing perspective, computing was no different from physics for T Division. Much of T Division's weapons-development research depended upon numerical methods rather than the empirical experiments that drove other LASL

the overarching auspices of the AEC. Los Alamos remained under the UC system until 2004; What is in 2018 Lawrence Livermore National Laboratory was founded in 1952 as the University of California Radiation Laboratory, but has undergone several name changes over its history. For the sake of clarity, it will be referred to as "Livermore" or "Lawrence Livermore" for the remainder of this dissertation.

¹⁰⁴ Ibid; Danny Cook, interview by Nicholas Lewis, Los Alamos, NM, August 13, 2015; Hammel, "Los Alamos Scientific Laboratory Energy-Related History," 3-5.

divisions. The most prominent example of LASL's empirical experimentation during the Cold War was its live testing of nuclear devices, but sub-critical and non-explosive testing also supported research efforts at Los Alamos. While live tests were the primary means of verifying weapon designs at Los Alamos, much of T Division's work required the computer-driven modeling of hydrodynamic and other natural phenomena too complex to be solved empirically, including the feasibility assessments for thermonuclear designs. The MANIAC's direct role in completing the first thermonuclear device in late 1952, after several years of stalled feasibility assessments and weapon-design calculations, cemented the place of general-purpose stored-program computing as a mission-critical tool at Los Alamos for T Division. It would be years before the broader scientific and engineering communities, and the less computing-dependent divisions at LASL, strongly embraced the general-purpose concept. The early evangelism of general-purpose computing that emerged from T Division staff and consultants, like Nicholas Metropolis and John von Neumann, stemmed from both a fascination with the state-of-the-art technology, and hands-on experience with its contribution to solving what, only a short time before, had been effectively unsolvable problems in hydrodynamics for nuclear-weapons development. Regardless of the warranted skepticism of the technology in the broader scientific community, actively promoting general-purpose computing R&D made particular sense for T Division.¹⁰⁵

Within LASL, T Division positioned computing as both a key technical resource and as a key area of fundamental research in the early 1950s, just as it and other divisions

¹⁰⁵ Owens, "Where are We Going, Phil Morse?" 35-36; Small, *The Analogue Alternative*, 2-8; Lazarus, "MANIAC II," 3.

positioned research reactors, particle accelerators, cyclotrons, or other high-dollar scientific equipment and instruments. These were resources that the AEC labs began to purchase and utilize in large numbers, especially after the emergency response to the advent of Soviet atomic weapons and the Korean War funneled increased levels of federal funding into the AEC labs. At the beginning of the crisis in 1950, the AEC allocated over \$80 million to research, the same year in which the newly founded National Science Foundation received \$350,000 total. In practice, the AEC labs often made little distinction between basic and applied research, given the popularly held belief within the AEC complex that the basic research of only a decade or so before had yielded the US nuclear enterprise. Believing that research was a "front line" defense for the future, the weapons labs allocated up to 15% of their programmatic funds to fundamental research efforts.¹⁰⁶

In a new field like general-purpose stored-program computing, T Division had great incentive and institutional support to pursue its own applied and basic R&D programs, both to develop solutions to problems few other scientific centers were yet facing, and to recruit and train experts in a field too young to have ready-made experts. During the same years when T Division brought the MANIAC into production use, and

¹⁰⁶ Westwick, *The National Labs*, 152, 157, 159; Metropolis, "The Los Alamos Experience," 242; High-dollar non-weapons research efforts at LASL that were funded (at least initially) with programmatic funds included the magnetic fusion program, which began in 1952 in the Physics (P) Division to research controlled thermonuclear reactors. Advanced nuclear reactor R&D began in 1954 to parlay Los Alamos' near unique expertise with plutonium metal into exploring the feasibility of plutonium-fueled reactors. The most famous (possibly infamous) non-weapons research program conducted at Los Alamos was Project Rover, beginning in 1955 to develop nuclear-powered rockets. The list of basic research projects, or projects being of only tangential connection to weapons development, that occurred at LASL only up to 1970 is extensive, and involved the construction of many of the largest facilities at LASL, including the Los Alamos Meson Physics Facility (LAMPF), and the Los Alamos Neutron Scattering Center (LANSCE), among others. Hammel, "Los Alamos Scientific Laboratory Energy-Related History," 3-5.

began to explore the capabilities of the new machine in both weapons and non-weapons applications, a computing industry began to take shape, largely due to military and other federally funded projects.

III. LASL and IBM

When T Division began its construction of the MANIAC in 1949, few general-purpose stored-program computers existed. The majority of those systems were one-off machines constructed for defense applications, like the ENIAC (after its conversion to partial stored-program capability), or as university projects, such as Cambridge University's EDSAC, the first full-scale stored-program computer. A number of copies, or near copies, of the Princeton Institute for Advanced Study computer were under construction in the late 1940s, including the MANIAC. The IAS-based machines were the first computers designed with the challenges of the AEC labs' thermonuclear-calculation work in mind, due to von Neumann's involvement and the AEC's support of the development program. The AEC likewise supported the construction of several other IAS-based computers to serve as advocacy tools to expand the range of "computable" sciences. The ENIAC was still being used through the early 1950s, but its resources proved too limited for the feasibility simulations that had stalled Los Alamos' work on thermonuclear weapons.¹⁰⁷

The ENIAC's principal designers, John Mauchly and J. Presper Eckert, formed the immediate predecessor company of the Eckert-Mauchly Computer Corporation

¹⁰⁷ Campbell-Kelly, *Computer*, 94, 98.

(EMCC) in 1946, the same year that Engineering Research Associates (ERA) formed in St. Paul, Minnesota out of a World War II US Navy code-breaking group. These companies released general-purpose stored-program computers, the UNIVAC I for EMCC, and the 1101 for ERA under Navy contract, in the early 1950s when almost no commercial competition existed and relatively little investment was required to enter the fledgling computer industry. The situation with the development of early mainframe computers was similar to the early personal computer era of the 1970s, when skilled enthusiasts created the systems they needed or desired, while also feeding into the formation (and often becoming part) of an industry to support the new technology. For potential mainframe customers of the late 1940s and early 1950s, there were few guarantees of when any computer-development project would be completed, or if a project would result in a working commercial system. Those realities influenced the decision in 1948 in Los Alamos' T Division to develop its own computer.¹⁰⁸

Remington Rand, IBM's most significant competitor in the US accounting-machine industry, purchased EMCC and ERA in 1950 and 1952 respectively to form its UNIVAC computing division, with a variety of UNIVAC-branded accounting- and scientific-oriented systems.¹⁰⁹ Most significant for Los Alamos, the notoriety the original UNIVAC had already gained in government and in the public eye prodded a hesitant IBM to enter the field of general-purpose stored-program computing. LASL's use of

¹⁰⁸ Campbell-Kelly, *Computer*, 107-108; Flamm, *Creating the Computer*, 45.

¹⁰⁹ The original UNIVAC computer was intended for data-processing activities, such as business and banking use. The UNIVAC I found use in scientific applications, which included Los Alamos borrowing time on one such system, but this was due to the scarcity of general-purpose stored-program computers in the early 1950s. The result was that scientific users found the original UNIVAC to be less capable for their purposes than machines like the ERA 1101 (later called the UNIVAC 1101, after the Remington Rand purchase of ERA) or the IAS-based computers.

IBM electromechanical and electronic punched-card equipment, as mentioned in the previous chapter, continued alongside the Lab's increasing interest in general-purpose stored-program machines. T Division's local reliance on the IBM equipment in the late 1940s was largely a pragmatic response to Bradbury and the AEC being unwilling to expand LASL's ranks of human computers significantly in the post-war years. The punched-card machines were a second-choice to the human option, but the IBM-sourced machine options were improving significantly for Los Alamos in the late 1940s. T Division played an active role in modifying and adapting the available technologies to its own purposes, which added to LASL's experience with scientific-computing equipment and methods during the thermonuclear crisis, and fed back into the nascent scientific-computing community and industry (in the case of IBM).¹¹⁰

A follow-on to the 604 Electronic Calculating Punch used at Los Alamos starting in 1948, IBM's 1949 Card-Programmed Electronic Calculator (CPC), a sophisticated, plug-board-programmable electronic calculator that employed 1,500 vacuum tubes, was the most powerful computing machine T Division employed locally before the MANIAC. The division ultimately installed six CPC units by 1953 at a monthly rental cost of \$1,600 each. Designed specifically for mathematical operations, the CPC was well-suited for hydrodynamics calculations where time limitations precluded the use of more capable off-site machines, like the SSEC or ENIAC. Like the ENIAC, the CPC was an important development tool for programming methods at LASL, providing a hands-on system for adapting the Lab's existing computational methods to machine processing. For example,

¹¹⁰ Fitzpatrick, "Igniting the Light Elements," 266; Campbell-Kelly, *Computer*, 98.

T Division wired a 604 control panel to create a special operating mode for the CPC. The new mode tracked decimal points automatically for users, making it possible for programmers to use floating-point arithmetic without the need to follow decimals manually. IBM's Technical Computing Bureau examined the LASL-modified CPC and developed a dedicated control unit that added T Division's automated-floating-point feature into new and existing units, which the head of IBM's CPC development project claimed had "revolutionized" CPC usage, and was thereafter used on most scientific jobs run on the CPCs. IBM later credited the CPC with preparing many of its scientific customers for machine-programming concepts years before the introduction of IBM's first stored-program computer.¹¹¹

Starting during World War II, and intensifying during the early Cold War years, LASL developed close ties to IBM representatives and engineers. Scientific customers from government labs and agencies, and from computationally demanding industries like aerospace, spurred IBM to develop its first dedicated, commercial calculating machines, the 604 and the CPC, even as the company remained reticent to pursue general-purpose stored-program systems. With the formal push to develop the H-bomb, T Division maintained a keen interest and engagement with IBM and its product line, even as LASL developed its own general-purpose computer. This was because even the general-purpose-computing evangelists in T Division did not expect or intend LASL's computing operations to move entirely to the MANIAC, hence the T-Division emphasis on the T-7 group as an R&D center to explore the capabilities of general-purpose computers. It was

¹¹¹ William J. Worlton, oral history interview with Bengt Carlson, Los Alamos, NM (October 30, 1969); Bashe, et al., *IBM's Early Computers*, 333-334.

also impossible for a single system to meet the Lab's computing demands, considering how quickly they were escalating with the development of thermonuclear weapons. IBM was one of the few sources of advanced calculating equipment then available to Los Alamos, and the one with which LASL was most familiar.¹¹²

IBM's reluctance to build a general-purpose computer began to crumble in light of several factors. One was that the CPC machines enjoyed what Cuthbert Hurd called "startling" popularity. Government agencies and labs ordered over 20 CPCs when first announced in 1949, with almost 700 entering service by the mid-1950s. IBM scientific customers also began to use or express interest in systems like the UNIVAC—LASL's Robert Richtmeyer borrowed time on NYU's AEC-owned UNIVAC I before the completion of the MANIAC. Coupled with the growing number of other first-generation computers being announced by other companies or research centers, Thomas Watson, Jr., son of IBM's CEO, in late 1950 began to take seriously plans for a general-purpose computer that Hurd and others at IBM had been drafting since July of that year.¹¹³

Hurd, a close friend of John von Neumann, had been a mathematician for the AEC at Oak Ridge National Laboratory, where he first used IBM punched-card equipment. He worked at Oak Ridge from 1947 to 1949, when he joined IBM as the founder and director of its Applied Science Department. The CPC was his first foray in the Applied Science Department with pressing IBM's capabilities in dedicated scientific-

¹¹² Bashe, et al., *IBM's Early Computers*, 62, 131; Cuthbert C. Hurd, "Computer Development at IBM," Nicholas Metropolis Papers, box 4, folder 32, LANL Archives, 12, 16; Although IBM machines like the IBM/Harvard Mark I and the SSEC were powerful scientific calculators, they were one-off systems not intended for commercial sale.

¹¹³ Hurd, "Early IBM Computers: Edited Testimony," 71-74.

computing machinery. Hurd's second foray inaugurated the company's first commercial general-purpose stored-program computer. Hurd remained consistently focused on advancing IBM's computing design and development capabilities over the 1950s, particularly in the face of Remington Rand's apparent headstart with the UNIVAC.¹¹⁴

Although the IBM 701 "Defense Calculator" was based on the IAS architecture, like LASL's MANIAC, the IBM system had significant deviations from the IAS computer. Modular vacuum-tube logic units meant ease of installation and maintenance, and improvements in I/O and architectural layout. One of the most significant changes in the 701 was a larger electrostatic memory of 2,048 36-bit words, twice the size of the original Princeton system, and moving to three-inch memory tubes, rather than the IAS computer's more troublesome five-inch tubes. The 701 system vocabulary was nearly identical to the IAS computer. One of IBM's most significant deviations from the original IAS design was to make the 701 compartmentalized in discrete units that were easy (relatively speaking) to ship, and to fit through doorways. These alterations were

¹¹⁴ Aspray, *John von Neumann*, 241-242; Hurd, "Early IBM Computers: Edited Testimony," 71-74; The UNIVAC I that Richtmeyer used at NYU was one of two UNIVAC I computers the AEC purchased. The UNIVAC at NYU was primarily intended to aid AEC-related computations, and to work on problems connected with NYU's Courant Institute, where the machine was based. The AEC also stipulated that 10% of the computer's time was to be set aside for basic research. The second AEC UNIVAC I went to Livermore for that Lab's programmatic research effort. Being a new lab, Livermore lacked its own equivalent of the MANIAC, and did not have the time to replicate T Division's in-house computer-development effort. Unlike Los Alamos, in which T Division, a science division, retained its control of LASL's computer resources due to its heavy computing requirements during World War II, Livermore formed from the beginning with a department dedicated to computing, its Computation Department, headed by Sidney Fernbach. The two AEC-purchased UNIVAC I computers, plus other UNIVACs around the US on which the AEC Labs borrowed time during the thermonuclear-development crisis, proved to be poorly suited for demanding hydrodynamic and neutron-transport problems, largely due to the UNIVAC I's slow acoustic memory systems and small 12-bit word size. The IAS-based systems, which utilized much-faster electrostatic-style memories, were approximately an order-of-magnitude faster for the weapons laboratories' applications than the data-processing-oriented UNIVAC I, leaving Livermore at a computing disadvantage in 1953 compared to LASL with the MANIAC, although the UNIVAC at Livermore became operational relatively quickly after arrival, compared with years of construction needed for the MANIAC.

largely beneficial from T Division's perspective, as most of LASL's large compute jobs were limited by either the memory capacity or I/O performance of a system. The larger the memory, the greater the complexity of the problems that could be run, or the fewer approximations or simplifications that were needed when coding the problem. Small memory size limited ENIAC's utility for Los Alamos, while the MANIAC's much-larger electrostatic memory (and drum-based secondary memory) made it the first system capable of thermonuclear feasibility simulations and fuel-mixture calculations elaborate enough to be conclusive. The 701 came with IBM's first commercial magnetic-tape storage system, which used a cellulose-based plastic tape, rather than the UNIVAC's more expensive bronze tapes, which were capable of lopping off poorly placed fingers.¹¹⁵

The 701 was IBM's first general-purpose stored-program computer and was a reflection of the changing re-orientation toward electronics within IBM. The formation of Hurd's Applied Science Department was a major step in that change, as IBM began hiring newly graduated electrical engineers, electronics majors, and applied scientists to develop expertise relating to the new field of computing, beyond IBM's traditional focus on electromechanical equipment. The process was similar to that in T Division at Los Alamos, which struggled to gather experts in these same fields.¹¹⁶

¹¹⁵ Bashe, et al., *IBM's Early Computers*, 134-135, 211; Worlton interview with Bengt Carlson; Hurd, "Computer Development at IBM," 15; Hurd, *Early IBM Computers: Edited Testimony*, 79; The 701's install-base had expanded greatly since Los Alamos received its first in March 1953, with thirteen systems delivered, and another six slated for delivery by early 1955. Although the install-base and user community for the 701 were small by later systems' standards, the interaction among the various users and installations was a significant asset of the 701. Their newsletter (begun in 1951) and exchange of utilities and experience formed the basis of the SHARE user group.

¹¹⁶ Hurd, "Computer Development at IBM," 15.

In September 1950, the IBM Naval Ordnance Research Calculator (NORC) project, a follow-on to the SSEC, began with Watson, Sr. and the IBM board's desire to aid in the US defense effort at the outset of the Korean War. At Hurd's recommendation, Watson, Jr. moved the 701 project forward in late 1950. The "Defense Calculator" moniker was an appeal to IBM management's patriotic drives (and/or its desire to build and maintain a strong business relationship with the DoD), attempting to maneuver around its reticence to produce general-purpose commercial systems. In 1951, Hurd, eager to bring von Neumann and his reputation to IBM's computing development effort, persuaded his friend to sign on as an IBM consultant. Offering one month of his time per year, von Neumann provided production, mathematical-feature, and architectural advice primarily on the NORC and Defense Calculator projects, which included his suggestion to add magnetic-drum secondary memory on the 701, much like the ERA-sourced drum used on the MANIAC.¹¹⁷

Los Alamos played a significant role in the 701's public unveiling. After J. Robert Oppenheimer served as guest speaker, Cuthbert Hurd ran a demonstration of the machine using an unclassified calculation Los Alamos provided. To much publicity from IBM and the AEC, the first 701 to leave IBM was installed at LASL in March 1953 on the first floor of "E" Building next to Ashley Pond. The 701 was the first general-purpose stored-program computer in the new T-1 IBM Computer Support Group, while the MANIAC remained under the auspices of the T-7 Computer Research and

¹¹⁷ Hurd, "Computer Development at IBM," 15-17; Bashe, et al., *IBM's Early Computers*, 134-135, 138-141, 164; Hurd, "Early IBM Computers: Edited Testimony," 71-76; Both the NORC and the SSEC are among the early computer systems that various industry insiders have retrospectively credited as being the "first" supercomputer.

Development Group. T Division was scattered across three buildings around the pond (at least until 1955), with the computing equipment on the first floor of E Building, and the groups of human computers occupying the second floor just above. T-1 formed in 1951 out of T Division's T-2 Computations and T-5 IBM Calculations Groups, which had housed the Lab's human computers and their machine counterparts, respectively. T-2's leader, the Swedish-born mathematician and Manhattan Project veteran, Bengt Carlson, transferred as the new leader of T-1. Carlson had joined the Manhattan Project in 1943 as part of the British Mission in Canada, transferring to Los Alamos from Montreal in 1945 with physicists J. Carson Mark and George Placzek to work on the neutron-transport problem. Carlson credited von Neumann with ensuring that Los Alamos became the first recipient of the 701, which LASL leased for \$15,000 per month. In late 1953, LASL received its second 701, out of only nineteen ever constructed.¹¹⁸

T Division began to prepare for the programming challenges of the 701 even before shipments of the machines began. Experience with the MANIAC and its similar, IAS-based architecture, and with customizing the CPCs before the MANIAC's completion, proved valuable in planning which programming aids, subroutine libraries, and other software tools were needed or advantageous for bringing the 701 into production use quickly, given the continuing federal pressure to develop and refine thermonuclear weapons. Many of the first operators and coders of the 701 at LASL had first joined T Division as human computers in the T-2 Computations Group between the

¹¹⁸ Bashe, et al., *IBM's Early Computers*, 162, 335; Worlton interview with Bengt Carlson; Worlton, "Hardware," 4; Willard G. Bouricius, "Operating Experience with the Los Alamos 701," in *Proceedings of the 1953 Eastern Joint Computer Conference* (New York: ACM, 1953), 45-47.

end of World War II and the early 1950s. Human-computers-turned-coders, like Marjorie "Marge" Devaney, Mary Menzel (later Tsingou), Mary Kircher, and Emily Willbanks, were recruited from the second floor of "E" Building to code for the MANIAC in Group T-7. Some of the MANIAC coders then transferred to T-1 to develop software for the 701 starting in 1952. The in-depth familiarity of coders and other staff with the MANIAC meant that LASL was the only 701 site to have extensive software and production experience with IAS-based systems before a 701 was installed locally. The RAND Corporation with its JOHNNIAC was the other site to have both a local IAS-based system and a 701, but the JOHNNIAC (1954) was not completed until after RAND's 701 arrived in late 1953.¹¹⁹

IBM was slated to produce its own collection of coding systems and utilities for the 701. Indeed, in August 1952, T Division participated in an IBM-hosted gathering of prospective 701 users, where IBM received suggestions for programming utilities customers desired for their upcoming systems. IBM's offerings for the 701 were only a few thousand lines of code in total, in strong contrast with the extensive software libraries IBM provided with later systems. T Division, like other 701 customers, set out to produce programming tools more tailored to its particular range of problems, just as T Division had done with its customization of the CPCs. As with the MANIAC, T Division sought to foster in-house expertise with the 701 that would ensure LASL maintained the

¹¹⁹ Edward A. Voorhees, "Recollections of the 701 at Los Alamos," *Annals of the History of Computing* 5, no. 2 (April 1983): 177-178; Marjorie "Marge" Devaney, oral history interview by Janet Abbate, IEEE History Center, Hoboken, NJ (April 4, 2002); Mary Menzel, oral history interview by Janet Abbate, IEEE History Center, Hoboken, NJ (April 4, 2002); Emily Willbanks, oral history interview by Janet Abbate, IEEE History Center, Hoboken, NJ (April 5, 2002); Aspray, *John von Neumann and the Origins of Computing*, 91; Hurd, "Early IBM Computers: Edited Testimony," 77.

capability to produce mission-critical software. Hardware support from IBM's resident support staff was a welcome benefit to T Division, because the 701's electrostatic memory and unreliable tape drives required frequent repair. T Division also wanted to use the 701 as a recruiting tool by commissioning cutting-edge software development projects for the new machine. T Division and LASL hyped the addition of a second (and later a third), world-class stored-program computer to the Lab's roster, as it was a resource that few other computer centers could claim in 1953. Several coders who worked on the MANIAC and then the 701 cited those cutting-edge machines as the primary reason they were drawn to work at Los Alamos.¹²⁰

Beginning in January 1952, over a year before the 701's arrival, T-1 assembled three programming sections to develop software tools and utilities for the new system. The assigned software projects were created based on T-7's experiences with software development since 1949. T-7's R&D effort had grown to emphasize the development of software that could take some of the programming burden off non-expert coders. Outside of a core, expert group primarily in T Division, non-expert coders constituted the bulk of potential users at LASL, a consideration that had to be balanced with developing programming aids that would not too heavily tax the limited system resources of the early computers.¹²¹

¹²⁰ Voorhees, "Recollections of the 701 at Los Alamos," 177-178; Bouricius, "Operating Experience with the Los Alamos 701," 45-47; Group T-1 701 Programming Section, "A Shorthand Coding System for the IBM 701 Calculator," Los Alamos Scientific Laboratory (June, 1953), Edward A. Voorhees Papers, box 9, folder 7, series 16, Computers and Computing, IBM 701 Principles of Operation, LANL Archives, 3; Bashe, et al., *IBM's Early Computers*, 334.

¹²¹ Metropolis, "The Los Alamos Experience, 1943-1954," 242-245; Voorhees, "Recollections of the 701 at Los Alamos," 178.

Based on the Lab's use of off-site general-purpose systems, its development of the MANIAC, and insights from collaboration with the RAND Corporation and the aerospace industry, T Division recognized that most prospective users at LASL were not yet comfortable with, let alone proficient at, using machine language (what T Division and others at the time called "longhand" coding). Many prospective users were also reticent to write their own subroutines, or otherwise invest large amounts of time to become experts on general-purpose systems. Given the time pressures surrounding the thermonuclear crisis, it was faster to send a possibly less-data-rich version of a problem to the human computers or the CPCs than attempting to code the problem for the general-purpose systems, whether onsite or offsite. One exception was T Division's own elaborate simulations that were often impractical for hand computation. For divisions and groups that relied more heavily on empirical methods and live testing than was possible for T Division, rough approximations in their calculations were often sufficient.¹²²

Where setup time for a problem was limited, or where approximations were adequate, scientists at LASL would typically default to traditional computing methods. This was not unique to Los Alamos. Historian of science Larry Owens noted that MIT,

¹²² Metropolis, "The Los Alamos Experience, 1943-1954," 242-245; Voorhees, "Recollections of the 701 at Los Alamos," 178; That is not to imply that simple approximations were not useful at Los Alamos. The most prominent example of a rough approximation being employed for weapons calculations at Los Alamos is Enrico Fermi's FERMIAC, which was an analog trolley and special paper trackway that used mechanical cogs and a list of random numbers to perform rudimentary Monte Carlo simulations for neutrons. The FERMIAC was built in 1947, when the ENIAC was being moved to Maryland and was temporarily unavailable for use. For a more detailed description of the FERMIAC and its uses, see Seidel, "From Mars to Minerva: The Origins of Scientific Computing in the AEC Labs." Los Alamos National Laboratory commissioned the production of an exact replica of the FERMIAC (exact down to a mislabeling found on the original) in 2017, and began the process of reverse-engineering how the FERMIAC was used, because no detailed instructions for its use have survived.

for example, had a "rule of thumb," in which problems that could be completed by hand in less than three "man-months," and would only be run once per year or less, would be sent to human computers. Programming and debugging were non-trivial expenditures of time with early general-purpose systems, which was why T Division invested in an R&D effort from the early 1950s onward to ease the programming burden, and to train users to become more proficient. Part of that effort was the intensive programming courses T Division staff began to offer for LASL personnel, which were taught at or near the Lab through the University of New Mexico. In addition to creating a library of subroutines for the MANIAC from which users could borrow, T-7 also developed over time a collection of interpretive systems, including assemblers, a high-level programming language (MADCAP), and "short-hand" coding schemes that, at the cost of computing performance, could ameliorate some of the complexities of "longhand" coding.¹²³

Expanding the roster of computer users within and beyond T Division, and developing a variety of tools optimized for LASL's production work pushed the broader agenda in T Division (and among other early evangelists) of improving general-purpose computing technology. In 1954, von Neumann, himself, as a member of the GAC, encouraged the AEC to invest even more heavily in computing technology, which put forward \$1-million per year for the purchase of those systems. The AEC also helped to establish several scientific-computing centers (in addition to those already at its labs and at NYU) across the US to promote computer use and development. Increasing the utility of general-purpose systems and enhancing their importance to LASL's programmatic

¹²³ Owens, "Where Are We Going, Phil Morse?" 36; Voorhees, "Software and Operations," 7-12.

work by engaging with new or otherwise hesitant users also helped to persuade LASL/AEC staff and management of the need to provide Los Alamos with the latest computing equipment and capabilities. This was an important position for T Division to promulgate in the early 1950s, when the live testing of rapidly changing weapons designs dominated LASL's programmatic approach, and general-purpose computers were only just making significant inroads as viable tools within the scientific community. Although those machines proved vital for T Division, few others at LASL had such extreme computing requirements.¹²⁴

Physics and physicists occupied a rarefied position at LASL after the Manhattan Project. They commanded much of LASL's funds, institutional direction, and esteem from the Lab's staff and supporting committees in Congress and the AEC. This was largely due to the perceived centrality of physics to LASL's past and then-current mission. Given the prestige of physics, it was not easy "selling" relevant cross-sections of the Lab, which was heavily focused on experimental methods and results, on new computing technologies and numerical methods. The reasons for this included the continuing perception of computing as a form of clerical work, and a resistance among many physicists to adopting new technologies, rather than employing more traditional calculating methods. Making general-purpose computers, which had a deserved reputation of unreliability and difficulty of use, both relevant and important within the experiment-oriented weapons community was a key battleground in T Division's effort to enhance the status and scope of computing at LASL. The long-term success and

¹²⁴ Voorhees, "Software and Operations," 7-12; Seidel, "From Mars to Minerva," 35-36.

consequences of that effort will continue to play out in the following chapters; it shaped the T-Division approach to developing and managing its computing resources in the early 1950s. Before T-1 had begun its in-house software-development program for the 701, T Division had already assembled a roadmap and guiding principles for such a program with T-7 and the MANIAC.¹²⁵

The 701 Programming Section, initially made up of six, and then nine people by the end of 1952, developed most of the system's utility programs. The section's single largest task before the 701 arrived was to develop the Los Alamos SHACO (Shorthand Coding) system, which made the 701 simulate a three-address, all-floating-decimal machine. With SHACO, users (or the coders who assisted in preparing their jobs) could prepare their problems in decimal numbers, while the computer automatically tracked floating points and monitored the locations of the program data and instructions in system memory, tasks that had to be performed manually when using an assembly language. SHACO-coded problems required about twice the processing time as problems coded in longhand, due to the added time for interpretation prior to execution. Problems written in IBM's own shorthand system, called Speedcoding, required about fifteen times longer for processing compared to the same problems written in assembly. T-1 estimated coding to be between ten- and thirty-times faster with SHACO, largely through the reduction of time spent debugging. Debugging had occupied nearly two-thirds of LASL 701 user coding time when the machine first arrived. Despite the performance drawback, problems coded with SHACO still ran nearly sixty-times faster on the 701 than similar

¹²⁵ Voorhees, "Software and Operations," 7-12.

problems on the CPCs. The difference in coding time provided a solid argument to potential LASL users in favor of transitioning (at least in part) to the general-purpose machines. SHACO was confined to smaller problems where performance was less of a concern, but significantly reduced the barrier-to-entry to general-purpose computing.¹²⁶

The second programming section dedicated to the 701, consisting of Bengt Carlson and Stewart Schlesinger, created another interpretive scheme, called Dual Coding. "Dual" allowed users to switch between fixed and floating-point in their problem, while still being able to call on "Print Memory," debuggers, and the 701's other existing utilities, because Dual was compatible with IBM's own Regional Coding system. Designed around the mathematical problems common to Los Alamos, Dual allowed users greater speed and flexibility in coding their problems than IBM's own system, while also retaining full access to pre-coded utilities that transferred information dynamically between the 701's electrostatic main memory and its magnetic storage, among other services. Dora Sweeney, who later worked for IBM, led the third section of five coders, which created LASL's "longhand" assemblers. Its 606 replaced IBM's SO2 assembler, which arrived with the 701, and its successor 607 was used for the remainder of the 701's service life.¹²⁷

About 35% of the problems that ran on LASL's 701s were coded in either Dual or SHACO, with the remainder (usually larger, more system-intensive problems) coded in

¹²⁶ Group T-1 701 Programming Section, "A Shorthand Coding System for the IBM 701 Calculator," 3-5; Bashe, et al, *IBM's Early Computers*, 338.

¹²⁷ "Regional Coding" referred to IBM's practice to divide the 701's memory into three main divisions, with one division each for input data, output data, and program code; Voorhees, "Recollections of the 701 at Los Alamos," 178.

either the 606 or 607 assemblers to make the best use of the 701's available resources.

According to T-1 staff member and program developer Willard Bouricius, LASL went from zero staff members being able to code for the 701 in January 1952, to twenty by the time the computer arrived in March 1953, with SHACO and the other LASL coding schemes bringing the number up to seventy-five by late 1953. As they became more proficient, due to the software tools T-1 developed, and the condensed 701 Coding Course that T-1 personnel taught to LASL staff members daily for ongoing, three-week intensive sessions, debugging 701 code dropped from an estimated two-thirds of the expended coding time to about half in less than a year.¹²⁸

T Division's efforts to produce proficient programmers among the scientific staff had another long-standing influence that shaped the division's approach to computer use and training. Ideally, T Division, or a scientist's own division, could have supplied coders to act as intermediaries, and translate a problem into the needed machine code,

¹²⁸ Bouricius, "Operating Experience with the Los Alamos 701," 45-47; "To Members of the 701 Class, Aug. 9-27, 1953," Edward A. Voorhees Papers, box 8, folder 8, series 16, LANL Archive; Voorhees, "Software and Operations," 9-10; The MANIAC and the IBM 701 provided roughly the same processing performance, but the 701's larger main memory made it better suited for larger, memory-limited problems than the MANIAC. Even so, the 701's system resources were still extremely limited, particularly in terms of the memory space available for variables in a problem, such as the dimensions represented in a simulation, or its resolution. As a consequence, although all simulations done at Los Alamos in the 1950s were in one dimension of space and with relatively few particles being represented, problems of any substantial size or complexity were typically written in an assembly language so as not to waste system resources on inefficient code or unnecessary software features. Using an assembly language versus using a "shorthand" interpretive system, depending on the problem, could mean the difference between a problem running and not running at all on a given system, as problems too large for main memory would have to move data between the fast onboard memory and a computer's much slower storage systems. The I/O cost of frequent data movement to and from the computer would effectively undermine the justification to have a powerful computer in the first place, as the system would spend much of its time idling in wait of data. Larger and more elaborate software utilities, like high-level programming languages and operating systems, exacerbated the issue, and were not practical for LASL's needs through most of the 1950s. The focus on processing performance and simulation detail and accuracy in scientific computing at LASL stood in contrast with data-processing oriented computing, which more quickly and readily embraced software utilities that made computers more user- and programmer-friendly, meaning that non-experts could program and operate a computer.

thus saving the scientist much of the burden associated with using early general-purpose computers. Such an arrangement would have eased much of the difficulty of incorporating early machines into the Lab's programmatic workflow, particularly if it could ease scientists who were non-expert coders, or were reluctant to begin to learn coding, into the use of general-purpose computers while those systems were still immature. Indeed, this was Livermore's primary approach, with a coder/mathematician assigned to every physicist who needed to construct a problem to run on a computer. As a consequence, Livermore scientists were better shielded from the more mundane and difficult aspects of early computer use. While T Division did have coders on staff who assisted scientists, and other divisions did hire coding assistants for scientists when needed and possible, the persisting Los Alamos housing shortage precluded the hiring of sufficient personnel to make such a service widespread. Accordingly, most Los Alamos scientists had to do their own coding, usually written onto a standardized coding form before a T-Division key-punch operator translated the written code onto punched cards.¹²⁹

While some issues, like delays in clearance processing for new hires, caused problems throughout the AEC, housing shortages were most acute for Los Alamos.

Housing in Los Alamos County was limited to the point that the Lab issued hiring freezes

¹²⁹ Abbate interview with Menzel; Mary Kircher, oral history interview by Janet Abbate, IEEE History Center, Hoboken NJ (2002); One of the more prominent instances of a dedicated "coder" being assigned to a scientific project at Los Alamos in the 1950s was of Mary Tsingou (formerly Menzel), who began working at LASL as a human computer, working with Enrico Fermi, Stan Ulam, and John Pasta to develop the first numerical experiment, the first experiment conducted entirely on a computer. On the MANIAC, in this case, but Tsingou later transitioned to coding for the 701. The problem related to the vibration of a string, and is widely considered an important early step in the development of Chaos Theory. The ability of T Division to dedicate a programmer to the experiment is not surprising, given the prominence of the scientists and mathematician involved. Such arrangements were also the exception.

at intervals during the early to mid-1950s, because there was simply nowhere to house newcomers. Some new hires, after being evicted from their temporary housing on the Hill once their allotted time had expired, commuted from Santa Fe or even further away (long, unpleasant drives in the 1950s). Given the dire housing situation, many prospective hires, particularly those with families, simply found employment elsewhere.¹³⁰

Housing issues caused particular concern during LASL's attempts to expand and diversify its lines of research in competition with Livermore, and the two competed for personnel and resources. Indeed, LASL plateaued and stabilized in the region of 3,800 employees between 1953 and 1959, after its population had risen sharply with the thermonuclear crisis. With only 625 employees in 1953, Livermore added approximately 500 new employees per year over the decade. Los Alamos County (which was still a closed, AEC-controlled community with no private ownership of property), responded to the housing crisis with new construction on the Los Alamos Mesa, and the reopening of the nearby White Rock bedroom community. White Rock had mostly housed construction workers in trailers during the Manhattan Project, and had closed after the war ended. Despite the new commitment from the AEC, construction took time, and could not keep pace with the pent-up housing demand.¹³¹

¹³⁰ Westwick, *The National Labs*, 79-80.

¹³¹ Housing shortages remain a problem in Los Alamos County in the early twenty-first century, particularly as LANL, as of 2018, seeks to increase its staff roster by several thousand new hires over the next several years; William J. "Jack" Worlton, oral history interview by Nicholas Lewis, St. George, UT (November 29, 2015); Memorandum from Henry R. Hoyt, Assistant Director, to Group and Division Leaders, "Housing, April 25, 1957," Edward A. Voorhees Papers, box 1, folder 15, series 1, LANL Archives; Memorandum from Bengt Carlson, T-1, to Carson Mark, "Housing for Computer Personnel," box 1, folder 15, series 1, LANL Archives; Westwick, *The National Labs*, 73, 240.

The housing problem impacted the computing operation significantly, as hiring authorization tended to happen only as housing vacancies appeared, and computing-related work was not the highest hiring priority outside of T Division. T Division faced shortages in a number of positions, particularly with programmers and computer operators. Programmers and operators were needed in large numbers, but recruiting staff from outside Los Alamos for what were generally low-paying, entry-level positions was difficult. In 1956, after the number of computer users at LASL reached over 250 (compared to only 70 two years before), T Division employed forty-six programmers and operators combined, at a time when LASL supported three major computer systems. Even after the hiring freezes ended in 1957, T Division could not consistently operate its computers on evenings and weekends (the preferred times to run large, time-consuming problems), because operators were sometimes difficult to schedule. A side effect of the operator shortages was that T Division was comparatively lenient in its hiring and work policies. T-Division preferred to hire the spouses or live-in relatives of existing staffers, as they did not require additional housing, and their clearance investigation would likely be easier and faster. T Division was willing to hire inexperienced operators (often with only a high-school diploma) and train them on the job, and, for both T-1 and T-7, allowed women to set their own schedules around the school hours of children and other life demands. Operators were also able to work as many or as few hours in a day or week as they desired. Half- and quarter-time operators were normal.¹³²

¹³² Anne Hayes, oral history interview by Nicholas Lewis, Los Alamos, NM (August 14, 2015); Abbate interview with Devaney; Westwick, *The National Labs*, 240; Memorandum from E. Voorhees, T-1, to C. Mark, T-DO, "Salaries of EDP Computer Operators, August 12, 1965," Edward A. Voorhees Papers, box 4, folder 15, series 9, LANL Archives; Seidel, "From Mars to Minerva," 35.

Operators also received free technical, college, and graduate training in programming, electrical engineering, computer science (after the formation of the CS field in the 1960s) and other fields from the University of New Mexico if they desired to move upward in the division, as some did. Women who worked as coders or operators experienced some measure of upward mobility in T Division, and could take advantage of the training opportunities available to all technical employees to become senior-level coders or operators, with the possibility of "converting" to a full staff position with better pay and benefits. Women required approximately five more years' experience to earn the same pay as men of the same "grade," the division's metric of relative experience, performance, and skill.¹³³

In addition to maintaining (with help from the onsite IBM technicians, whose service contract was part of the machine's lease agreement) and developing software for the 701s at LASL, T-1 also administered access to users. Usage of the 701s was on a "self-service" basis, where a user scheduled time with the T-1 dispatcher who allotted an assigned time for the user (typically between 20 and 30 users per day on the 701) either to debug or run their program. The operator who sat in the computer room then made sure that the correct user had access to the system at the assigned time, while making logs

¹³³ "Some" being the keyword for the upward mobility women experienced at Los Alamos in the 1950s and 60s. It was not until the late 1970s that a woman would become leader of one of the computing-related groups at Los Alamos, despite the computing-related groups and projects within T and then C Division having a larger representation of women compared with the norm at LASL, due to the early and growing presence of women in the computing field between its formation and the mid-1980s. Becoming a mid- to senior-level technical employee (rather than a full staff member) was usually the highest level a women who joined T Division as a computer operator could reach in the 1950s and 60s; Continuing pay discrimination towards women at Los Alamos was not directly addressed until a lawsuit sparked some changes to the imbalance in the early 1980s; A "staff" member was considered faculty within the University of California system, with the benefits associated with that position, which were considerable at Los Alamos; "Computer Operator Salaries, 1964," Edward A. Voorhees Papers, box 1, folder 15, series 1, LANL Archives; Voorhees, "Salaries of EDP Computer Operators, August 12, 1965."

for the computer's usage statistics. Users could either load their program on their own with punched cards, or provide a magnetic tape to the operator, who would load the tape for the user (unlike the MANIAC, the 701 did not have paper-tape read/write capability). Because the 701's tape drives were notoriously unreliable in their first iterations, most users opted to load from slower, but more reliable punched cards.¹³⁴

The self-service system originated with the CPCs and then the MANIAC, and grew organically out of personnel shortages, and T Division's place as both the operator and primary user of the Lab's computing systems. Access was, by default, typically hands-on for the small group of expert users who utilized most of the machine time at LASL. In late 1953, when the 701 had become LASL's primary workhorse for hydrodynamics calculations, T Division accounted for over 49% of the 701's total operational time, about 193 of the machine's almost 400 working hours in a given month. Other LASL divisions represented only about 14% of the 701's time, with non-LASL users at 20%, and maintenance occupying the other 17%. The MANIAC continued to be used for weapons work, but its time was also divided into a variety of other, non-programmatic research applications. Non-weapons research on the MANIAC included programming languages and the development of other forms of software (such as early work in AI with chess-playing programs), and time dedicated to LASL and outside researchers investigating the application of computers in genetics, musicology, astrophysics, early work on chaos theory, and other lines of research. By the time of the MANIAC's replacement with the MANIAC II in 1957, less and less programmatic work

¹³⁴ Bouricius, "Operating Experience with the Los Alamos 701," 45-47; Voorhees, "Recollections of the 701 at Los Alamos," 178.

took place on the home-grown machine, and the MANIAC II began its long service life primarily as a dedicated computer science R&D platform, still within the T-7 group. The growing range of users of the IBM equipment who were also developing and refining software tools, support from IBM itself, and the greater memory capacity of the 701, meant that the IBM line quickly became the focus of LASL's production computation.¹³⁵

After the proof-of-concept Ivy Mike detonation in late 1952, Los Alamos continued to press forward at a rapid pace to achieve practical and deliverable weapon systems between 1952 and 1954. T Division's role in the fundamental feasibility work of any critical assembly (which was only testable through numerical methods), and the greater complexity of thermonuclear designs meant that the MANIAC and the 701 were in high enough demand to warrant the lease of a second 701 in late 1953. Atmospheric live testing was LASL's primary method of verifying new and refined designs before they entered production for the US stockpile, with thirty-seven live US tests taking place primarily at the Nevada Test Site and in the South Pacific between 1951 and 1953.

Starting in 1954, a growing number of thermonuclear tests (including the infamous Castle Bravo test that irradiated large swaths of the Marshall Islands) added to an escalating

¹³⁵ Bouricius, "Operating Experience with the Los Alamos 701," 45-47; "Monthly Summary of 701 Time, August 1953," Edward A. Voorhees Papers, box 63, folder 1, series 48, 701 Machine Time, LANL Archives; Voorhees, "Recollections of the 701 at Los Alamos," 178; Worlton interview with Bengt Carlson; According to Bengt Carlson, most of the 701's non-operational time per month (of which there was a lot) was due to its electrostatic memory system. The general reliability of first-generation vacuum-tube computers at LASL, according to Carlson, rose dramatically with the switch to magnetic-core memories; A famous image features Nicholas Metropolis and Paul Stein playing chess with the MANIAC. To contend with the MANIAC's limited resources, the chess game was played with a simplified "anti-clerical" board, in which the bishops had been removed, and the board size shrunk to accommodate; Although MANIAC II was primarily used for CS R&D over its lifespan, there were researchers both at Los Alamos and elsewhere who preferred to use the MANIAC II even after it was no longer state-of-the-art, due to its easy-to-use programming environment, which included the MADCAP programming language featuring characters like built-in superscript and subscript, which made programming ideal for some scientific applications.

frequency of live tests in general, with 151 taking place between 1954 and 1958. The theoretical development work that primarily occupied T Division physicists, design and material-physics simulations from W and other divisions, and the need to process the larger volumes of data produced through a more rapid rate of testing, required a greater use of general-purpose computers. Indeed, while T Division's use of the 701s expanded from 49% to 56% of the available machine time by the mid-1950s, other divisions experienced an even larger jump, from only 14% to 36%, spread across nearly 1,000 machine-hours each month. Non-LASL users effectively vanished, as AEC and other research facilities received their own computers in the intervening years.¹³⁶

Although Los Alamos and Livermore were reliant on live tests for verification, the extreme forces of a detonating weapon made it difficult to monitor and understand what was happening inside the critical assembly as it destroyed itself, particularly as a conventional fission weapon, called the "primary," was used to generate the energy needed to produce a fusion detonation, the "secondary." Being able to view to a meaningful degree the interaction of variables like heat, pressure, and radiation on casing materials, and to gauge the characteristics of different fuel mixtures (a significant issue, as highlighted in the Castle Bravo test), on minute timescales and levels of detail required computer modeling even during live testing. Live tests were then used to calibrate what the computer-based simulations had predicted, which contributed to the next round of refinements and tests. In 1952, Bengt Carlson led T Division in the development of the

¹³⁶ Fitzpatrick, "Igniting the Light Elements," 166-167; "Monthly Summary of 701 Time, March 1956," Edward A. Voorhees Papers, box 63, folder 3, series 48, 701, 704 Machine Time, LANL Archives; *United States Nuclear Tests: July 1945 through September 1992*, DOE/NV-209, Rev. 14 (December 1994), vii.

"Sn" deterministic transport method to interpret the measurements taken from two critical assembly tests. First used on the MANIAC, the Sn method enabled T Division to model the single, observed outcome of a test, rather than a distribution of possible outcomes, as a Monte Carlo simulation would provide. This aided in T Division's reconstruction after the fact of a critical assembly's performance and behavior, and its ability to make adjustments before the next round of testing. The Sn method grew to become an expansive new field of deterministic modeling in physics and mathematics at and beyond Los Alamos in the decades that followed.¹³⁷

Developing new features, like safety systems, was an important aspect of weapons design work, but improving the efficiency of fissile-material use was the highest priority. Even at \$15,000 per month for a 701 in 1953, leasing a computer system was a bargain given the modest efficiencies in the extremely expensive fissile materials. High-speed computing had become inseparable from nuclear-weapons development to the point that LASL shut down its central hand-computing and CPC operations between 1954 and 1955, focusing entirely on general-purpose systems for its main production work. Other forms of computing, like analog, special-purpose systems remained in use for specialized scientific applications through at least the 1960s outside of T Division.¹³⁸

¹³⁷ Worlton interview with Bengt Carlson; Bengt G. Carlson, "Solution of the Transport Equation by Sulfur Neutron Approximations," Los Alamos Scientific Laboratory, LA-01599 (October 1953), 1-2; J. E. Morel, "Deterministic Transport Methods and Codes at Los Alamos," Los Alamos National Laboratory, LA-UR-99-3566 (1999), 1-3; The Castle Bravo incident in February 1954 was the result of misunderstanding how the particular fuel mixture in the device could result in a detonation several times the intended yield.

¹³⁸ Fitzpatrick, "Igniting the Light Elements," 166-167; MacKenzie, 186; Hurd, "Early IBM Computers: Edited Testimony," 77; Hurd, "Computer Development at IBM," 18; Office of Public Relations, "Automatic Digital Computers," Los Alamos Scientific Laboratory, Los Alamos, NM (March 1965), 3.

The second 701 arrived in December 1953, and, owing to the experience of installing the first system earlier in 1953, and the existing collection of IBM- and LASL-developed programming aids and LASL code libraries, the computer went into full production quickly, in early 1954. Several groups of LASL scientists and engineers had already completed T Division's intensive 701 programming courses. This significantly expanded LASL's 701 user base to over 70 trained users just as the political and research environment changed markedly after a live-testing incident placed Los Alamos in an unwanted spotlight. The Castle Bravo test of late-February 1954 was several times more powerful than intended due to a new and inadequately understood fuel mixture. The blast turned the test island in Bikini Atoll into a deep crater. Fallout rendered the rest of the previously evacuated Bikini Atoll uninhabitable, and irradiated other atolls in the Marshall Islands. The inhabitants of the affected islands were evacuated, but not until several days later. Fallout also irradiated the Japanese fishing boat *Lucky Dragon*, one crew member of which died from radiation poisoning. The incident, along with the continued US testing of multi-megaton hydrogen weapons in the South Pacific drew international concern and backlash directed at the AEC and Los Alamos specifically, and to the escalating US and Soviet atmospheric atomic tests in general.¹³⁹

The Castle Bravo backlash had multiple consequences for Los Alamos computing. The first consequence was a recognized need for more and better fuel mixture simulations before live testing occurred. The second consequence was Norris Bradbury's conscious effort to "broaden" LASL's research efforts to other forms of

¹³⁹ Westwick, *The National Labs*, 215-216.

nuclear power, fearing that a possible testing ban might have put Los Alamos scientists out of work. Bradbury had already feared that the successful development of a US atomic and thermonuclear stockpile in the 1950s might render Los Alamos unnecessary within five to ten years. It seemed difficult to "improve" upon multi-megaton weapons, aside from making gradual changes in features and efficiency. Bradbury pushed Los Alamos heavily into reactor development for power and propulsion as a means of leveraging LASL's expertise with plutonium for the sake of keeping experts employed in the event that weapons became a dead end.¹⁴⁰

The weapons labs had resisted the Project Sherwood fusion-reactor research effort that began in the AEC in 1951, but quickly embraced it with rising fears of a test ban and diminishing weapons research funding. Los Alamos also embraced space exploration and rocketry with Project Rover, beginning in 1955. Rover, an attempt to develop a surface-to-surface atomic-powered rocket, was also Los Alamos' introduction to research into power reactors, once the province primarily of Oak Ridge and Argonne among the AEC labs. Parts of the AEC later objected to LASL conducting power-reactor research, by which point the investment was too great and the research too established to be shut down.¹⁴¹

The transition at Los Alamos from the rapid-pace of the thermonuclear crisis to increasing concerns over safety and public image began to routinize weapons

¹⁴⁰ Norberg interview with Bradbury; Westwick, *The National Labs*, 79-80, 215-217; The growing interest in power-reactor research was also related to the passing of the 1954 Atomic Energy Act, which revised the limitations on commercial access to classified AEC reactor information. Public backlash against the Oppenheimer security hearings and the associated McCarthy-era communist witch hunts spurred the US Congress to revisit the AEC's sometimes extreme security requirements, including access to information relating to reactor development for peaceful purposes.

¹⁴¹ Westwick, *The National Labs*, 215-217.

development and testing, which meant that computing assumed a larger role in the Lab's mission. Weapon-yield estimates, for example, were a significant function of computer simulations before Castle Bravo, and became even more significant following the incident, requiring simulations of greater resolution and particle densities. The computers were also needed for the design and material-physics calculations behind the new range of power and propulsion reactor projects, expanding the range of computer uses and users at LASL. T Division was heavily involved with the mathematics and code-development process for the reactor-research effort. As part of that effort, T Division produced several new reactor codes (software packages that drove computer simulations) that soon entered regular use at LASL. The DSN transport code, for example, was designed to simulate spaces, like an infinite plane or cylinder, in one dimension. The TDC transport code was the first regularly used code at LASL to represent a two-dimensional space, such as a finite cylinder. Based on the earlier development of the Sn method in T Division, these transport codes were part of a growing collection of increasingly sophisticated simulation tools at Los Alamos as its computing effort and programming methods matured in the 1950s.¹⁴²

¹⁴² B. Carlson, C. Lee, and J. Worlton, "The DSN and TDC Neutron Transport Codes," Los Alamos Scientific Laboratory, LAMS-2346 (October 16, 1959), iii, 1-3; Bengt G. Carlson, "Numerical Formulation and Solution of Neutron Transport Problems," Los Alamos Scientific Laboratory, LA-2996 (November 11, 1963), 7-8; Limiting the number of dimensions and particles, and the represented time span in a simulation was necessary to keep from overwhelming a computer system by absorbing its runtime for unacceptably long periods of time. It was (and remains) possible to program a simulation unsolvable to computers decades hence, simply because of the scales and scope involved with the simulation. A two-dimensional code can radically increase the complexity of a problem, and the time needed to compute the results. The result was that T Division had to be cautious about how such codes were used, and the level of detail that could be incorporated into a problem using those codes. The nature of simulations and the codes used for constructing them will be addressed in greater detail in Chapter 3.

The addition of the second 701, and the rapid training of groups of LASL scientists through T Division's programming courses, meant that Los Alamos' computing operation was reasonably well-prepared for the transitional period of the mid-1950s. Carson Mark and the rest of T Division's leadership set an acquisition goal of doubling LASL's computing capacity approximately every two years, to ensure that computing would not become a limiting factor in the research at LASL, as had been the case at the onset of the thermonuclear crisis. Although LASL, in T Division's view, required a rapid escalation of computing power and more sophisticated codes and methods on an ongoing, regular basis to support LASL's mission, Los Alamos' computing operation was better prepared for the future than that of its main, non-Soviet competitor.¹⁴³

Los Alamos' first IBM 701 provoked a short-term computational-performance imbalance between Los Alamos and Livermore. The AEC ordered a UNIVAC I for Livermore in 1952, because the UNIVAC was among the few ready-made computer systems available close to Livermore's opening date in 1953. Installed in March of 1953, the same month as LASL's first 701, the UNIVAC I was primarily business-oriented, and left Livermore at a significant computational disadvantage. LASL's two IAS-based systems, which sported electrostatic memories many times faster and larger than the UNIVAC I's mercury delay-line technology, were better suited to the Labs' physics-modeling work, which was largely memory-limited. This disparity continued after Castle

¹⁴³ Norberg interview with Bradbury; Worlton interview with Bengt Carlson; Westwick, *The National Labs*, 215-217; Worlton, "Hardware," 5; A street in White Rock bears the name "Rover Boulevard" in honor of the project; A running joke at Los Alamos during the Cold War, recalled by a senior staff member at a LANL student conference, encapsulates the rivalry between the two labs: "The Soviets are the competition. Livermore is the enemy."

Bravo, when computing demands on the weapons-design process, and on the new range of energy-related projects, escalated at both labs. Livermore received the first of its two 701s in April 1954, but still lagged behind LASL with its two 701s and the MANIAC already in place, all while Livermore worked to bring its programmatic effort to full staffing and capacity to act as a competitor to Los Alamos. The disparity spurred something of an arms race in computing power between the two labs and their partners in industry.¹⁴⁴

IV. The Stretch Project

In late 1954, the Lab at Livermore was contending with the same goals of preventing further incidents like Castle Bravo and of diversifying its atomic-energy research. Like Los Alamos, Livermore ordered a range of IBM systems between 1954 and 1958, which were intended to meet the increasing computing demands of both weapons and atomic-energy research. At the same time, Livermore's staff and computer-user population grew rapidly. Working quickly to meet the short-term computing demand, Edward Teller, the Associate Director of Livermore, and Sidney Fernbach, the head of the Computation Department, also worried about their medium- and long-term computing requirements. Teller had recruited Fernbach, a former student of

¹⁴⁴ Westwick, *The National Labs*, 214-217, 277-279; Voorhees, "Recollections of the 701 at Los Alamos," 178; Norberg interview with Bradbury; Charles Folkner, oral history interview by Nicholas Lewis, Los Alamos, NM (August 13, 2015); A memory-limited problem is one that can be scaled to fill the available memory of a given machine. The more memory, the more detailed or expansive the computer simulation can be. Small-memory machines tend to constrain memory-limited problems, sometimes to the point that the results may be less than useful for drawing meaningful conclusions about the problem at hand. It was largely the ENIAC's limited, vacuum-tube-based memory that precluded conclusive results from the feasibility calculations Los Alamos ran on it.

Oppenheimer at Berkeley, away from a postdoc position in physics at Stanford to manage the acquisition of Livermore's UNIVAC and the recruitment and training of its operators and programmers. He became head of the newly formed Computation Department shortly thereafter.¹⁴⁵

Hoping to push industry to produce a system more in-line with Livermore's projected computing demands, Teller and Fernbach created a Request for Proposals (RFP) for the LARC (Livermore Automatic Research Computer), which they distributed to over a dozen companies in the nascent computer industry. Among the more significant companies invited to bid were International Telemeter Corporation, IBM, and Remington Rand (which became Sperry Rand after merging with the Sperry Corporation in 1955). Remington Rand's UNIVAC Division, having acquired ERA in 1952, had just delivered the ERA-designed, core-memory-based 1103 to the NSA.¹⁴⁶

¹⁴⁵ Pugh, *Memories that Shaped an Industry*, 160-162; Marilyn Ghausi, oral history interview with Sidney Fernbach (February 8, 1989). Accessed October 3, 2016. <http://www.computer-history.info/Page1.dir/pages/Fernbach.html>; In 1958, Fernbach also became the Acting Theoretical Physics Department leader, after its original leader died in a plane crash. Much like Norris Bradbury holding his "interim" directorship of LASL for two decades, Fernbach remained the department head for twenty years, while still leading the Computation Department until the early 1980s.

¹⁴⁶ Pugh, *Memories that Shaped an Industry*, 162; Arthur Lawrence Norberg, *Computers and Commerce: A Study of Technology and Management at Eckert-Mauchly Computer Company, Engineering Research Associates, and Remington Rand, 1946-1957* (Cambridge, MA: MIT Press, 2005), 253-254; Bashe, et al., *IBM's Early Computers*, 248; Livermore would not begin researching power reactors until 1961, and on a much smaller scale than Los Alamos. According to Peter Westwick in *The National Labs*, p. 218, Los Alamos immediately invested about 12% of its annual operating budget (\$46.4 million in 1956) into reactor research in the mid-1950s, with about half of that 12% (\$2.8 million) going to Project Rover. About one-quarter (\$1.5 million) went to power reactors. Livermore, which had a much smaller budget in 1956 (\$22 million), still invested over 21% of its annual budget to reactor research, with two-thirds of the allocation going to Project Sherwood. Livermore's budget would, by 1960, greatly exceed that of Los Alamos (\$66 vs. \$58.4 million), a disparity that would only get wider by the 1970s. This budgetary factor would influence the acquisition of computer resources and the allocation of support among the respective staffs of each lab. The disparity fueled no small resentment at Los Alamos toward Livermore, as Los Alamos' designs represented the overwhelming bulk of the weapon systems in the US nuclear arsenal, despite having a smaller budget than its competitor.

The ERA-UNIVAC 1103 was comparable in performance with the IBM 701, but had superior memory reliability, even if that memory was small and not particularly fast. Remington Rand continued to sell its EMCC-based UNIVACs to business customers, much as IBM sold the 701 to scientific customers and its 702 variant to business customers. At that point, IBM did not command a leading position in general-purpose computing as it did in the accounting-machine industry, in large part due to its late start in electronics. Remington/Sperry Rand's strong presence in the business sector with the original UNIVAC, and its leap into scientific computing with the acquisition of ERA, spurred several responses from IBM. The IBM 650 drum-based computer, announced in 1953, was a significant contender in the low-cost data-processing space, while the AEC complex, as the largest single purchaser of scientific mainframe computers in the US, was a high-end target in the computer market. As the companies vied for market share, the AEC nuclear-weapons labs exercised considerable influence on the emerging industry.¹⁴⁷

Teller and Fernbach's RFP called for a two "mega-pulse" computer, meaning an internal clock that operated at two-million cycles per second. They also stipulated that the machine should use transistors rather than vacuum tubes. The two mega-pulse metric implied the use of transistors approximately twice as fast as vacuum tubes. Teller and Fernbach's goal was to provide Livermore with a powerful scientific computer by 1957 or 1958 that could meet that lab's computing demand as its staff count grew, while also enabling more sophisticated calculations. The Livermore request was intended to

¹⁴⁷ Pugh, *Memories that Shaped an Industry*, 162; Bashe, et al., *IBM's Early Computers*, 248.

produce a machine that could push computational boundaries, but not pose a development challenge so significant that the computer failed to materialize. Specifying the use of transistors was somewhat risky (although experimental transistorized computers had existed since 1953), given the significant technical challenges in reducing their cost while improving their performance and reliability. The lead time in the RFP for a completed system was intended to give time for those challenges to be addressed.¹⁴⁸

The transistor was not a new invention in the mid-1950s, but the development of a reliable, relatively inexpensive transistor technology was a significant challenge for the electronics industry. It then became a problem for computer manufacturers to secure a sufficient supply for transistor-based computers, as the US military absorbed most of the transistor supply in the earlier part of the decade. The US military had also funded much transistor R&D with the generous funding arrangements typically found with military contracts.¹⁴⁹

¹⁴⁸ Ibid; Bashe, et al., *IBM's Early Computers*, 425; Note that computer performance was difficult to quantify and compare in the 1950s, as there were no standard benchmarks or metrics among manufacturers or users. Attempts to compare performance were usually convoluted, and based on assumptions about what constituted normal workloads for a given computer user or center. Assumptions about "normal" workloads could vary wildly, depending on whether a given problem was memory-, compute-, or I/O-limited, and even if a user's workload relied more heavily on one particular operand than another (addition vs. division, for example). By the 1960s, T Division at LASL primarily focused on operations per second (OPS) as its metric, and successfully helped push the supercomputer community to accept FLOP/S (floating-point operations per second) as the standard metric of scientific-computer performance, thanks to LASL's C-3 Benchmarking Group producing industry-circulated benchmark guides. The 1950s and 60s were a different situation. LASL during that time typically measured performance in relation to other systems in use at the Lab, such as gauging a system in powers of the MANIAC, or the later IBM 7090. The MANIAC I was, for example, about 0.10x the power of the IBM 7090. The performance measures of computers will become more consistent in later chapters, as more and better direct comparisons become available from the documentary record.

¹⁴⁹ For a detailed discussion of the role of military contracts in the development of the Silicon Valley tech sector, including the DoD heavily subsidizing the transistor during its first decade, see: Thomas J. Misa, "Military Needs, Commercial Realities, and the Development of the Transistor, 1948-58," in Merritt Roe Smith, ed., *Military Enterprise and Technological Change* (Cambridge, MA: MIT Press, 1985), 253-87, and Christophe Lecuyer, *Making Silicon Valley: Innovation and the Growth of High Tech, 1930-1970*

The Livermore RFP reflected Teller's belief that transistors were the most viable means of supplanting the size, power, and packaging problems of first-generation vacuum-tube computers. The RFP sparked IBM's interest, with the company putting together two proposals. The first was a simpler, two-megapulse design using well-tested but obsolete Philco "surface-barrier" transistors. The IBM team stressed the second, more daring proposal, which was based on concepts from Gene Amdahl's Datatron Project, a conceptual system outline meant to put forward new, advanced features and capabilities. Planning to stretch IBM's existing technical capabilities, particularly with faster "drift" transistors, the proposal Amdahl, Cuthbert Hurd, and the rest of the IBM group pushed at Livermore was for a system approximately five-times more powerful than specified in the bid, with a larger and faster memory (two fast memory systems, in fact), and advanced architectural features, like the ability to "lookahead" and preload upcoming instructions. The IBM proposal was part of that company's goal to become a dominant force in scientific computing, with the LARC potentially providing a halo product for showcasing IBM's design prowess.¹⁵⁰

Teller and Fernbach, not interested in waiting, were seeking a computer that would be available sooner rather than later. The IBM proposal, in light of its advanced features, had requested at least a three-year development period. Instead, Livermore awarded the LARC contract to Sperry Rand, which had proposed a system to be delivered in only twenty-nine months. Although the LARC project was far from a simple

(Cambridge, MA: MIT Press, 2006); Flamm, *Creating the Computer*, 15-18, 78-79; Bashe, et al., *Early IBM Computers*, 417.

¹⁵⁰ Pugh, *Memories that Shaped an Industry*, 160-162; Bashe, et al., *IBM's Early Computers*, 424-426.

or profitable undertaking for Sperry Rand, the proposed design featured well-tried surface-barrier transistors and a slower core-memory specification than the IBM proposal. That made Sperry Rand a safer choice for Livermore, as fewer technologies would have to be dramatically improved or developed entirely in a limited timespan than with IBM.¹⁵¹

The Livermore rejection in favor of Sperry Rand concerned IBM management, particularly Thomas Watson, Jr. (the president and soon-to-be CEO), and spurred a refinement of the Datatron concept to potential customers as a means of securing for IBM a dominant place in the scientific-computing market. IBM proposed the refined plan for a ten-megacycle computer to the NSA in June 1955, but the NSA declined the complex project. Negotiations resulted in NSA R&D funding to support high-speed core-memory development. IBM began its search for a customer with a willingness to take a calculated risk on an advanced computer. More importantly, that customer also had to support a large proportion of that system's development. Los Alamos, with its long history with IBM, its own experience with computer hardware and R&D projects, and its AEC funding, was one of the few clear possibilities.¹⁵²

In the fall of 1955, Los Alamos began searching for a new computer system that could meet its projected computing requirements through the late 1950s. T Division was amenable to a longer development period, if it meant receiving a more capable computer

¹⁵¹ Bashe, et al., *IBM's Early Computers*, 424-426; Norberg, *Computers and Commerce*, 253-254; Flamm, *Creating the Computer*, 110; Sperry Rand's LARC contract with Livermore stipulated \$2.85 million, while the final machine was closer to \$19 million. The machine's development also exposed schisms between the EMCC and ERA contingents of Sperry Rand's UNIVAC division, and helped to foster the flight of ERA alumni, including Bill Norris, who left to form the Control Data Corporation.

¹⁵² Pugh, *Memories that Shaped an Industry*, 164.

at its conclusion. That was particularly true if the longer development period also meant T Division could participate in designing a system tailored to the Lab's requirements. As important as the IBM 700-series computers already were to LASL in the mid-1950s, systems that were purchased rather than built locally lacked some of the Lab-specific refinements and features that characterized the MANIAC and the upcoming MANIAC II. Vendors typically did not know what the applications for their general-purpose computers would be, which also made predicting and gauging performance difficult before a system was put into use.¹⁵³

The most important refinements for the Los Alamos constructed computers were the built-in instructions tailored for lab use (an area where the MANIAC deviated from the IAS system), and the close integration between hardware and software development that T-7 had emphasized from the beginning of the MANIAC Project. Purchasing (or leasing, in the case of IBM) systems was, indeed, more cost-effective than building in-house systems by the mid-1950s. From T Division's perspective, purchasing systems rather than building them in-house had come at the price of inadequate vendor software and inadequate software integration with the system itself. LASL computer users often complained of vendor software (when it existed) lacking in flexibility, which spurred much of T Division's custom software development. The opportunity to help guide the system's software development in conjunction with its hardware could provide T Division

¹⁵³ "Technical Evaluation of Proposals Submitted to Prepare an Advanced Computer for LASL, April 2, 1956," Edward A. Voorhees Papers, box 16, folder 6, Stretch Solicitations and Selection, series 22, LANL Archives; Worlton interview with Carlson; Stephen Dunwell, OH 153, oral history interview by William Aspray, Poughkeepsie, NY (13 February, 1989). Charles Babbage Institute, University of Minnesota, Minneapolis, <http://hdl.handle.net/11299/107272>.

with the best of both worlds: IBM-built hardware with LASL-influenced software and architectural integration.¹⁵⁴

IBM's informal proposal in late 1955 to gauge LASL's interest in supporting such a development struck a chord within T Division, in part because of the IBM representatives' willingness to collaborate with LASL. In repeated meetings, the Sperry Rand representatives expressed little interest in a collaboration with LASL. The same was true in Sperry Rand's meetings with Livermore's Computation Department on the LARC. That arrangement, according to Fernbach, worked well for Livermore, as the Computation Department did not desire a collaboration on the LARC's development, largely due to Livermore's lack of experience with computing hardware. Instead, Fernbach's department and potential users made requests to Sperry Rand of features they desired, but the company did not alter the computer more than the system designers were willing to accommodate. Despite T Division's interest in collaborating, Los Alamos could not simply enter a contract with IBM for the proposed computer. The AEC required LASL to submit its own RFP, and go through a formal bid process from vendors, as Livermore had done for the LARC. T Division wrote and submitted its RFP for an "Advanced Computer" to the AEC in early January 1956.¹⁵⁵

¹⁵⁴ "Technical Evaluation of Proposals"; Worlton interview with Carlson.

¹⁵⁵ Ibid; W. H. Brummett, Jr., and H. G. Kolsky, "Summary Report of Contract Panel Meetings and Visits to Vendors, Contract Panel for Purchase of Advanced Type Computer for LASL, March 1, 1956," Edward A. Voorhees Papers, box 16, folder 6, Stretch Solicitations and Selection, series 22, LANL Archives; Ghausi interview with Fernbach; It should be noted that the Oppenheimer Security Clearance Hearings took place in 1954, roughly at the same time Livermore began its computer search for the LARC. As Norris Bradbury and other contemporaneous LASL staff members recalled, this was a high point for tensions between Los Alamos and Livermore, as Teller's testimony contributing to the revocation of Oppenheimer's clearance was widely considered to be a betrayal at Los Alamos, where Oppenheimer was, and still is, held in high regard.

The LASL RFP closely resembled IBM's informal proposal, requesting a machine with a 2-microsecond main memory (meaning the computer could access a word held in memory in as little as 2 microseconds) of at least 50,000 50-bit words, built-in floating-point capability, the inclusion of fast I/O equipment, and an instruction execution time of between one-half microsecond and two microseconds. Multiplication and division, for example, required longer execution times than addition and subtraction. Most of the specified features were given cautious wiggle room for performance and means of execution, and some were not necessarily required, but simply requested. For example, T Division included in its RFP a second, faster core-memory system for fast loops in code (an original feature in the IBM Datatron concept), but simply stated that such a feature "may be desirable."¹⁵⁶

The RFP called for the system to be delivered by late 1959, and for a price of \$3.5 million. As with the LARC, the AEC wanted a fixed-price contract. IBM had initially wanted a cost-plus-\$1 contract for the LARC, similar to that the company had secured from the US Navy for the NORC. IBM's hope was that a cost-plus contract with the AEC could offset development costs of the new system. Instead, with only \$3.5-million from the AEC, IBM hoped that volume sales of the new computer could help recoup money lost (IBM estimated it would lose about \$10 million, but ultimately spent closer to \$35 million) to develop the system for Los Alamos. The LASL RFP did not include the

¹⁵⁶ "Los Alamos Computer Proposal, January 9, 1956," Edward A. Voorhees Papers, box 16, folder 6, Solicitations and Selection, series 22, LANL Archives; "IBM Proposal for a Multiplex 10 Megacycle Computer, Oct. 1955," Edward A. Voorhees Papers, box 16, folder 6, Solicitations and Selection, series 22, LANL Archives; Bashe, et al., *IBM's Early Computers*, 424; "Company Details," Edward A. Voorhees Papers, box 16, folder 6, Solicitations and Selection, series 22, LANL Archives.

performance estimate IBM provided in its first informal meeting at Los Alamos, which Ed Voorhees described as "100 or more times faster than previous computers." Various figures at IBM, including Watson, continued to cite this and similar rough performance metrics during the computer's development (usually in reference to the IBM 704), which Voorhees' notes indicate appeared early in the proposal process.¹⁵⁷

Between January and April of 1956, seven members of T Division, including Bengt Carlson, Max Goldstein, Harwood Kolsky, Roger Lazarus, Nicholas Metropolis, Ed Voorhees, and David Woods, plus a representative from the Lab's Accounting Office, comprised the Technical Advisory Committee that reviewed proposals from vendors. Four vendors responded to the RFP, including the Reeves Instruments Corporation, IBM, Sperry Rand, and RCA. Only RCA, after initial interest, opted not to submit a proposal, citing the \$3.5-million price as being unrealistic. RCA estimated about a \$10-million loss in development costs for such a computer, which was in-line with IBM's predicted losses.¹⁵⁸

Because T Division drafted its RFP with IBM's informal proposal in mind, there was little surprise that IBM's written proposal in March of 1956 was the only one to match LASL's request. Sperry Rand was second closest, proposing both a copy of the LARC, and a slightly faster LARC variant at an unspecified higher price. Reeves proposed a system far below the specifications that T Division had requested, and for a higher price. In April 1956, the T-Division Advisory Committee recommended, as

¹⁵⁷ "Los Alamos Computer Proposal, January 9, 1956"; "IBM Proposal for a Multiplex 10 Megacycle Computer, Oct. 1955"; Bashe, et al., *IBM's Early Computers*, 424.

¹⁵⁸ "Technical Evaluation of Proposals"; Brummett, and Kolsky, "Summary Report"; Pugh, *Memories that Shaped an Industry*, 165.

expected, acceptance of the IBM proposal to the AEC. In his report to the AEC of the Committee's recommendation, Carson Mark emphasized, out of his six major reasons for choosing IBM, the fact that representatives of that company had more interest in a LASL influence on "the machine logics" than did Sperry Rand. Mark stated that the Committee considered LASL involvement in the design process to be a significant advantage, and an opportunity to "improve the final machine for LASL needs."¹⁵⁹

The AEC accepted IBM's proposal, stipulating a machine to be delivered in forty-two months, and for \$4.3-million, a higher price that IBM had negotiated from the AEC. With the new contract came a new name for the development effort, becoming the "Stretch Project" to signify the technological stretch it represented to IBM. Formal planning meetings began in early September 1956, where the initial six members of the Los Alamos Stretch Committee, Bengt Carlson, Harwood Kolsky, Roger Lazarus, David Woods, Roger Frank, and Edward Voorhees traveled to Poughkeepsie, New York, to meet with the fifteen primary IBM representatives of the Stretch Project. The first concern was over the need to obtain "Q" clearances for the IBM reps, before the combined group could discuss classified information. Keeping to unclassified discussions before the clearances were completed, the Los Alamos committee, with Bengt Carlson as its chairman, brought forward a collection of issues that LASL users had noted about IBM's existing systems and wanted rectified with the new computer, such as providing a larger number of index registers to make resource management easier for programmers. The LASL group also wanted IBM to report any expected changes to

¹⁵⁹ "Technical Evaluation of Proposals"; Brummett, and Kolsky, "Summary Report."

the system instruction set to T Division before being implemented, in order to determine if the changes would be detrimental to LASL's production-code development for Stretch. T Division was already planning development of a Stretch emulator for the IBM 704, so that programmers could begin and test code work before Stretch was completed.¹⁶⁰

The majority of the first formal IBM/LASL Stretch planning meeting, spanning three days, focused on the nature of LASL's mathematical challenges and requirements, emphasizing I/O strategies for the new system. Roger Lazarus, for example, provided a neutron-transport equation which was common to Los Alamos problems, but posed a particular issue with I/O performance. The assembled group outlined a rough I/O configuration and system vocabulary that could improve the I/O throughput for transport problems. Harwood Kolsky did much the same for hydrodynamics, outlining how hydrodynamics equations functioned, and the issues they posed for computer systems. It was from that early discussion on I/O that Kolsky assembled a twelve-page formal breakdown of the major I/O problems as they existed at Los Alamos for the potential system design. The LASL argument called for an I/O "philosophy" for Stretch that was based on the Lab's experiences with existing IBM systems. The philosophy related to how the computer's architecture and external devices should be organized, designed, and implemented to ensure I/O was tailored to the scientific computing problems at LASL. While the IBM 700-series comprised LASL's production workhorses, I/O handling on those systems was not well-suited to scientific problems, as I/O devices were often slow,

¹⁶⁰ "Technical Evaluation of Proposals"; Pugh, *Memories that Shaped an Industry*, 165-166; D. W. Sweeney, "Stretch Mathematical Planning Meeting, 9/11/56," Edward A. Voorhees Papers, box 16, folder 8, STRETCH Math Planning Meeting, Sept. 1956-Aug. 1957, series 22, LANL Archives.

and system calls were poorly timed/managed for efficient operations. Another cause for T Division's concern over I/O performance was the increasing speed at which faster computers would generate data, and the related problem of main system memories growing in size to many thousands of words.¹⁶¹

The throughput speeds of computer I/O systems (primarily punched cards, impact printers, magnetic tapes, and magnetic drums) were not keeping pace with system memory capacities and processor output speeds. This was a problem for scientific-computer users on several fronts, including the simple act of loading program data from storage into a computer's memory. Problems could vary in their dependence on I/O access by as much as 1,000 times, so a computer had to be flexible in its handling of widely varying I/O requirements, while still offering adequate throughput on all jobs. Another significant I/O performance issue was the need to transfer the contents of the system's main memory to storage, later known as a "core dump," either to recover a job's progress after a potential computer fault, or to resume an ongoing job at a later time. Tape drives, for example, provided ample storage in aggregate, but offered slow throughput that would leave the computer sitting idle for extended periods while the transfer took place. Magnetic drums provided better read/write speeds, but insufficient storage space for the contents of a large memory. Kolsky's outline formally documented the possibility of utilizing IBM's new magnetic, moving-head hard disk drive technology

¹⁶¹ Sweeney, "Stretch Mathematical Planning Meeting, 9/11/56"; Harwood G. Kolsky, "Project Stretch: Some Preliminary Thoughts on Input-Output Problems," Edward A. Voorhees Papers, box 20, folder 2, STRETCH notes, reports, correspondence, series 22, LANL Archives.

(first found on the RAMAC computer) as an intermediary between fast memory and permanent tape storage.¹⁶²

The IBM RAMAC (Random Access Method of Accounting and Control) accounting computer introduced the first commercial hard disk drive, the IBM 350, in 1956. According to Kolsky, reading a single tape reel into a computer's memory required *6.6 minutes*. The RAMAC's disk drive could read the same amount of data in four *seconds*. Write speeds were always slower than read speeds, but multiple hard disk drives and read/write heads per drive could be utilized in parallel to increase the disparity between hard disk and tape storage performance whether reading or writing. A significant reduction in Stretch's idle time while it waited for I/O operations would reduce operating costs, and increase the productivity of the machine's users. Resolving the most pertinent I/O throughput issues also helped LASL to justify the computer development project itself, as there was little point in constructing a machine with Stretch's computational capabilities if it remained idle for an excessive amount of its operational time. An early result of the Stretch collaboration, the repurposing of the hard disk as an intermediary storage system, remains an underrecognized development in the history of high-performance scientific computing.¹⁶³

The collaboration process continued later that same month at the second meeting, also over three days, with Stephen "Red" Dunwell, an electrical engineer involved with the design of most of IBM's early calculating machines, requesting greater elaboration on

¹⁶² Ibid; John N. Savage, "Understanding Stretch: A Primer on Automatic Digital Computers," Los Alamos Scientific Laboratory, Los Alamos, NM (June 15, 1961), 5.

¹⁶³ Kolsky, "Project Stretch: Some Preliminary Thoughts on Input-Output Problems."

problem types. The Los Alamos group then outlined key mathematical problems (in their unclassified form), such as Monte Carlo simulations and matrix problems, and brainstorming with the IBM staff about potential solutions that could be integrated into the computer. The most significant results from the second formal meeting were, aside from the move to a 64-bit word length (with an additional eight parity bits), to "stay away" from numbers at that preliminary design stage, and Dunwell's suggestion to break into smaller groups in future meetings. Gene Amdahl, a key system architect of the early Stretch proposals, whose newly shipped IBM 704 design was meeting with great praise from customers, left IBM in 1956 after Dunwell was appointed head of the Stretch project instead of himself.¹⁶⁴

The co-ordinated planning meetings continued through April of 1959, alternating between Los Alamos and IBM facilities. AEC delays in officially approving the Stretch contract meant that the machine's delivery date had also slipped to May of 1960, to the advantage of the development effort. As Dunwell had suggested in the second meeting, the larger contingents broke into six smaller groups, focusing on performance, input-output, instruction format and indexing, arithmetic operations, automatic programming (programming aids and languages), and multiple precision. Willis Ware noted the early progress of the Stretch collaboration after visiting IBM Poughkeepsie in late 1956. An electrical engineer, Ware completed his PhD at Princeton in order to work with von Neumann on the IAS computer. He joined the RAND Corporation in 1952, where he

¹⁶⁴ "Transport Equation," Edward A. Voorhees Papers, box 16, folder 8, STRETCH Math Planning Meeting, Sept. 1956-Aug. 1957, series 22, LANL Archives; "2nd Stretch Meeting with IBM, Sept. 19-21, 1956," Edward A. Voorhees Papers, box 16, folder 8, STRETCH Math Planning Meeting, Sept. 1956-Aug. 1957, series 22, LANL Archives; Bashe, et al., *IBM's Early Computers*, 431.

aided in constructing the IAS-based JOHNNIAC, named after von Neumann. Ware visited IBM to investigate the status and viability of the Stretch Project for RAND. During his visit, Ware met with the manager or group leader of each major technology or system under development for the Stretch computer. He noted that the "status of the engineering requirements of the Stretch program is in excellent condition," concluding that "there is no question in my mind that the engineering phases with high probability will comfortably make the delivery dates required." The one deficiency he perceived was that Los Alamos and IBM had, as of late 1956, only two people assigned to the system's functional design. That deficiency vanished quickly, with the ballooning, for better or worse, of Stretch Project staff at IBM, reaching over 300 by 1959.¹⁶⁵

In early 1957, the Stretch Project was transferred from IBM's research organization to product development. Over that year, the IBM/LASL group solidified the system's architecture and instruction set, with Stephen Dunwell sending progress reports of those planning meetings to the AEC Los Alamos Field Office. IBM concentrated on developing and integrating the needed hardware technologies with the system design, and LASL's Mathematical Planning Group focused on software design for the new computer. As Willis Ware noted, IBM had multiple lines of development for each major technology needed for the computer. That strategy was to avoid becoming overly committed to any given technical path. That approach worked well for much of Stretch's development. For

¹⁶⁵ Bashe, et al., *IBM's Early Computers*, 432-433; "Stretch, October 5, 1956," Edward A. Voorhees Papers, box 16, folder 8, STRETCH Math Planning Meeting, Sept. 1956-Aug. 1957, series 22, LANL Archives; Willis H. Ware, "Summary of Trip to the International Business Machines Corporation, Poughkeepsie, New York, 10-19-56," Edward A. Voorhees Papers, box 16, folder 6, Solicitations and Selection, series 22, LANL Archives; Willis H. Ware, OH 356, oral history interview by Jeffrey R. Yost, RAND Corporation, Santa Monica, California (11 August, 2003). Charles Babbage Institute, University of Minnesota, Minneapolis, <http://hdl.handle.net/11299/107703>.

example, the transistor fabrication process was modified several times to improve production speed and unit reliability. To increase the speed of development, IBM used computer-aided design (on the 704) whenever possible, the first time computer-aided design was used in the development of another computer system. Test assemblies were used when computation could not readily provide feedback on component design, like on portions of the arithmetic unit.¹⁶⁶

In addition to the early transfer of the system to IBM's product development line, several other factors brought growing problems to IBM's Stretch group over the next several years. The most significant were directly related to pressures from IBM's management. Management exerted pressure on the Stretch group to standardize component design across IBM's other product line (despite causing compatibility and component production issues). The group also experienced pressure to expand the system design to suit a wide range of users, upsetting existing design elements, and making the system more complex and prone to technical delays. Tensions developed between IBM management and the Stretch design staff, in addition to those from LASL and the AEC, as emergent technical issues conflicted with stated performance goals of the new system. Most notably, IBM's agreement in 1957 to develop simultaneously a version of Stretch for the NSA, called "Harvest," complicated Stretch's design, and delayed its completion and delivery between the late 1950s and early 1960s.¹⁶⁷

¹⁶⁶ Pugh, *Memories that Shaped an Industry*, 437; Ware, "Summary of Trip to the International Business Machines Corporation"; S. W. Dunwell, IBM, to Paul Wilson, Area Manager, Los Alamos Field Office, "Meeting of Members of the Joint LASL-ISB Mathematical Planning Group, August 8, 1957," Edward A. Voorhees Papers, box 17, folder 6, IBM Progress Report, May 1957-Dec. 1959, series 22, LANL Archives; Hurd, "Early IBM Computers," 86.

¹⁶⁷ Pugh, *Memories that Shaped an Industry*, 437.

The software development effort in LASL's mathematical planning group tended to have fewer stumbling blocks than the hardware development. The hardware and software design efforts remained apprised of each other's progress through the regular collaboration meetings, where the two also negotiated refinements in both hardware and software implementation as the delivery date drew closer. During that process in 1957, Harwood Kolsky, the LASL physicist, T-Division associate group leader, and member of the Mathematical Planning Group, moved to Poughkeepsie to aid in the design of Stretch's "Lookahead" unit, and eventually joined IBM for the remainder of his career. Originally a feature of Amdahl's Datatron concept, the Lookahead unit greatly reduced I/O latency by preloading into memory the upcoming instructions needed for a program and ensuring that the required system resources had been allocated before a process called for them. Recruited by Dunwell, Kolsky worked with IBM physicist and computer scientist John Cocke to determine the number of "levels" of lookahead that would offer the best performance, while minimizing the complexity of the Lookahead unit. Running simulations on a 704, Kolsky and Cocke discovered that four levels, that is, loading up to four upcoming instructions while simultaneously running a fifth, were optimal for LASL codes. This process also determined that the second, faster memory system found in the Stretch proposal was not needed due to the performance advantages of the lookahead unit.¹⁶⁸

¹⁶⁸ Voorhees, "Software and Operations," 13-14; "MCP Personnel," Edward A. Voorhees papers, box 26, folder 3, STRAP and MCP Working Papers, series 22, LANL Archives; J. L. Garrity, G. R. Hira, J. F. O'Connell, and C. A. Folkner, "STRETCH Master Control Program," Edward A. Voorhees papers, box 26, folder 3, STRAP and MCP Working Papers, series 22, LANL Archives; Eric Knutsen, *Wasn't That a Time: Stretch/Harvest Retrospectives*, Harvest/Stretch Reunion, September 28 and 29, 2002, Poughkeepsie, NY (Riverside, CT: LookBack Publishing, 2002), 31, 110-111.

The LASL group spent much of the late 1950s designing Stretch's longhand assembly program, STRAP, and the computer's operating system, MCP (Master Control Program). The two were then coded in a collaboration between IBM and LASL. MCP was T Division's first effort at developing a full operating system, which was a relatively new concept in the late 1950s. Limited computer resources up to that point had placed severe constraints on the size and complexity of software controls that could also be placed on a system. While T Division and its computer users typically resisted added software that could significantly impact performance or resource availability, the complexity of systems like Stretch, and the pressures of managing an increasing workload of user jobs at the Lab, meant that T Division needed greater flexibility in its computer management capabilities. Other prospective Stretch customers would also benefit from the better control and management software tools that an operating system offered. MCP grew, in part, out of T Division's experience writing a relatively elementary "monitor" program for the IBM 704, which automated the queueing and output of multiple batch jobs preloaded by an operator. Having a monitor meant that the owner of a job did not need to remain in attendance while the job was being executed. The automated system streamlined the overall process of submitting smaller jobs (the majority of jobs run at Los Alamos) and having them returned. MCP served a similar purpose, managing large numbers of smaller jobs to be run primarily during the day.¹⁶⁹

Management of small jobs was of particular importance for a system like Stretch. A "small" job at Los Alamos typically required fewer than twenty minutes to complete on

¹⁶⁹ "MCP Personnel"; Garrity, et al., "STRETCH Master Control Program"; "Stretch, October 5, 1956."

an IBM 704. On Stretch, the same job would shrink to less than a minute: too short a time for individual users to load and execute a job efficiently on their own in a self-service fashion. Large, more resource-intensive jobs were to be run primarily at night, and without an operating system, in order to maximize system resource availability. Unlike the 704 monitor, MCP offered automatic fault detection and recovery, and extensive operator feedback and control of the system. The sophisticated automation and operator control of Stretch stood in contrast with the rudimentary software tools and literal hands-on operator control found with the MANIAC.

In addition to developing an operating system sufficiently versatile for all potential Stretch customers, the goals behind MCP were both technical and organizational for Los Alamos. The technical goals were to minimize system downtime due to computer and programming errors, and to reduce wasteful idle time when running the most common and voluminous jobs at LASL. Continued staffing shortages also made ease of operation a priority. The easier it was to operate Stretch, the less training and expertise potential operators required, easing T Division's task of maintaining a sufficient number of operators to keep Stretch operational around the clock. IBM's onsite maintenance staff were invaluable for resolving technical issues. MCP's organization and control features were designed to take advantage of Stretch's built-in recovery, lookahead, and high-speed I/O capabilities, and were negotiated as part of the collaboration process at the joint LASL/IBM meetings. The meeting notes and technical reports on MCP documented the influence of T Division's experiences operating and using the MANIACs and the IBM 700-series computers on MCP's feature set. A large

part of that influence came from feedback T Division had requested from users about the capabilities needed in the system software package, including the programming aids developed by LASL and IBM. T Division had, as Carson Mark suggested in his 1956 committee's recommendation to accept the IBM contract, embraced Stretch as an opportunity to develop a system oriented around LASL's computing needs.¹⁷⁰

In addition to STRAP and MCP, T Division independently designed and coded for Stretch a programming language, COLASL (Compiler of LASL), and a "load-and-go" assembly language, IVY. COLASL and IVY were used extensively at Los Alamos, but found little traction elsewhere. IVY was not geared for use outside LASL, because it was tailored to LASL-specific jobs. For COLASL, the long development period of Stretch meant that, by the time the 7030 (Stretch's official IBM designation) was shipping to LASL and other customers, FORTRAN, first released in 1957 for the IBM 704, had a head start of between four and five years among other scientific-computing users and centers. Although eventually available on all of LASL's IBM systems (available in 1962 for Stretch), Los Alamos computer users largely avoided FORTRAN until the mid-1960s, due to perceived performance deficiencies in the compiled FORTRAN code over assembly languages.¹⁷¹

¹⁷⁰ "MCP Personnel"; Garrity, et al., "STRETCH Master Control Program"; "Stretch, October 5, 1956"; Housing shortages were not as dire in the late-1950s and early 1960s at Los Alamos as had been the case in the early-to-mid 1950s. The housing problem never completely disappeared for LASL, even after the AEC permitted private ownership of property starting in 1963. T Division had to contend with operator and programmer shortages through the 1960s. Tools designed to reduce the difficulty of operating T Division's computers, despite ongoing objections to the resource utilization of large software packages, like operating systems, were a necessity as LASL's computing operation grew in size, and user demand increased in the late 1950s.

¹⁷¹ Voorhees, "Software and Operations," 13-14; Wells, "Evolution of Computer Software," 3.

When LASL users began to adopt FORTRAN, it was typically in the form of LRLTRAN, a Livermore-modified version of FORTRAN. LASL users typically used the STRAP (which co-compiled for the 704) and later STRAP-II assemblers when coding for more intensive jobs on Stretch. In the collaboration's division of labor, IBM developed FORTRAN for Stretch, while T Division developed the machine's mathematics-subroutine library, allowing users to draw from a precoded selection of subroutines (similar to that T Division created for the MANIAC in the early 1950s), rather than developing their own while coding. The LASL software-development effort for Stretch did not emerge as an entirely new endeavor for T Division, but as part of the Lab's own computing R&D program, its experiences with its own evolving software requirements, and T Division's interaction with the small but growing scientific-computing community.¹⁷²

LASL grew increasingly invested in IBM's scientific computer platforms across the 1950s. At the time of the Livermore RFP, IBM had recently announced its vacuum-tube-based 704 computer, which was a more powerful variation on the 701 (although the two systems were not intercompatible), but with built-in floating point hardware, and magnetic ferrite-core memory.¹⁷³ The 704, which was slated to begin shipping in 1955,

¹⁷² Voorhees, "Software and Operations," 13-14; Wells, "Evolution of Computer Software," 3; "Computing Requirements of the ERDA Advanced Computer Laboratories," Edward A. Voorhees Papers, box 1, folder 17, Historical Articles, series 2, LANL Archives.

¹⁷³ Pugh, *Memories that Shaped an Industry*, 160-163; Ferrite-core memory was made up of small iron rings with control wires threaded through them. The control wires set and read the direction of magnetization of each ring, representing a zero or a one saved in memory. Core memory grew out of MIT's Project Whirlwind, which grew into the US Air Force's SAGE early warning system. IBM provided the computer hardware that powered SAGE, including the magnetic core memory that replaced the electrostatic memory common to many first-generation computers. The primary advantage of core memory over previous technologies was reliability. Bengt Carlson at Los Alamos noted that the transition from the 701 to the 704 marked an enormous increase in computer "up time," due to memory issues disappearing

was developed in response to the success of the 701. With Gene Amdahl as its principal designer, the 704 represented IBM's commitment to become more than just another player in the scientific computer market, as Big Blue planned to produce the 704 in large quantities (ultimately 140) for a mainframe computer, compared with only nineteen copies of the 701. The 704 was IBM's technical and corporate commitment to be a leader in the new computer industry. What made the 704 stand out was the inclusion of traditional IBM service tailored to individual customers, and features meant to appeal to its most demanding scientific users, including hardware-based floating point, making the 704 the first commercial system with that feature. John von Neumann, as a consultant for IBM, actually discouraged the company from including built-in floating point on the 704, as he believed it to be a waste of system resources.¹⁷⁴ At Los Alamos, Jack Worlton joked, "I suppose if we were all as brilliant as von Neumann, we wouldn't need floating point!"¹⁷⁵

Between late 1955 and late 1956, LASL acquired three IBM 704s to replace the two 701s as its main production systems. The acquisition represented a significant investment and commitment for LASL beyond the machines themselves, as it meant that T Division had to redevelop existing codes and utilities from the incompatible 701. The 704's superior performance, reliability, and memory capacity meant that T Division could develop more sophisticated programming aids, such as the FLOCO and FLOCO II "load-

almost entirely. Core memory remained in common use into the 1970s, particularly in supercomputing, as it offered significant performance advantages over early semiconductor-based memories until around 1975.

¹⁷⁴ Worlton interview with Bengt Carlson; Bashe, et al., *IBM's Early Computers*, 384, 390-393; Hurd, "Computer Development at IBM," 27-29.

¹⁷⁵ Worlton interview with Bengt Carlson.

and-go" assembly programs. These assemblers allowed a program and its data to be loaded onto the computer simultaneously, with the computer being able to run parts of the program even before all of the program and its data were loaded onto the system. The heavily used DSN and TDC transport codes, for example, were both written in FLOCO. Experience developing the FLOCO assemblers, and the resulting advantages in programming speed and ease of use (relatively speaking) they offered to LASL computer users, contributed directly to the IVY "load-and-go" language for Stretch. T Division based Stretch's COLASL programming language and compiler, which supported natural algebraic notation, on the ALGAE programming language the Division developed for the 704.¹⁷⁶

The 704s and the MANIAC II were the first T-Division computers installed on South Mesa, across the new "Omega" Bridge in TA-3, in the new Administration Building. LASL dismantled most of the original Technical Area buildings during the late-1950s. Hasty wartime construction meant that most of the original TA-1 buildings were no longer adequate or safe to use. Still wholly owned by the AEC, the Los Alamos townsite became an open community in 1957, but, out of necessity, LASL's computing operation had always been engaged with the outside world. That engagement began with the need to find or borrow computing power during the Manhattan Project and early Cold

¹⁷⁶ "Chronology of LASL Computers," Edward A. Voorhees Papers, box 67, folder 8, LASL Computer Inventories 1961-1973, series 49, LANL Archives; Bengt G. Carlson, Chester S. Kazek, Jr., and Clarence E. Lee, "FLOCO II Manual," Los Alamos Scientific Laboratory, LAMS-2339 (August 1959), 1-2; "Computing Requirements of the ERDA Advanced Computer Laboratories."

War years. The MANIAC was itself a product of sharing system development plans and strategies with other computing sites.¹⁷⁷

A community of users developed around the IBM 701 in the mid-1950s, in part because of IBM-organized user installation meetings. Coding utilities first used on the 701 at Los Alamos were the first pieces of software made available to users at other 701 sites. Many centers with a 701 continued to focus on locally developed software solutions, causing redundancy in software development across those sites. The announcement and release of the 704 spurred a more active community, starting with an IBM-initiated gathering of the seventeen centers planning to acquire a 704. At that IBM

¹⁷⁷ "Family Days, July 16-17, 1960," Edward A. Voorhees Papers, box 2, folder 2, Family Days, 1955, 1960, 1965, series 5, LANL Archives; The AEC's decision to open the Los Alamos townsite was controversial among some residents, according to the local newspaper of the time. While easier for guests of residents to enter the town, some residents feared that outside crime might enter as well. Most residents seemed to welcome the transition to a more conventional living arrangement, as the controlled access to Los Alamos brought a heightened sense of isolation to many who lived there in the 1940s and 1950s; The original MANIAC was not moved across the new bridge, out of fears that it might become a fire hazard, given its delicate construction. That fear of the machine catching fire did not prevent LASL from donating the MANIAC to the University of New Mexico, where it remained in use for several years; The MANIAC II was the beginning point at which LASL could not develop systems in-house that were a close match of those sold commercially, particularly not for the hardware-development and production costs involved. This was due to the massive investments companies like IBM received largely through defense contracts to develop systems for military applications, like the SAGE Project, the technologies and systems developed for which made their way into IBM's product lines. The 704 was a consequence of this largely DoD-funded technical and organizational maturation in the computer industry. While the MANIAC II featured about five-times the performance of the MANIAC I, and about ten-times the memory capacity, its built-in floating point capability was an ad-hoc system added after the 704's floating-point system was announced, in order to boast similar capabilities as the 704. The addition of magnetic core memory was, likewise, an adjustment to the MANIAC II's original design, based on the published details of the 704, as the new LASL computer was originally slated to feature an updated electrostatic memory. According to Nicholas Metropolis, the decision to build the MANIAC II was not out of necessity, but out of the realization that, if T Division did not build the new system when it did, there would be little justification to build one later, once the maturation of the computer industry had completely rendered locally developed systems outmoded. Developing the new system provided T-7 with both a platform and official sanction to continue predominantly as T Division's in-house R&D group. In 1963, T-7 transistorized the MANIAC II, which began a process of reinventing the MANIAC's hardware on almost an annual basis to reflect the latest lines of hardware R&D taking place in the computing industry until the machine's decommissioning in 1977; Atsushi Akera, *Calculating a Natural World: Scientists, Engineers, and Computers During the Rise of U.S. Cold War Research* (MIT Press, 2006), 254.

gathering, representatives of the prospective 704 centers agreed to meet in a formal user group of their own making, hoping to work together to standardize and share many of the basic utilities that the new system would require. With Los Alamos and Livermore in attendance, the first meeting took place at the RAND Corporation in August 1955, with the group dubbing itself "SHARE" as a description of the group's purpose: sharing software, experience, and knowledge of the 704 among other group members. Membership of SHARE grew to 62 before the end of the first year, and the voluntary programming assignments distributed within the group during the first meeting resulted in a growing collection of over fifty programs during that year. SHARE expanded around the growing number of IBM 704s, and then around other IBM computers released over the decades, with SHARE becoming a model of user-group interaction and exchange for other user communities.¹⁷⁸

A founding member of SHARE, Los Alamos later joined the user groups for every major system it acquired, playing an active role in acquiring and distributing software among those groups. While the classified pieces of LASL's software base remained out of public circulation, unclassified science and simulation codes, utilities, mathematical routines and libraries, and other software repositories were freely

¹⁷⁸ Lazarus, et al., "MANIAC II," 3, 7, 9; Worlton interview with Bengt Carlson; Bashe, et al., *IBM's Early Computers*, 347-349; Paul Armer, "SHARE—A Eulogy to a Cooperative Effort," *Annals of the History of Computing* 2, no. 2 (April 1980): 125-127. For additional information about SHARE and computer-user groups and services, see: Atsushi Akera, "Voluntarism and the Fruits of Collaboration: The IBM User Group, Share," *Technology and Culture* 42, no. 4 (2001): 710-736, William F. Vogel, "'The Spitting Image of a Woman Programmer': Changing Portrayals of Women in the American Computing Industry, 1958–1985," *IEEE Annals of the History of Computing* 39, no. 2 (April-June 2017): 49-64, Jeffrey R. Yost, *Making IT Work: A History of the Computer Services Industry* (Cambridge, MA: MIT Press, 2017).

exchanged between the Lab and the outside community. Willis Ware at RAND was a prominent figure in the SHARE community for decades. The maturing computer industry of the 1950s, and the concurrent growth of scientific-computing centers utilizing the systems coming out of that industry, meant that established centers like Los Alamos were not as unusual or as isolated in their computing use as had been the case in the 1940s and early 1950s. Just as the maturation of the computer industry allowed Los Alamos to transition from developing its own computer systems to purchasing the majority of its production machines, the increasingly widespread and interactive scientific-computing community of the late-1950s onward eased some of the Lab's software-development burden.

Sharing utilities, subroutines, and other commonly used software tools provided T Division with greater freedom to focus its programming and software-design efforts on projects tailored specifically to LASL's needs. While those needs included programmatic code work, they also included areas of software development that remained immature. Operating systems, for example, comprised a relatively new field of software R&D in the 1950s and 1960s. During that time, T Division found most of its computer vendors to be lacking in the OS builds, features, or hardware compatibilities that LASL required. As a consequence, T Division focused a significant amount of its software production effort and funding on custom OS builds that suited the Lab's requirements. While T (and then C) Division's OS strategy fell into question during the 1970s, the Divisions repeatedly transitioned from building custom hardware and software solutions, to purchasing those solutions as particular computing technologies matured among vendors. The transition

was not always voluntary on T or C Division's part, and did not always result in the acquisition of solutions better-suited to LASL's computing requirements, as subsequent chapters will demonstrate, but the transitions were among the most impactful points of negotiation between the computing divisions and users over the development and philosophy of computing at Los Alamos.

V. The Home Stretch

In April of 1960, the month in which the original Stretch Project contract between IBM and the AEC stipulated the arrival of the IBM 7030 at Los Alamos, the machine was still not complete. The major system components were only being fully assembled and interconnected for the first time. The arduous process of troubleshooting the final assembly and conducting acceptance testing still awaited, while significant performance issues with the memory units remained. Slipping computer delivery dates were nothing new in the computing industry. The Sperry Rand LARC, for example, would not be delivered to Livermore until June of 1960, three years later than the Sperry Rand contract with the AEC originally stipulated. In the case of advanced scientific-computing systems like the LARC and Stretch, vendors and customers alike tended to underestimate the difficulty of completing the needed R&D and preparing to manufacture the new technologies involved. At the same time, vendors could easily overestimate the ease of integrating those technologies to work as desired once they were completed. Underestimating the difficulty of software development became its own, long-standing issue in scientific and most other forms of computing over the subsequent decades.

Unlike the hardware delays, IBM's delays with its software coding work for Stretch, such as FORTRAN for the 7030, had little impact on Los Alamos, because T Division developed in-house the most critical utilities, including physics-code work, needed to bring Stretch into production.¹⁷⁹

As early as 1958, Stephen Dunwell and his IBM development team realized that the company's optimistic performance estimates for Stretch were likely not possible, given the emerging constraints of the core-memory system, and compromises involved with developing one system for LASL, the NSA, and other customers. The IBM group also had extensive difficulties packaging the computer's components densely enough into IBM's standard module sizes for high-performance while maintaining acceptable reliability. As a consequence, IBM and the AEC renegotiated the Stretch contract in 1959 to allow for slower processing speeds and the elimination of the small 0.5-microsecond memory unit. The renegotiation included a larger 2-microsecond memory (up to 98,000 words) and faster I/O accessories to help compensate. T Division reports from Bengt Carlson just before and after the renegotiation tempered the Lab's performance expectations starting in 1959, but continued technical issues also delayed

¹⁷⁹ R. E. Merwin, "Letter from R. E. Merwin, Advanced Computing Systems, to Paul Wilson, Area Manager, Los Alamos Field Office, Atomic Energy Commission, April 4, 1960," Edward A. Voorhees Papers, box 17, folder 3, IBM Progress Report, Jan. 1960-Nov. 1960, series 22, LANL Archives; Bashe, et al., *IBM's Early Computers*, 436; Ghausi interview with Fernbach; Lewis interview with Folkner; Although the 7030's delivery date and performance expectations were delayed and altered toward the end of the machine's development, LASL in particular found great reassurance in the overall design and potential of the Stretch computer when, in 1959, IBM released its 7090 system. The 7090 owed much to the Stretch design, but had several of Stretch's more complex features stripped out, such as the lookahead unit and the larger memory. The 7090's simpler design and great ease of production meant that it shipped before Stretch, becoming the first commercial transistorized (or second-generation) computer. The 7090 provided excellent performance, and proved well-suited to LASL's workload (not a surprise, given the connection between the 7030 and the 7090's hardware and software development. Much of LASL's Stretch software was made to be compatible with the 7090.), with T Division ordering two 7090s to be delivered and used alongside Stretch.

delivery for an indeterminate period. For over a year, Dunwell sent progress reports to LASL and the AEC, stating that the 2-microsecond main memory was nearly operational, but the memory unit itself did not become fully operational until the late summer of 1960. LASL users had begun early tests of Stretch's arithmetic unit in April of 1960, but IBM halted testing, due to the machine's incomplete condition. At IBM's request, LASL users resumed more elaborate system tests starting in October, once Stretch's hardware was fully assembled, and many of its most pressing issues were resolved. Dunwell's team extensively used the Stretch emulator for the 704 that Harwood Kolsky helped IBM to develop to test the performance and operational impact of various adjustments to the Stretch system, without going to the time and expense of experimenting with different settings on the 7030 itself.¹⁸⁰

Using test codes and programs from T Division, LASL and IBM testing at Poughkeepsie through late 1960 and early 1961 revealed that Stretch's average performance was approximately half of what was intended, at about four times the performance of the smaller IBM 7090. The goal for Stretch had been closer to eight times the 7090. Stretch could maintain a significant performance advantage over the 7090 on high-precision problems, and problems where the lookahead unit could most successfully mitigate I/O latencies. IBM's management was not initially concerned with the technical problems associated with Stretch, its cost overruns (development costs were

¹⁸⁰ Bashe, et al., *IBM's Early Computers*, 441, 444-445; S. W. Dunwell, "Letter from S. W. Dunwell to Paul Wilson, Area Manager, Los Alamos Field Office, Atomic Energy Commission, April 1, 1959," Edward A. Voorhees Papers, box 17, folder 3, IBM Progress Report, May 1957-Dec. 1959, series 22, LANL Archives; "Summary of Stretch Usage by LASL Personnel, 1/26/61," Edward A. Voorhees Papers, box 17, folder 6, Stretch Acceptance Test Procedures, series 22, LANL Archives.

over twice the original estimate), or its early performance issues, believing that the time and associated costs would contribute to other future products. The continuing performance discrepancies leading up to delivery were a different story. Thomas Watson, Jr. delayed announcement of Stretch's completion for several months while testing continued. Speaking for T Division, which, from the beginning of the collaboration, had been less concerned than IBM with performance estimates, Bengt Carlson emphasized the importance of features and reliability in the Lab's acceptance testing. A reliable computer with features and capabilities designed, at least in part, for LASL's users, was of greater value to T Division than the highly variable performance figures. Neither Dunwell nor Carlson appeared to have doubts that the computer would prove more than adequate for LASL's needs, leaving final price negotiations to IBM Sales and the AEC.¹⁸¹

An important problem estimating Stretch's performance before the computer's completion was the system's unrivaled complexity. Despite Dunwell's team simplifying the original hardware design in 1958, the 7030 was quite easily the most complex computer ever constructed up to that point. The collaborative effort with LASL, and IBM's inclusion of additional cryptographic hardware and other features for the NSA version of the machine, called Harvest, resulted in an instruction set of over 730 instructions. Stretch's oversized instruction set expanded the computer's footprint to match, because of the added hardware. For comparison, Seymour Cray's first

¹⁸¹ Bengt Carlson, "Letter from Bengt Carlson to Stephen Dunwell, IBM, LASL Acceptance Test Programs, February 13, 1961," Edward A. Voorhees Papers, box 17, folder 6, Stretch Acceptance Test Procedures, series 22, LANL Archives; Lewis interview with Folkner; Bashe, et al., *IBM's Early Computers*, 446, 450-451; Pugh, *Memories that Shaped an Industry*, 183.

supercomputer design, the Control Data 6600, entered the design phase in 1960, and featured only sixty-four built-in instructions. A computer's instruction set was a collection of machine-language commands available to a programmer to tell the processor what to do. Instructions could be built-in to a computer's hardware architecture, or emulated through a software interpreter. Built-in instructions were far more efficient than using software, which was why LASL users and programmers were eager to contribute to Stretch's instruction set, because built-in instructions provided them greater flexibility and ease of programming, while still having the performance advantage of using an assembler. Instructions tailored to LASL's most commonly used or needed operations could significantly ease the programming burden.¹⁸²

As a consequence of its complexity, Stretch was a physically large machine. Its Central Processing Unit alone, spread across multiple cabinets, was over 30-feet long, seven-feet wide, and six-feet high, containing nearly 170,000 discrete transistors. The two hard disk units, capable of storing 2 million words across 39 platters each, were described as looking like very large juke boxes. With over 300,000 transistors spread across the entire computer, including its six memory boxes (which contained Stretch's core memory planes in an oil immersion to maintain a consistent temperature), and its disk and tape storage units, printers, card readers, and its operator, storage controller, and maintenance consoles, Stretch produced about 370,000 BTUs of heat per hour. As a

¹⁸² "CDC 6600, LASL Configuration," Edward A. Voorhees Papers, box 35, folder 4, Computer Characteristics, series 27, LANL Archives; "Stretch Computer (IBM 7030), 7/13/65," Edward A. Voorhees Papers, box 2, folder 2, Family Days, 1955, 1960, 1965, series 5, LANL Archives; The CDC 6600 is typically hailed as the first "RISC" (Reduced Instruction Set Computer). While Seymour Cray enjoyed the engineering simplicity of having a very small instruction set, many users at Los Alamos were not pleased with the concept, as a limited collection of hardware instructions placed a greater burden on programmers.

result, the computer required an extensive, purpose-built air conditioning system.

Purchasing a 7030 required a significant commitment of space and resources, including the space required for the IBM technicians who remained in residence with each machine.¹⁸³

Unable to fit, air condition, or adequately power Stretch in the Administration Building already filled with the MANIAC II and the IBM 704s, Los Alamos constructed a state-of-the-art facility next door. Officially designated as Building 132, or the Central Computing Facility (CCF), it was generally known at Los Alamos as the "Stretch Building" until the computer was removed in 1971. Part of a separate \$1.7 million allocation from the JCAE, the Stretch Building featured a self-supporting accordion-pleated roof that left the interior floor space free of load-bearing walls and columns. In addition to the 7030, LASL constructed the Stretch Building to house the computer's support staff, its storage tapes and data-plotting equipment, and an IBM 7090 that replaced one of the leased 704s. The Stretch Building, as the Lab's first facility dedicated entirely to computing, represented LASL's significant commitment to Stretch, and the growing recognition of computing as a unique resource at Los Alamos, rather than as simply an extension of T Division's physics work.¹⁸⁴

In March of 1961, thirteen LASL scientists conducted Stretch's final round of fifteen pre-shipment reliability tests in Poughkeepsie. As Carlson had predicted the previous month to Stephen Dunwell, Stretch performed well in what was T Division's

¹⁸³ Savage, "Understanding Stretch," 7-8; Kenneth J. Johnson, ed., "Stretch, the Super Computer is Gone," *The Atom* 8, no. 6 (July-August, 1971): 23-26.

¹⁸⁴ Savage, "Understanding Stretch," 7-8; Johnson, ed., "Stretch, the Super Computer is Gone," 23-26.

most important delivery criteria. According to Carlson, running tests from five minutes to over an hour long, the 7030 averaged 90% "good time." That is, 90% of the time the computer was switched on, it was running LASL problems successfully, easily beating the contracted minimum of 80% "good time." In dozens of crates divided among six moving trucks, Stretch made the 2,000-mile journey from Poughkeepsie to Los Alamos in April 1961. The delivery made news across trade publications and dozens of New Mexico and national newspapers. The May 4, 1961 issue of LASL's internal *LASL News*, attempting to distinguish the latest computer from the Lab's existing collection of large mainframe systems, referred to the new machine as a "super computer."¹⁸⁵

After Stretch's reassembly at Los Alamos in late April, post-delivery reliability tests began in May. After a false start on May 8, testing began in earnest on May 10, with the computer returning 97% and then 99% "Good Time" on multiple, continuous runs over the next week, leading to LASL's formal acceptance of the system from IBM. The IBM 7090, installed also in the Stretch Building in early April, by contrast, was still only returning an 80% "Good Time," which T Division considered "marginally acceptable." A possible unintended consequence of the heightened level of technical attention Stretch received at IBM Poughkeepsie, the 7030 had relatively few teething issues at Los Alamos, and entered production use quickly for a computer of the time period and of its complexity. By July, users at LASL were already making extensive production runs, operating without a system fault for up to seventeen hours in some

¹⁸⁵ Johnson, ed., "Stretch, the Super Computer is Gone," 23-26; "Computer Activity Report, Los Alamos Scientific Laboratory, May 29, 1961," Edward A. Voorhees Papers, box 2, folder 3, AEC Computer Meetings 1956-1961, series 6, LANL Archives.

cases, an impressive number in high-end scientific-computing at the time. The system logs reported that core dumps to the hard disk units worked as planned, and made recovery after system errors extremely fast. The extensive site preparations and LASL's role in developing much of the system software (most of which was tested through emulation on 704s or 7090s at IBM or LASL) also meant that a core of needed system utilities, and a variety of user jobs, were ready for use before or upon the computer's arrival.¹⁸⁶

Stretch enjoyed a relatively smooth introduction to Los Alamos, with mostly favorable reports from early users about its reliability, speed, and ease of programming. Most users at LASL were accustomed to the limited memories, and reliability and performance issues of the first-generation systems in use at the Lab since the first MANIAC. One physicist testing the 7030 reported his dramatic introduction to Stretch's capabilities when running an eigenvalue problem that typically required ten minutes to complete on the 704. Stretch loaded the card deck via the reader, and, immediately after spitting out the final card, the "halt" light illuminated on the operator's console. After half-an-hour of troubleshooting, the operator discovered that nothing was wrong. Stretch had simply finished the problem.¹⁸⁷

While Stretch made short work of jobs like these, the computer's primary advantage to LASL was its ability to perform far more complex problems in the same amount of time in which previous machines processed simpler tasks, while making the

¹⁸⁶ "Computer Activity Report, Los Alamos Scientific Laboratory, May 29, 1961"; "Logs for IBM 7030 Stretch Acceptance Testing, March, May 1961," Edward A. Voorhees Papers, box 18, folder 2, Stretch Acceptance Testing Logs, March-May 1961, series 22, LANL Archives.

¹⁸⁷ Knutsen, *Wasn't That a Time*, 40; Bashe, et al., *IBM's Early Computers*, 452.

overall programming effort less difficult for users than on the earlier computers. The typical one-billion arithmetic calculations needed for a single design study for a thermonuclear warhead in the early 1960s required about ten hours to complete on an IBM 704. Stretch could perform a design study of far higher resolution, requiring approximately thirty times the number of calculations, in the same ten hours. At the same time, Stretch's large instruction set, tightly integrated programming aids, large core memory, and excellent reliability meant that users could code problems of greater computational complexity than before, but faster and with fewer errors. Stretch could also detect and recover from errors automatically, reducing the time operators and the job queue lost to troubleshooting. IBM's focus on arbitrary performance benchmarks, while of less interest to T Division than data handling, reliability, and ease of programming, dominated external depictions of Stretch, even before its delivery to Los Alamos. From T Division and LASL's perspective in 1961, Stretch was an almost unqualified success.¹⁸⁸

After Watson, Jr's disappointment with the 7030's performance figures in late 1960, he announced that IBM would lower the selling price from about \$13.5 million to \$7.8 million, to reflect the approximate 40-50% performance shortfall. The change in price was connected to Livermore's expressed interest in purchasing its own Stretch, after the extremely late arrival of the LARC. After Livermore agreed to the reduced price, Watson extended the reduction to the other customers that had already entered negotiations to purchase a system, and canceled further Stretch orders. Because of LASL's involvement with Stretch's hardware and software design, the AEC and IBM

¹⁸⁸ Savage, "Understanding Stretch," 7-8.

negotiated a final price of about \$5 million for LASL's 7030, including the large assortment of ancillary devices included with the computer. Watson readily accepting culpability for the performance shortfall, and IBM's significant reduction in price on the pre-existing Stretch orders, surprised industry insiders. The immediate consequences of Watson's reaction monopolized much of the subsequent industry discussion of the 7030, which changed only slightly after Watson later retracted or appended some of the blame and dispersions he had cast in 1960 and 1961. Watson's reaction to Stretch bore longer-term consequences for IBM, Los Alamos, and the scientific-computing industry over the 1960s, just as the "super computer" began to solidify into what became the category of "supercomputing."¹⁸⁹

Stretch was the largest single leap in computing capability for Los Alamos since the introduction of the ENIAC, but it embodied over a decade of T Division's collected experience in developing and adapting general-purpose computer technology specifically to problems in nuclear-weapons development, and scientific-computer modeling in general. Unlike the LARC, which was almost entirely a vendor-designed system, the components and abilities that mattered most to Los Alamos bore T Division's fingerprints from the beginning of the formal design phase, as LASL's Mathematical Planning Group represented some of the field's most experienced computer hardware and software experts then available in the relatively young computing field. The result was a system that conformed closely to LASL's particular computing requirements, from its numerous,

¹⁸⁹ Pugh, *Memories that Shaped an Industry*, 184; Bashe, et al., *IBM's Early Computers*, 452-453; "The Shrinking of Stretch," *Datamation* (June 1961): 46; After its 1956 Consent Decree with the US Department of Justice, IBM began selling computers to customers, in addition to leasing them.

physics-oriented built-in instructions, to its advanced job-management routines, to its fast I/O and job recovery capabilities.

The development of Stretch and its installation in 1961, only nine years after the completion of the MANIAC I, grew out of and highlighted several processes underway at Los Alamos over the 1950s. First, it witnessed T Division's transition from operating a variety of computing technologies—including human, electromechanical, and general-purpose computing—to the dominance of general-purpose digital systems as the Lab's primary production computers. This transition occurred as the computational demands on nuclear-weapons development, and on the AEC labs' research-diversification efforts, increased with the thermonuclear crisis. Proponents of general-purpose computing at Los Alamos and elsewhere also played a role in advocating and advancing their preferred technology over more traditional methods.

The Stretch Project was also the culmination of a transition from locally constructed systems to vendor-sourced computer hardware, amid the maturation of the computer industry. Unlike the LARC with Livermore, Stretch represented a melding of LASL's unusually long and complex experience and expertise in scientific-computing with a vendor's design and manufacturing capabilities. The extensive integration of T Division's design input into Stretch's hardware and particularly its software contributed heavily to the differing interpretations of the 7030 between Los Alamos and IBM. For Los Alamos, Stretch meshed with the Lab's computing priorities and demands more closely than any system then available, including the MANIAC II. For IBM, a profit-oriented computer company on its ascendancy in the computing market, Stretch had

meanings extending beyond how the machine was integrated into a demanding workflow. The following chapter examines the long-term consequences of Stretch for the shaping of Los Alamos' computing use and philosophy, and how Stretch contributed to the form of the emerging supercomputer industry of the 1960s.

-Chapter 3-

The Ties that Divide: Supercomputer Selection and the Formation of C Division

I. Chapter Introduction

In summary of chapter 2, the collaborative development of the IBM 7030 "Stretch" completed the transition of LASL's central-computing operation to general-purpose digital computers sourced from IBM. Once a diverse collection of technologies and methods, LASL computing focused increasingly on general-purpose computers due to their demonstrated utility for the particular problems facing Los Alamos. The computing transition took place during LASL's evolution from the sole AEC weapons-design lab, to being one of two competing design labs expanding their research focus in the midst of a maturing Cold War and expanding US nuclear enterprise. LASL's T Division embraced the Stretch collaboration as a means of producing a scientific computer tailored to the Lab's computing needs, while benefiting from IBM's design and manufacturing prowess. Los Alamos users regarded Stretch highly for its programming and architectural features, and T Division hoped to repeat its collaboration with IBM on a successor machine. In contrast to LASL, IBM's reaction to Stretch's perceived faults prompted IBM management to issue a series of lasting changes within that company. IBM's reaction, combined with new oversight and budgetary constraints from the AEC, forced T Division to abandon its post-Stretch computer-development plans.

This chapter examines how T Division's search for other computing options in the mid-1960s, taking place during an intensification of computer use at Los Alamos, and in

the middle of a notorious corporate rivalry, reshaped fundamental aspects of LASL's computing operation. That reshaping, which included a radical and contested technological shift to a new computer vendor, spurred the formation of LASL's first dedicated computing division. The computing division's formation, a reflection of the growing importance of computing to LASL's mission, was an attempt to unify—and to ameliorate T Division's historically contingent dominance over—the Lab's computing resources and R&D focus. The reorganization brought greater unity to many of the previously diverging computing efforts at Los Alamos, but the continuity in personnel and philosophy between T Division and the new Computing Sciences and Services Division left unaddressed many of the underlying tensions between a new contingent of programmatic computer users and T Division's approach to computing. As a consequence, the LASL computing operation, no longer contained within a powerful science division, found itself increasingly exposed to the demands and critiques of the broader Los Alamos user community by the end of the 1960s.

II. "Super Stretch" and the Supercomputer Market

The Stretch computer met with a favorable reception at Los Alamos, garnering accolades from LASL computer users as it quickly became the centerpiece of the Lab's programmatic work. Stretch's unusually smooth integration into production use after its arrival in 1961 owed much to T Division's role in shaping the computer and its software specifically to LASL's requirements. The computer's large, core-based memory hosted problems of greater complexity and resolution, which users could code more easily and

with less debugging hassle thanks to the large set of instructions, T Division's software programming aids, and the machine's built-in error-correction and recovery capabilities. Stretch's hard disk drives, which could copy or restore the entire contents of the main memory in less than one second, vastly reduced time wasted to I/O cycles on larger jobs, or on queues of smaller jobs managed by the MCP operating system, itself designed at LASL. For Los Alamos, these and other cutting-edge capabilities, such as Stretch's multi-processing ability (in which the OS could run and prioritize multiple programs simultaneously) and its lookahead system preventing CPU-wait time, made up for the computer's late arrival and its specifications shortfall.¹⁹⁰

Outside Los Alamos, Stretch was largely known for the issues Thomas Watson, Jr. publicly announced at the 1961 Western Joint Computing Conference, primarily regarding its production delays, slower-than-predicted performance (mostly of system memory), and development costs greatly exceeding initial predictions. As a consequence, despite supplanting the Sperry Rand LARC as the world's fastest computer, and bringing IBM into a technical lead over its competitors, Stretch did not receive the fanfare upon release that surrounded other large IBM computing projects, like the SSEC or the NORC. Along with reducing its price, and canceling the computer's production beyond the eight pre-existing orders after LASL's 7030 (making nine 7030s total) due to

¹⁹⁰ Bashe, et al., *IBM's Early Computers*, 451-453; Stretch's error detection and correction capabilities were of greater importance to Los Alamos than most other contemporaneous computing sites due to LASL's high-altitude location in the Jemez Mountains of New Mexico. Located at approximately 7,100 feet above sea level, the Lab's computing operation was more susceptible to errors from cosmic background radiation, which could flip bits in system memory. As computer memories and data volumes grew larger, errors from both coding mistakes and environmental causes also became more likely. For additional information about the technical and environmental basis of computer errors at Los Alamos, see: Eugene Normand, et al., "First Record of Single-Event Upset on Ground, Cray-1 Computer at Los Alamos in 1976," *IEEE Transactions on Nuclear Science* 57, no. 6 (December 2010): 3114-3120.

each machine being sold at a loss, Watson assigned much of the blame to Stephen Dunwell, the leader of the Stretch Project. Dunwell remained at IBM, but with a tarnished reputation. Although the Stretch Project had accomplished its goal of making IBM a leader in cutting-edge computer design and technologies, Watson's embarrassment with the 7030 had far-reaching consequences. The most significant consequence for Los Alamos was Watson's drive to avoid repeat embarrassments, causing him to diminish the company's investment in high-end scientific computing in the early 1960s, just as LASL and Livermore together began to seek options for Stretch's successor.¹⁹¹

Before discussing IBM's changing business directions following Stretch, it is necessary to understand the Cold War developments that were concurrently shaping the AEC's lab complex in the early 1960s. Those developments directly influenced how Los Alamos perceived its computing requirements, and how it acted upon its options. In spring of 1960, a year before Stretch's delivery at LASL, the AEC's Computer Advisory Group (of which Robert Richtmeyer and Willis Ware were members, and Nicholas Metropolis was chairman) produced a report that stressed the "critical importance" of computers to most facets of research within the AEC complex, particularly nuclear-weapons design and evaluation. The report stated that the AEC already supplied the greatest support for scientific computation of any institution in the United States, but further measures were needed to avoid becoming complacent, and to avoid falling behind the rate of change in computing technologies. A key recommendation of the report was

¹⁹¹ Bashe, et al., *IBM's Early Computers*, 450-452, 455-456; The NSA's Harvest, which included one of the nine total Stretch systems constructed, was well-received upon its installation in 1962. The Harvest module roughly doubled the computer's size over the Stretch system. Harvest was in use at the NSA for fourteen years, while LASL's Stretch was in use for ten.

to increase AEC funding and resource allocation at universities to foster fundamental R&D into computer uses relevant to the AEC labs, such as hydrodynamics and plasma physics. The other recommendation was to encourage AEC labs to support a research and feasibility study in collaboration with vendors in order to produce computers with "capabilities significantly greater than the Stretch machine." The advisory group recorded in its April 1960 meeting, but not in the resulting report, that it feared computer designs were becoming standardized around manufacturers' economic interests, rather than being optimized for the computational problems at the AEC labs. Direct involvement with component and computer-system research, the group argued, would ensure that the AEC's interests remained relevant among scientific-computing vendors.¹⁹²

The formation of the Computer Advisory Group in 1958 and the content of its Spring 1960 report were in response to Cold War political developments that shaped how the AEC nuclear-weapons complex had operated through much of the 1950s. In 1958, the United States and Soviet Union mutually agreed to a moratorium on full-scale nuclear testing. The moratorium began in response to growing international unease with the superpowers' continued and increasingly powerful atmospheric nuclear tests following the Castle Bravo incident. Bradbury and other AEC lab directors had feared the possibility of a moratorium, and had committed to diversifying the research efforts of

¹⁹² AEC Computer Advisory Group, "Report on the Status of Computer Activities, April 28, 1960," Edward A. Voorhees Papers, box 32, folder 1, SPARC Correspondence, series 25, LANL Archives; "AEC Computer Advisory Group Meeting, April 7, 1960," Nicholas Metropolis Papers, A-98-011, box 4, folder 4, AEC Computer Advisory Group, LANL Archives.

their respective labs to avoid redundancy or the potential loss of personnel if the AEC attempted to scale back its nuclear-weapons programs.¹⁹³

In anticipation of a moratorium, LASL had quickly embraced atomic- and fusion-reactor research to leverage its existing nuclear expertise, with the launch of Sputnik in 1957 offering political and financial incentive to the Lab's "Rover" atomic-powered rocket program. Los Alamos also developed a particle-accelerator program, for which Norris Bradbury went over the heads of the AEC General Advisory Committee's leadership to obtain congressional funding. As a consequence of the "Eisenhower" moratorium (1958-1961), Los Alamos went from one-quarter of its research effort being devoted to non-weapons programs in 1957, to two-thirds being devoted to non-weapons work by 1961. The Lab's weapons program changed considerably during that time, with many staff members transferring to other research roles. The remaining weapons projects were scaled back due to testing restrictions, and over 300 staff members connected to the weapons program were let go from LASL completely. Nuclear weapons development had already been transitioning from the rapid rollout of new designs (sometimes several per year) in the early to mid-1950s, to a slower iterative process by the late 1950s. Incremental efficiency refinements and the improvement of safety features occupied more of the design effort as the 1960s began. This was particularly true at Los Alamos, which prided itself in observing a more conscientious, risk-averse design philosophy than

¹⁹³ Westwick, *The National Labs*, 215; Norberg interview with Bradbury; Richard G. Hewlett and Jack M. Holl, *Atoms for Peace and War, 1953-1961: Eisenhower and the Atomic Energy Commission* (Berkeley: University of California Press, 1989), 449-451.

Livermore's. The moratorium accelerated many processes that were already underway at LASL, including the increasing reliance upon computing.¹⁹⁴

Full-scale atmospheric testing had been the primary means of verifying the performance and functional characteristics of weapon designs before and after they entered the US nuclear stockpile. At the heart of the Computer Advisory Group's call to action was the fact that existing computer simulations and mathematical models were not predictive representations of physical systems. Due to the limitations of computer technologies and the state of the field of numerical analysis in the 1950s and 1960s, the computer models in use within the computational sciences were based on approximations. Although extremely useful for guiding the design process, and necessary when empirical methods were not practical or possible (as was often the case with T Division's weapons work), these approximations had unknown consequences for the accuracy of simulations, particularly where approximations interacted with one another. Live tests were used to calibrate the computer simulations, but the calibrated simulations did not result in a better representation of the underlying physics within a detonating weapon, just a closer approximation of the observed results. As John Pasta summarized at the April 1960 advisory group meeting, with the testing program curtailed, weapons scientists were doing their best to answer, through interpolation on computers, the Department of Defense's questions about weapons systems, but the weapons scientists could not extrapolate based on those models.¹⁹⁵

¹⁹⁴ Westwick, *The National Labs*, 216-217; Norberg interview with Bradbury; Hewlett and Holl, *Atoms for Peace and War*, 449-451.

¹⁹⁵ AEC Computer Advisory Group, "Report on the Status of Computer Activities, April 28, 1960"; "AEC Computer Advisory Group Meeting, April 7, 1960."

During the moratorium, the diminished weapons programs at LASL and Livermore became more dependent on computers, but the computers and numerical methods available were not sufficient for completing the process of bringing a weapon into the arsenal. As a consequence, the design work slowed. This was not perceived as a pressing issue within the DoD, as LASL and Livermore had already supplied a diverse collection of nuclear weapons that were in rapid production for the US stockpile. New design work was largely focused on making incremental changes. Non-weapons work, particularly reactor design supported in part by NASA and ARPA, experienced a rapid increase in need for computing resources, but the result of the moratorium at LASL was an overall slower rate of growth in computing demand than had been anticipated when the Stretch Project began, when T Division planned an approximate doubling of computing power every two years. The installation of Stretch and two IBM 7090s in 1961 provided Los Alamos with a temporary excess of computing capacity, which further eased LASL's acceptance of Stretch's reduced performance figures.¹⁹⁶

For T Division, the slower-than-expected increase in computing demand, and the excess capacity the new IBM computers provided, did not indicate a diminished need for advanced computer systems at Los Alamos in the future. Rather, the moratorium signalled that LASL had to remain flexible in its research focus and capabilities if it were to remain relevant amid changing Cold War demands. With the weapons program struggling in its attempt to supplant live testing with computer simulations, and non-weapons research becoming a large, and even dominant part of LASL's mission, T

¹⁹⁶ AEC Computer Advisory Group, "Report on the Status of Computer Activities, April 28, 1960"; Westwick, *The National Labs*, 186, 216-217.

Division recognized that it needed to seek far more capable computer systems and programming tools, regardless of when or if the moratorium might end.¹⁹⁷

In early 1960, T Division had already begun discussing with IBM the possibility of a follow-on collaboration on a so-called "Super Stretch." When T Division requested bids from computer vendors, in accordance with the Computer Advisory Group's call for the formation of collaborative R&D studies, IBM's early proposal for the Super Stretch became the basis of that company's bid for an "Ultra High Speed Computer Study." The goal behind IBM's proposed study was to research technologies that would contribute directly to a computer approximately 100 times faster than Stretch. The proposal fit within T Division's hopes of repeating the Stretch collaboration, but IBM's expected price tag on the development effort was far higher than the maximum of \$250,000 LASL had hoped to distribute among up to three vendors for the first year of its study contracts.¹⁹⁸

After the increasing budgets of the 1950s, the new Kennedy Administration refocused US national security policy toward a flexible-response strategy, diminishing the US focus on nuclear weapons and the need for the nuclear-weapons labs. As a consequence, the AEC labs' budgets, after peaking in 1960, decreased year-over-year into the mid-1970s. LASL had to be more conscientious with its computing funds than ever

¹⁹⁷ Westwick, *The National Labs*, 269-270. The development of truly predictive computer models for physical systems remains a goal at Los Alamos and in the US national lab complex in the early twenty-first century, urged on by the comprehensive ban on full-scale nuclear-weapons testing that began for the US in 1992. The computational sciences did not begin to transition toward the creation of predictive models until the rise of massively parallel computing technologies in the 1990s, due to a recognition of the processing power and advanced programming techniques that would be needed for predictive capabilities to be possible. The US Department of Energy and other organizations have invested many billions of dollars since the 1990s to aid in the advancement of computer-based simulations for scientific research. In the 1960s, the goal was to generate the best possible approximations through better codes and computers, even if truly predictive models were not feasible.

¹⁹⁸ Edward Voorhees, "IBM "Super-Stretch" Proposal, May 6, 1960," Edward A. Voorhees Papers, box 32, folder 4, SPARC Proposal, series 25, LANL Archives; Westwick, *The National Labs*, 277-279.

before, but funding was only one way in which the Lab's traditional freedom in computer selection was being scaled back. The cost and delivery overruns of both the LARC and Stretch gave the AEC pause over the labs' effective autonomy in their computer collaborations and purchasing decisions. The desire to encourage computer research beneficial to the AEC weapons lab drove the Computer Advisory Group's call for directed R&D investments among multiple computer vendors. Budget-conscious AEC management viewed the smaller-scale investments as a means of limiting a lab's financial commitment to any one vendor before a working computer was available. IBM returned to Los Alamos with a more financially palatable, smaller-scale proposal. T Division accepted the bid, which became the core of LASL's first study contract.¹⁹⁹

Bengt Carlson reduced development costs of LASL's study contracts further through coordinated research with Project Lightning, a computer-component R&D program funded by the NSA under the guise of the US Bureau of Ships. The NSA had contracts with several vendors to develop "ultra high speed" computer circuits and memory, although not complete computers, as the AEC weapons labs explicitly sought in their contracts. Carlson wanted to avoid duplicating NSA-funded research, and even hoped to build upon it. The Bureau of Ships was initially reluctant to disclose its computing contracts and level of investment, but Carlson was, by 1961, able to coordinate LASL's computer-component R&D investments at RCA, IBM, and Sperry Rand with those of the NSA. By spring of 1961, the joint LASL and Livermore study-

¹⁹⁹ IBM Research Center, "Proposal for An Ultra High Speed Computer Study," Yorktown Heights, NY (May 2, 1960), 1-3, 5-6, Edward A. Voorhees Papers, box 32, folder 4, SPARC Proposal - IBM - Super Stretch, series 25, LANL Archives; Westwick, *The National Labs*, 278.

contract effort formally became the SPARC Project, the plan for continued R&D investments with several vendors, including IBM, over the course of several years. The goal, coordinated with Project Lightning's component R&D investments, was to use the joint lab funding to advance computer-component and design technology to the point that vendors could produce bids on scientific-computer designs within two years, and deliver a completed system within five. IBM was the favorite for the final computer contract, with LASL and Livermore holding that company's scientific-computer design proposal (discussed in detail below) as the standard for the finished system.²⁰⁰

The SPARC Project collaboration was LASL's best option for a system tailored to its requirements, but it could not be an in-depth design collaboration as LASL had with IBM on the Stretch Project. Financial concerns from the AEC aside, Livermore lacked sufficient hardware experience, and the technical demands of computer design, having become far greater than had been the case with Stretch in the 1950s, meant that LASL could not take sole responsibility for a large-scale design collaboration.²⁰¹

Even before Stretch's delivery to LASL, IBM was a ready participant in the study contracts, which would, given IBM's newfound design prowess with the IBM 7000-

²⁰⁰ Letter from Bengt Carlson, T-1, to Carson Mark, T-DO, May 23, 1960, Edward A. Voorhees Papers, box 32, folder 1, SPARC Correspondence, series 25, LANL Archives; Letter from Bengt Carlson, T-1, to P. M. Petersen, SP-1, May 24, 1960, Edward A. Voorhees Papers, box 32, folder 1, SPARC Correspondence, series 25, LANL Archives; Letter from P. M. Petersen, SP-1, to Bengt Carlson, T-1, May 27, 1960, Edward A. Voorhees Papers, box 32, folder 1, SPARC Correspondence, series 25, LANL Archives; Memorandum from C. R. Lee, Jr., to W. F. Jenike, "Project Lightning," June 15, 1960, Edward A. Voorhees Papers, box 32, folder 1, SPARC Correspondence, series 25, LANL Archives; "Minutes of Meeting on SPARC, Held in Office of Dr. S. Fernbach at Livermore, April 18, 1961," Edward A. Voorhees Papers, box 32, folder 1, SPARC Correspondence, series 25, LANL Archives; Letter from S. Fernbach to Bengt Carlson, T-1, April 11, 1961, Edward A. Voorhees Papers, box 32, folder 1, SPARC Correspondence, series 25, LANL Archives; Letter from Bengt Carlson, T-1, to Darol Froman, TAD, May 12, 1961, Edward A. Voorhees Papers, box 32, folder 1, SPARC Correspondence, series 25, LANL Archives.

²⁰¹ Letter from Bengt Carlson, T-1, to Darol Froman, TAD, May 12, 1961.

series, likely have resulted in IBM winning a future bid to build the Super Stretch (or a variant thereof) that interested both LASL and Livermore. In the Spring of 1960, Emanuel "Mannie" Piore, the new director of IBM's research laboratories, had rehired Gene Amdahl, the noted computer architect who had left IBM after Stephen Dunwell became the Stretch Project leader, to be manager of the Experimental Machines lab. Amdahl's first assignment was to determine how IBM should approach its next large-scale scientific computer. The Stretch Project had vastly exceeded its projected budget (even before the performance issues were known), and Piore contended that IBM needed to be more conscientious about its design goals, and about the promises it made to customers. Piore argued that IBM needed to maintain a strong presence in the scientific-computing market, but it also needed to do a better job of not overselling the unrealized products of ongoing, state-of-the-art research.²⁰²

The result was the formation of two advanced-scientific-computer projects at IBM. The first, under Amdahl, was Project X, whose engineering personnel had experience with the Stretch Project. The Project X computer was intended to be between ten and thirty times more powerful than Stretch, and to be deliverable within five to six years. This was the computer that LASL and Livermore expected to emerge from the scaled-back IBM study contract. Project Y was an even-more ambitious system-design effort, with a computer intended to be about 100 times more powerful than Stretch,

²⁰² Bengt Carlson, T-1, "SPARC Project, May 12, 1961," Edward A. Voorhees Papers, box 32, folder 1, SPARC Correspondence, series 25, LANL Archives; Letter from Hugh H. Spencer, Radio Corporation of America, to Sidney Fernbach, Lawrence Radiation Laboratory, April 26, 1960, Edward A. Voorhees Papers, box 32, folder 1, SPARC Correspondence, series 25, LANL Archives; Norberg interview with Amdahl; Pugh, *IBM's 360 and Early 370 Systems*, 370-371, 374.

similar to IBM's original Ultra High Speed Computer Study. Project Y's computer, being the more advanced of the two, had fewer concrete expectations, and a more nebulous projected delivery date.²⁰³

Just as the two projects began to assemble resources in late 1961, Watson reorganized IBM's business priorities. Watson began to focus IBM's design labs around developing a unified product line, while scaling back the company's investment in scientific-computer design until the technical issues behind Stretch's disappointment were remedied. Caution became a growing emphasis among IBM engineers. As a consequence, the Project X group was small and initially operated only out of Piore's small discretionary budget. Component-research funding from the NSA and the AEC labs soon augmented Project X's operation, but IBM's internal support remained modest. Working under greater high-level scrutiny, the IBM Research and Development laboratories halted unpromising lines of component research quickly, while leaving other projects with minimal funding and staff. Project X was not Gene Amdahl's sole managerial responsibility at IBM, and his time was soon further divided after being drawn into the design of the new, unified system, which would become the System/360. Scientific-computer development was a low priority at IBM following Stretch's troubled rollout.²⁰⁴

In September 1961, the Soviet Union ended the moratorium when it resumed full-scale nuclear testing, catching the US weapons labs off guard. LASL and Livermore quickly moved personnel back to their respective weapons programs to resume live

²⁰³ Pugh, *IBM's 360 and Early 370 Systems*, 370-371, 374.

²⁰⁴ Pugh, *IBM's 360 and Early 370 Systems*, 374-378; Norberg interview with Amdahl.

testing. Los Alamos also hired new staff to replace those Bradbury had previously let go. The US initially confined itself to underground testing, but both the US and Soviet Union were conducting a flurry of full-scale atmospheric tests by spring of 1962. Pressure from the growing environmental movement over the dangers of fallout from atmospheric tests, and increased political pressure for a treaty following the Cuban Missile Crisis, spurred negotiation of the Limited Test-Ban Treaty (LTBT), signed in August of 1963. The LTBT prohibited nuclear testing in the atmosphere, underwater, and in space, moving all live tests underground. Underground testing, which had become more common in the US even before the LTBT, was more expensive and time-consuming than atmospheric testing, but it also presented practical advantages for taking measurements and reducing the visibility of tests in progress.²⁰⁵

As a consequence of the LTBT, LASL's testing moved exclusively to the Nevada Test Site, away from remote Pacific islands. The transition contributed further to weapons development becoming more routinized, moving from a rapid, free-form process surrounding the previous atmospheric-testing spectacles, into a series of predictable, scheduled phases, from concept, to tested, deliverable product. The routinized design process, formally known as PERT, or Program Evaluation and Review Technique, descended from the 1958 development of the US Navy's Polaris missile system. PERT required detailed planning for each stage of development, with feedback controls on scheduling and finances. The success of PERT on the Polaris project spurred

²⁰⁵ Westwick, *The National Labs*, 218; Norberg interview with Bradbury; Lewis interview with Folkner; Joseph Masco, *The Nuclear Borderlands: The Manhattan Project in Post-Cold War New Mexico* (Princeton, NJ: Princeton University Press, 2006), 68-71; Hugh Gusterson, *Nuclear Rites: A Weapons Laboratory at the End of the Cold War* (Berkeley, CA: University of California Press, 1996), 140-144.

the federal government to press other agencies, including the AEC, to adopt PERT in 1962. A warhead design in the 1960s required up to a decade to move through the various development stages, from conception to production, with multiple weapons undergoing development concurrently at LASL and Livermore. Across the 1960s, under the PERT-driven process, only a single new warhead design entered the US arsenal.²⁰⁶

Under the LTBT, US nuclear testing in the mid-to-late 1960s increased rather than decreased, but under a schedule-driven process that escalated the amount of design and computer-based simulation work required to extract ever-smaller improvements out of warhead designs. The new PERT routine and diminishing returns on incremental improvements also resulted in smaller and smaller differences between LASL and Livermore's respective weapon designs. The growing similarity in the two labs' design problems and constraints contributed to greater financial and programmatic pressures to collaborate on computer purchasing and software development. Those pressures mostly resulted in attempted collaborations in the 1960s and 1970s, rather than successful joint efforts, but computing became an increasingly important point of connection and contention between LASL and Livermore.²⁰⁷

With their budgets from the AEC effectively being reduced year over year, and with the resumption of live nuclear testing, Los Alamos and Livermore refocused on their core missions of weapons development. Nuclear-weapons work once again expanded to represent the majority of LASL's operating budget, up to 68%, and even more for

²⁰⁶ Westwick, *The National Labs*, 273-274, 277-279; Masco, *The Nuclear Borderlands*, 68-71; Gusterson, *Nuclear Rites*, 140-144.

²⁰⁷ Westwick, *The National Labs*, 273-274, 277-279; Norberg interview with Bradbury.

Livermore. Non-weapons work, although scaled-back, still occupied a far greater place at LASL than it had before the Eisenhower moratorium. Non-weapons work for military applications actually expanded following 1961, as DoD discretionary spending replaced some of the reduced AEC funding. Laser research, for example, became an important, DoD-funded line of research at both LASL and Livermore, as the technology had a number of potential military and energy applications. As a consequence, laser research became one of the significant new users of computer resources at LASL starting in the 1960s.²⁰⁸

During the SPARC Project between 1961 and 1963, LASL's fleet of IBM scientific computers, consisting of Stretch and two 7090s, left Los Alamos' total computing power ahead of the increase in computing demand that T-Division had expected before the moratorium. The resumption of live testing did increase demand from the weapons program, but only as quickly as that program could bring itself back up to speed after experiencing layoffs, rehires, and personnel transfers, and while adapting to the more stringent requirements of underground-only testing. Although non-weapons research became a significant user of computer power starting in the late 1950s, it was not yet as computer-dependent or intensive as weapons design and simulation, largely because of the greater effort expended in the AEC labs on applying computers to weapons over any other application. The computational sciences (scientific fields that employed computing for modeling and information processing) expanded steadily in the 1960s, with the introduction of a broader range of scientific computers of all sizes, and

²⁰⁸ Westwick, *The National Labs*, 277-279.

software that made them more usable for non-expert users and operators. Even so, the barrier-to-entry in terms of hardware costs and software availability remained high for non-experts until the 1970s, when the computational sciences expanded quickly.²⁰⁹

Expecting vendors to provide bids in 1963 on an "Ultra High Speed" (UHS) computer to be delivered in 1965, T Division focused its computing budget on improving its suite of programming aids and simulation codes. Improved and expanded software libraries were in far greater demand at Los Alamos amid the changing Cold War environment of the early 1960s than the acquisition of new computer hardware. The code effort was of immediate aid to LASL scientists in their use of the existing IBM hardware, and was intended to be of long-term benefit when T Division prepared for the Ultra High Speed computer after it was selected. Vendors still expected scientific customers to provide most of their own software. Upgrading the two IBM 7090s in the Central Computing Facility (CCF) to the improved 7094s in 1963 was intended to bridge the gap between Stretch and the Ultra High Speed computer. The upgrade also kept LASL within the IBM family, where the Lab had invested most of its software development and ancillary-equipment purchases. The IBM 7000-series, with its well-tailored suite of LASL-developed software, remained popular (and sometimes preferred) among Lab computer users into the early 1970s, even after the acquisition of more powerful supercomputers.²¹⁰

²⁰⁹ Pugh, *IBM's 360 and Early 370 Systems*, 373; "LASL Computer Inventory, 2/12/71," Edward A. Voorhees Papers, box 67, folder 8, LASL Computer Inventories 1961-1973, series 49, LANL Archives; Charles Folkner, oral history interview by Nicholas Lewis, Los Alamos, NM (August 13, 2015).

²¹⁰ "LASL Computer Inventory, 2/12/71"; Lewis interview with Folkner; Reid Worlton, oral history interview by Nicholas Lewis, St. George, UT (November 29, 2015).

In August of 1963, the Minnesota-based Control Data Corporation officially announced its CDC 6600 supercomputer. Founded in 1957, CDC stocked its ranks with well-chosen expertise drawn from ERA and Univac, which allowed for a specialized focus on high-end scientific computing that IBM and other large manufacturers did not possess, providing CDC an edge in this segment of the market. The CDC 1604 mainframe, one of the earliest transistorized computers, entered the market only months after the IBM 7090, but at one-third the price, selling more than 50 units at over \$1 million each. Following on this success, the 6600 was Seymour Cray's first supercomputer at Control Data. Design work began in 1960, with the 6600 announced only two months before IBM publicly announced its disappointment with Stretch's performance. Among its more unusual features, the 6600 was the first of what is now called a Reduced-Instruction-Set Computer (RISC), with its CPU employing only sixty-four instructions, compared to the 735 instructions built into Stretch. Instructions for housekeeping tasks and I/O were offloaded onto the 6600's ten Peripheral Processing Units (PPUs), which left the central processor free and optimized for heavy-duty calculation. The 6600 was quite easily the most powerful supercomputer in the world upon its release.²¹¹

The 6600 vastly exceeded any other computer then on the market, readily surpassing Stretch. The IBM Project X computer was expected to be more powerful than

²¹¹ James E. Thornton, "The CDC 6600 Project," *Annals of the History of Computing* 2, No. 4 (October 1980): 338-339, 343; James C. Worthy, "Control Data Corporation: The Norris Era," *IEEE Annals of the History of Computing* 17, No. 1 (1995): 47-49; Charles J. Murray, *The Supermen: The Story of Seymour Cray and the Technical Wizards behind the Supercomputer* (New York: John Wiley, 1997), 74, 92; Reminiscences of Computer Architecture and Computer Design at Control Data Corporation, OH 321, oral history interview moderated by Neil R. Lincoln (May and September 1975). Charles Babbage Institute, University of Minnesota, Minneapolis. <http://purl.umn.edu/104327>.

the 6600, but IBM's slow rate of development on its advanced scientific-computer projects meant that the expected ship date of the Project X system had slipped to 1967, while the CDC 6600 was to ship in 1964. Surprised that a company of less than three-dozen employees could manage such a commanding lead over IBM, Watson ordered an acceleration of Projects X and Y, despite the resources IBM was already devoting to the System/360 strategy. IBM still could not hope to match CDC's offerings in scientific computing for several years, and resorted to slashing the price of its 7094 model to discourage purchases of the new mid-range CDC 3000 series and preorders of the 6600 among existing and potential IBM customers.²¹²

LASL and Livermore were also aware of IBM's growing scientific-computing problem in 1963. T Division hardware experts had followed IBM's progress with its SPARC component-development contracts, and noted a disparity between IBM's reported and actual progress with component technologies, particularly core-memory systems, meant for use in the Project X computer. The slipping of the Project X computer's formal delivery date to 1967 only officially confirmed what LASL and Livermore had already realized. IBM could not meet the proposed 1965 delivery time frame in the SPARC contract, at least not with the computer LASL and Livermore believed to be the most likely successor to Stretch.²¹³

Livermore already had experience with Control Data as a vendor, after purchasing a CDC 1604 as a development platform for time-sharing OS research. Disappointed with

²¹² Lewis interview with Folkner; Donald M. Hall, *Generation of Wealth: The Rise of Control Data and How it Inspired an Era of Innovation and Investment in the Upper Midwest* (Minneapolis, MN: Nodin Press, 2014), 92-94; Pugh, *Memories that Shaped an Industry*, 248-249.

²¹³ Lewis interview with Folkner.

IBM's responses regarding a follow-on to Stretch, or an intermediary machine in the meantime, Livermore ordered two CDC 3600s as cheaper alternatives to the IBM 7094.²¹⁴ Livermore was also the first to order and receive a CDC 6600, but T Division at Los Alamos did not give up entirely on IBM in 1963.²¹⁵ T Division staff continued to work unsuccessfully with IBM engineers on the possibility of producing a computer sooner rather than later out of the SPARC contract research. Growth of computer demand at the Los Alamos Central Computing Facility (CCF) was accelerating, but time remained for the division's leadership and technical experts to consider LASL's options before putting out a bid for a large computer to be delivered in the mid-1960s. Despite losing ground to Control Data, a relative newcomer to the computing market after its founding in 1957 by former ERA staff, IBM still commanded unparalleled resources in the industry, and had a long history as LASL's preferred vendor. With Watson reinvigorating scientific-computer development following the CDC 6600 announcement, it appeared possible in T Division's estimation for IBM to improve its offerings within the following couple of years, even if it could not match Control Data on raw performance. Processing power, as witnessed with Stretch, was not T Division's sole

²¹⁴ With Livermore's larger budget, its Computational Department was, in addition to the two CDC 3600s, also able to purchase two IBM 7094s in 1963, as Los Alamos had. T Division at LASL did not have the funds to purchase, or the personnel necessary to operate, the variety of computer systems available to Livermore in the 1960s. Part of Livermore's ability to begin development on a time-sharing OS in the early 1960s was having the funds and programming personnel needed for an unproven, and not yet technically necessary (at least, not necessary for immediate programmatic purposes), line of software development. Choices in operating-system development and deployment strategies would become a significant differentiating point between the two labs in the decade to follow.

²¹⁵ Insider rumors claim that, for each high-end system Control Data released, the NSA claimed serial no. "0," that is, a computer constructed before the official "first" of a given model, serial no. 1. See: Thomas J. Misa, *Digital State: The Story of Minnesota's Computing Industry* (Minneapolis: University of Minnesota Press, 2013), 128; Edwin D. Reilly, *Milestones in Computer Science and Information Technology* (Westport, CT: Greenwood Press, 2003), 41; and Jitze Couperus (Control Data Systems), at cryptome.org/jya/sadd.htm (April 2019).

interest with its procurement of computing resources. Ease of use, avoiding an expansion of local operating requirements with a new vendor, and preserving LASL's software investment with IBM products were important considerations that T Division was hesitant to sacrifice.²¹⁶

III. The Committees

By 1964, LASL's weapons program had largely adapted to its rebirth as the Lab's primary research pursuit, and to the new underground-testing routine. Coupled with non-weapons research, unclassified and otherwise, intensifying its own use of computer resources, LASL developed a pressing need to augment the computing power of Stretch and the 7094s in the Central Computing Facility with a new, large-scale computer in the near term. The rate of increase in computing demand coincided with the predicted need for a large computer in the mid-1960s. In September of 1964, Bengt Carlson appointed a selection committee of three staff members charged with identifying the Lab's computing requirements, and devising the criteria that a subsequent advisory committee would use to choose a new supercomputer.²¹⁷

The computing market, and computing technologies, had changed considerably since the completion of Stretch. New companies, in particular Control Data, had

²¹⁶ Ghausi interview with Fernbach; Edward Voorhees, "The 7095: IBM Proposal," Edward A. Voorhees Papers, box 32, folder 7, IBM Correspondence and Notes, 1963, series 25, LANL Archives; Edward Voorhees, "Poughkeepsie Meeting, Nov. 22, 1963," Edward A. Voorhees Papers, box 32, folder 7, IBM Correspondence and Notes, 1963, series 25, LANL Archives; Edward Voorhees, "Poughkeepsie, 2/6/64," Edward A. Voorhees Papers, box 32, folder 7, IBM Correspondence and Notes, 1963, series 25, LANL Archives.

²¹⁷ B. Carlson, T-1, to Distribution, September 14, 1964, "Advisory Committee, Selection of Computers," Edward A. Voorhees Papers, Box 35, Folder 2, Series 27, LASL Computer Needs, LANL Archives.

emerged in the scientific-computing space, and new concepts, like concurrency, multiprogramming, and even multiprocessing, were beginning to gain traction as means of accelerating computing performance. Los Alamos needed to assess how those changes impacted its computing options in the mid-1960s, given that neither IBM's product offerings nor its business position in scientific computing were what T Division had hoped they would be following the SPARC contracts. As of 1964, T Division remained the single largest user of the Lab's computer facilities, consuming nearly 40% of their computing time. Edward Voorhees, who administered T-1's computer systems and operations, headed the selection committee's survey of the Lab's existing computing capacity in the summer of 1964, and provided the advisory committee (which advised the AEC on the final purchase decision) the criteria for choosing a supercomputer that could carry Los Alamos through the second half of the 1960s.²¹⁸

The survey concluded that the Lab possessed the equivalent computing power of 2.4 Stretch systems, with the majority of that power spread across the two IBM 7094s, Stretch itself, and the upgraded MANIAC II. The MANIAC II was primarily dedicated to R&D projects, such as programming language development, not running production code. The remainder of LASL's computing capacity was spread across an assortment of smaller computers from IBM and a few other vendors, and were housed in T Division and several other divisions at Los Alamos, and not necessarily in the CCF. As of 1964, the Lab's computer systems were fully deployed, with most all available computing time

²¹⁸ Dept. of Supply and Property, "Procurement Request and Technical Study for Advanced Computer for the Los Alamos Scientific Laboratory," February, 18, 1965, Edward A. Voorhees Papers, Box 35, Folder 9, Series 27, 1801 Study for IBM/360, LANL Archives; "Use of Computing Time by Division," Edward A. Voorhees Papers, Box 67, Folder 3, Series 48, Monthly Summary Sheets 1961-1964, LANL Archives.

being in use. T Division attempted to operate its computers around the clock, seven days per week, but housing in Los Alamos still complicated the hiring situation in the 1960s. Hiring freezes no longer existed, but T Division struggled with the dearth of housing restricting the number of operators the division could hire. The earlier preference for hiring the spouses or live-in relatives of existing employees became less practical during the 1960s, due to the growth of the computer operation, and the need for employees with stronger backgrounds in computing who could operate the increasingly sophisticated machines. In 1962, the US government amended its property rules for AEC communities, permitting the sale of private property for the first time in Los Alamos. The construction of private homes began in Los Alamos in 1963, steadily improving the housing situation from the mid-1960s onward.²¹⁹

The Advisory Committee's survey estimated that, based on the Lab's historical precedent, demand for computing power would increase by approximately 50% each year, doubling every two years, as it had before the testing moratorium. D. V. Susco calculated the doubling trend in terms of Stretch equivalents, estimating that the Lab

²¹⁹ R. Lazarus, T-7, to Whom It May Concern, December 2, 1964, "Computer Selection," Edward A. Voorhees Papers, Box 35, Folder 3, Series 27, Selection Studies, LANL Archives; Bengt Carlson, T-1, to Carson Mark, July 23, 1965, "Housing for Computer Personnel," Edward A. Voorhees Papers, Box 1, Folder 15, Series 1, Salaries, Housing, LANL Archives; J. Carson Mark, T-DO, to N. E. Bradbury, November 30, 1964, "Comments on Computer Capacity," Edward A. Voorhees Papers, Box 35, Folder 7, Series 27, Selection Notes Dec. 1964 – Feb. 1965, LANL Archives; Chambers, "Technically Sweet Los Alamos," 373-375; "Purchase-Lease, Dec. 19, 1966," Edward A. Voorhees Papers, box 67, folder 8, LASL Computer Inventories, 1961-1973, series 49, LANL Archives. In the early 1960s, the T-7 Computer Research and Development group transistorized the MANIAC II, beginning an ongoing process of continually upgrading the MANIAC II to serve as a testbed for new component technologies. This allowed T Division to maintain hands-on experience with the latest technologies in the scientific-computing space, even if LASL was no longer building its own computer systems. This experience proved particularly valuable with the rapid development of computer networking in the late 1960s and 1970s. Like the development of computer systems in the 1940s and early 1950s, a dearth of vendors supplying ready-made solutions meant that Los Alamos had to create most of its own networking hardware, at least until the high-speed computer-networking field began to mature in the 1980s.

would require the power of 3.6 Stretches by the end of 1965 and up to 5.4 Stretches by the end of 1966. The continuing trend would reach 10.8 by 1968 and 21.6 by 1970 (Susco considered the projection to be a severe overestimation of the power that would be needed). The projected trend is what J. Carson Mark, the T-Division leader, described to Lab director Norris Bradbury as a continuing, smooth increase in computing demand. Mark stated to Bradbury, "It is apparent that demand for computing cannot increase very sharply and suddenly in the context of an organization of fixed size with an established program." Both Mark and Roger Lazarus, the advisory committee chairman, predicted a smooth growth in demand for the foreseeable future, noting in their individual reports to Bradbury and the advisory committee that part of the projected growth was an anticipated increase in complexity of codes to be run on the machines.²²⁰

In the mid-1960s, the majority of weapons-simulation code runs operated in just *one* dimension of space, that is, simulating the motion of subatomic particles across a single-dimensional axis, with a smaller number of less-frequently used codes operating in two dimensions. T Division expected the introduction of three-dimensional codes to increase computer-processing demands markedly, but not for several decades. The difference in processing times based solely on changing a problem's number of dimensions was immensely significant. For example, a one-dimensional code that required four minutes of processing time on Stretch, and 100 minutes in two dimensions,

²²⁰ William J. Worlton, "Survey of LASL Computing Requirements," Fall, 1964, Edward A. Voorhees Papers, Box 35, Folder 2, Series 27, LASL Computer Needs, LANL Archives; D. V. Susco, to Whom It May Concern, December 2, 1964, "Computer Selection," Edward A. Voorhees Papers, Box 35, Folder 3, Series 27, Selection Studies, LANL Archives; Dept. of Supply and Property, "Procurement Request," 17-19.

would require a projected forty hours in three dimensions. Each dimensional jump resulted in about a twenty-five-fold increase in required processing time. The problem would only grow more acute with the added variable of time, and with finer and finer simulation resolutions. One- and two-dimensional codes remained the state-of-the-art in the mid-1960s. Computers of adequate power to process some of the predicted three-dimensional codes were not expected to exist for another twenty years, but T Division was already preparing the theoretical groundwork for their development with its in-house software R&D.²²¹

The eighteen-point selection criteria Ed Voorhees developed out of the survey emphasized the need for a computer that would not only be available soon enough to reduce the immediate strain on the Central Computing Facility (CCF), but could also meet the anticipated near-term requirements of increasingly complex one- and two-dimensional codes. Armed with the selection criteria, and a sense of urgency, the advisory committee met in September of 1964, soliciting bids from twelve computer manufacturers, and stipulating the Lab's requirements in terms of power, price, and delivery date. The new machine, the committee stipulated, needed to have at least 65K of memory, be at least as powerful as Stretch, have a long word length (preferably at or near sixty-bits), offer parallel input/output capability, be a floating-point machine, and be available for delivery in 1965.²²²

²²¹ Worlton, "Survey of LASL Computing Requirements"; Susco, "Computer Selection."

²²² Susco, "Computer Selection"; Dept. of Supply and Property, "Procurement Request," 6-8; "Use of Computing Time by Division"; Lazarus, Oral History, 1995.

The advisory committee comprised ten T-Division staffers (some served as alternates to the six permanent T-Division members) and a member each from the Explosives (GMX) and Weapons (W) Divisions, which were the second and third largest users, respectively, of the Lab's computers. Representatives of the Supply and Property division attended all meetings, and AEC legal counsel Le Roy Warren attended many. Other members of T-1 were invited to vendor presentations, and their comments were taken into consideration. All members of both committees had extensive programming experience, and each had served at least seven years with the Lab. Four members had served on the Stretch Planning Committee, and two others had served on the LASL-IBM Mathematical Planning Committee that guided the logical design of Stretch. Roger Lazarus served as committee chairman.²²³

In 1964, the Joint Committee on Atomic Energy allocated \$7 million for the new Los Alamos computer and its associated peripherals. The AEC held the funds in reserve until the LASL advisory committee could make its recommendation, which the AEC would then review before releasing the funds. Following the cost overruns, delays, and apparent disappointments with Stretch and Livermore's LARC, the AEC increased its control in the 1960s over the computer selection and purchasing process in the lab system, imposing new guidelines for justifying and requesting the purchase of a new system. The most significant guideline fell under the AEC Manual Chapter 1801, which required that all computing companies be permitted an equal opportunity to provide equipment. Los Alamos still retained considerable latitude in choosing its computer

²²³ Susco, "Computer Selection"; Dept. of Supply and Property, "Procurement Request," 22, 40; Mackenzie, "The Influence of the Los Alamos and Livermore National Laboratories," 191.

systems, as no guidelines yet existed for the selection of machines that were not fully tested. The lack of guidelines provided the Los Alamos committee with greater freedom in its selection, but it also left with the Lab final responsibility for evaluating manufacturers' performance and availability claims. The criteria for selection were drafted at Los Alamos, and T Division invited advisory committee members to serve based on their individual experience with adapting and programming their particular branches of weapons research and design for calculation on high-speed computers, rather than any formal guidelines for committee composition.²²⁴

IV. The Search

Of the twelve companies solicited in the advisory committee's bid process, six responded with proposals meeting the basic criteria, including Burroughs, Control Data Corporation (CDC), General Electric, IBM, Philco, and Sperry-Rand Univac. A series of visits and presentations followed over the next two months, as representatives from manufacturers presented bids to the Los Alamos committee, or committee members traveled to sites where computers (at least, those that did not exist only on paper) could be examined and evaluated in person. The committee also provided a series of non-classified test codes, the "P-series" codes, to the manufacturers with working machines. The tests were designed to evaluate each machine's performance on a set of problems the

²²⁴ Fitzpatrick interview with Lazarus; W. J. Worlton, et al., "Reevaluation of Proposals for an Advanced Computer for Los Alamos Scientific Laboratory, December 22, 1965," Edward A. Voorhees Papers, Box 37, Folder 5, Series 27, Request to Purchase CDC-6600, Dec. 1965, LANL Archives; R. B. Lazarus, et al, "On the Selection of ADP Systems – A Position Statement," November, 1970, Edward A. Voorhees Papers, Box 51, Folder 10, Series 38, Selection Criteria, LANL Archives.

committee deemed representative of those in use at LASL. If a manufacturer did not possess a working system, it was instructed to estimate that machine's performance using methods disclosed to the committee, and, if possible, to run the tests on a machine of similar architecture.²²⁵

T Division ran the same P-series codes on Stretch and an IBM 7094 to provide benchmarks. Manufacturers reported that the Burroughs D851 offered an estimated 6.3 times the power of Stretch, a working CDC 6600 scored 5.8, and the IBM 360/70 scored an estimated 1.7. Because the test codes were programmed in long-hand (T Division's name for machine language), D. V. Susco, the W-Division representative, conceded that the results for codes programmed in FORTRAN (the availability of which was a committee requirement for each machine) would be quite different. The test codes were meant as guidelines, rather than as absolute determinants.²²⁶

After a two-week round of additional presentations at Los Alamos, in which manufacturers presented their performance results and provided more detailed configuration and pricing bids to the combined selection and advisory committees, the CDC 6600 (one of only two computers, along with the GE 635, that was operational for testing at that point) and the Burroughs D851 were the advisory committee's early

²²⁵ Susco, "Computer Selection"; L. M. Linson, T-1, to W. J. Worlton, T-1, November 5, 1964, "Completion of timings for 7030 benchmark version of Computer Study Test Problems," Edward A. Voorhees Papers, Box 35, Folder 3, Series 27, Selection Studies, LANL Archives; Dept. of Supply and Property, "Procurement Request," 161-191; Lazarus, T-7, to Whom it May Concern, Dec. 2, 1964, "Computer Selection."

²²⁶ Susco, "Computer Selection"; L. M. Linson, "Completion of timings for 7030 benchmark version of Computer Study Test Problems"; Lazarus, T-7, to Whom it May Concern, Dec. 2, 1964, "Computer Selection"; Running the P-series tests on a working machine did not necessarily result in a more reliable representation of performance than the same tests performed on an extrapolated machine, as CDC and GE both ran the tests incorrectly at first, resulting in lower scores than before Los Alamos personnel stepped in to aid in the process.

favorites. The Burroughs D851 was an evolution of the modular D825, considered to be the first multiprocessor computer. Like the GE 635, which could support up to six CPUs, additional CPU modules could make the D851 a multiprocessor supercomputer. The remaining machines were eliminated due to inferior performance compared to the Burroughs or CDC machines, inadequate word length (in the case of the GE), or projected delivery dates that were too distant in the future.²²⁷

With the Project X computer delayed and becoming increasingly doubtful, IBM had proposed the 360/70, a still-unproduced, large model in its new System/360 series of computers. Announced in April of 1964, the S/360 family became IBM's primary product line, replacing multiple lines of incompatible systems. Proposing the 360/70, with its expected capabilities falling far short of systems that were already, or would soon be, available, was an indicator of the company's struggles in scientific computing. While all-out performance was not LASL's sole concern, the disparity between the Model 70 and its competitors was too large to justify. In addition to its specifications, the S/360's orientation toward business data processing gave some of the Los Alamos committee members pause over its suitability for scientific computing. It also brought into question whether IBM wanted to remain a contender in the scientific-computing space, or was effectively abandoning it for the faster-growing business sector.²²⁸

²²⁷ Susco, "Computer Selection"; Lazarus, T-7, to Whom it May Concern, Dec. 2, 1964, "Computer Selection"; *Programming for the D851 Modular Data Processing System*, Burroughs Corporation (March 1964), 1-2, at archive.org/details/bitsavers_burroughsm51ProgrammingMar64_7923083/page/n1 (accessed April 23, 2019); Philip Enslow, Jr., "Multiprocessor Organization—A Survey," *ACM Computing Surveys* 9, no. 1 (March 1977): 103-129.

²²⁸ Pugh, *Memories that Shaped an Industry*, 248-249; Susco, "Computer Selection"; Pugh, *Building IBM*, 275-276, 298; Lazarus, T-7, to Whom it May Concern, Dec. 2, 1964, "Computer Selection." Scientific computing, which had been the early driver of general-purpose computing development, utilized computers as calculators to solve problems in physics, engineering, biology, and other fields in the sciences and

IBM likely grew concerned when it did not receive a follow up request from Los Alamos about its proposal for the Model 70. Watson, Jr. had committed substantial effort into a larger S/360 model, the Model 90, but in 1964, the 360/90 was at least several years off, even further than the Model 70, and far beyond the delivery date Los Alamos stipulated. With its CEO pushing aggressive tactics to sell the still-ephemeral Model 90, even to the point of selling the machine at a loss to curtail sales of the rival CDC 6600, IBM returned to Los Alamos with a new proposal, offering to install a series of S/360 computers over the following two years, with a more powerful model to replace the previous unit as the larger computers became available. Most of the manufacturers revised their proposals during the selection process, updating the specifications, terms, and pricing on their bids between September and December. IBM's changes were the most dramatic.²²⁹

From October to December 1964, IBM revised its proposals several times, promising ever-shorter delivery times, faster variants of the 360, better deals on service, trade-ins, and peripherals, and significantly reduced prices. IBM's strategy was to buy time for the Model-90 range using financially attractive trade-in offers on lower-powered machines that would be available sooner. William W. Wood, leader of GMX-10, the Statistical Mechanism and Detonation Theory group, was the sole GMX-Division representative on the committee. Wood was suspicious of IBM's revised performance

mathematics. Business computing, by contrast, utilized computers to aggregate and manipulate data relevant to business interests, such as accounting and payroll. Although business generally began to adopt general-purpose computers later than the sciences and mathematics, the pool of potential customers was much larger in business than among scientific users.

²²⁹ Susco, "Computer Selection"; Pugh, *IBM's 360 and Early 370 Systems*, 383; Lazarus, T-7, to Whom it May Concern, Dec. 2, 1964, "Computer Selection."

estimates, with a proposed 360/91 (which replaced the Model 90 on offer, while a 360/92 replaced the 91 in yet another proposal) purporting a ten-fold performance increase over the Model 70, despite the newer system's memory being largely unchanged from the 70. Jack Worlton, on the selection committee, and the T-1 group leader, expressed similar concerns about IBM performance estimates, due to the machine's memory configuration. The selection committee itself made a point of examining the adequacy of the proposed Model 91 memory, but could not conclude its investigation before time constraints meant the committee members had to provide their recommendations in December.²³⁰

On November 16, days after IBM made its revised proposal offering an interim 360/70 to trade for the Model 90, Burroughs changed the D851 delivery date from 1965 to March of 1966, eliminating it from consideration. With the field down to two contenders, the advisory committee invited Control Data to revise its bid in response to IBM. CDC responded with several package deals, much like IBM's, offering more and better disk storage and peripherals, and, in one configuration, the option to trade the 6600 at a favorable rate for the 6800 once it became available in 1967. IBM and CDC made a series of counter-proposals into December of 1964, with IBM's plan for a series of increasingly powerful S/360 models plus peripherals offering less power in the intervening years than the CDC 6600 before installation of the Model 92 in 1966, but

²³⁰ W. W. Wood, GMX-10, to Roger B. Lazarus, T-7, November 25, 1964, "Your Memorandum dated November 23, 1964, Subject: "Computer Selection" Directed to "Whom it May Concern," Edward A. Voorhees Papers, Box 35, Folder 3, Series 27, Selection Studies, LANL Archives; Lazarus, T-7, to Whom it May Concern, Dec. 2, 1964, "Computer Selection"; J. Worlton, T-1, to B. Carlson, April 16, 1965, "Substitute computing equipment in the event of 360/92 failure in speed," Edward A. Voorhees Papers, Box 35, Folder 7, Series 27, Selection Notes Dec. 1964 – Feb. 1965, LANL Archives. IBM's aggressive push of the 360/90, as Emerson Pugh notes, spurred CDC to seek litigation against IBM, resulting in an out-of-court settlement.

costing approximately \$1 million less than the revised Control Data offerings, assuming the trade of the 6600 for a 6800 in 1967. At the end of 1964, the advisory committee members submitted reports to Roger Lazarus with their suggestions, preferences, and, in a few cases, displeasure with the selection process and committee, which were combined into a summary that outlined the far from unanimous purchase recommendation of the committee.²³¹

V. The Selection

The seven official committee reports (one committee member, R. Thomas, served as an alternate for Harry C. Hoyt, but did not report on his own) ranged from one to sixteen pages in length, with most authors summarizing the selection process, and a few voicing their dissatisfaction with it, before providing their recommendation. A programmer and analyst of problems in weapons development from W-4, D. V. Susco's sixteen-page report provided detailed, itemized points where the selection process was apparently deficient in its task, arguing that the selection criteria were not sufficiently specific or definitive, to the point of causing confusion among vendors (IBM, for example, expressed initial confusion over what constituted an "interim" computer), and that it specified hardware and software preferences that were not in-line with the Lab's needs or with what was practical to implement. The lack of specificity over when in 1965 delivery should take place, for example, was one point of contention for Susco, and

²³¹ Lazarus, T-7, to Whom it May Concern, Dec. 2, 1964, "Computer Selection"; "Summary of the Computer Proposal, Dec. 23, 1964," Edward A. Voorhees Papers, Box 35, Folder 3, Series 27, Selection Studies, LANL Archives. The CDC "6800," renamed the 7600, did not enter the market until 1969.

the clause stating that a multiprocessor computer should be a consideration was another, given the extreme challenge of integrating a multiprocessor system into production use at the Lab while the technology remained in its infancy. William Wood included with his report a response to Roger Lazarus' memo of November 23. Lazarus had expressed disappointment that the advisory committee's recommendation would clearly not be unanimous.²³²

Although less vitriolic than Susco's report, Wood cited Lazarus' selection criteria as a possible cause of the perceived disappointment. Wood did not consider a lack of unanimous agreement as a disappointment or failure, nor did Susco. Instead, the lack of consensus, in Wood's view, was a consequence of the speed with which the selection and advisory committees were expected to decide, and a general lack of adequate data about the machines in question, which made it difficult to evaluate the three major systems under investigation. Susco stated that disunity also reflected the diversity of interests and individual experiences within the committee. Contrary to Lazarus' statement in his November 23rd memo that the advisory committee had come to the end of its usefulness, Susco and Wood contended that the difficulty of the selection process meant that a computer study group, of some sort, was still needed. Susco indicated that a computer study group, dedicated to remaining "continually aware of the current status of computing machines and their assorted accessories" should become a permanent fixture to aid future decision processes. (In the late 1970s, the Computing Division formed a dedicated

²³² Susco, "Computer Selection," 3-9; Lazarus, T-7, to Whom it May Concern, Dec. 2, 1964, "Computer Selection"; W. W. Wood, GMX-10, to Roger B. Lazarus, T-7, Nov. 25, 1964, "Computer Advisory Committee, September through November 1964."

supercomputer benchmarking group, tasked with evaluating the performance of computers at the Lab and among vendors. The reports of the benchmarking group not only aided Los Alamos in assessing supercomputer performance, but became a staple resource within the scientific-computing community and its vendors, reducing the kinds of evaluation difficulties the Los Alamos computer selection committees had faced in the 1960s.) The remaining reports, all from T Division personnel, mostly focused on the difficulty of choosing between roughly equal proposals, with a few suggesting a refinement of the criteria specific to the two remaining systems.²³³

Written after Burroughs withdrew, the reports chose between the later IBM and CDC proposals, although some outlined the reasons for eliminating various other machines. Most of the reports offered justification for their decisions, providing perspectives on how the various advisory committee members perceived the technologies on offer, and how they interpreted the attributes of those technologies in relation to the Lab's computing needs and environment. It is this part of the selection process where the divergence of interpretations among committee members regarding the role of computing at the Lab, both current and in the envisioned future, is most discernible. Harry Hoyt, who was responsible for supervising and programming weapons design and numerical analysis problems in T-5, selected the IBM proposal. Stating that, while the cost difference between the two proposals was small, he believed IBM, with its promise of a substantially more powerful Model 92 in 1967, offered the Lab the greatest computing

²³³ Susco, "Computer Selection," 14-16; Lazarus, T-7, to Whom it May Concern, Dec. 2, 1964, "Computer Selection"; W. W. Wood, GMX-10, to Roger B. Lazarus, T-7, Nov. 25, 1964, "Computer Advisory Committee, September through November 1964."

capacity in the long term. Because of the similarity of the existing proposals, Hoyt recommended that the selection criteria be rewritten, so that IBM and CDC could offer more pointed proposals to differentiate their cost/performance differences. Hoyt's recommendation was the shortest, and offered the least justification, focusing entirely on cost versus performance.²³⁴

Tom Jordon, who supervised applied mathematics and problem programming in T-1, acknowledged the attractive economics of package deals, but feared they removed too much flexibility from future procurement decisions. In a reference to the much-changed supercomputer market following IBM's effective withdrawal in the early 1960s, and the commensurate rise of Control Data and others to fill that vacancy, Jordon stated that no one only three years earlier could have foreseen the computers available to the advisory committee. He warned that committing to a package deal at that point in the past would have left Los Alamos with substantially inferior systems in the mid-1960s. He stated his hopes that, if the work towards a common automatic-programming language continued outside the Lab, and caught on inside, switching manufacturers in the future would be less burdensome.²³⁵

Based on his stated viewpoints, Jordon selected the CDC 6600, as it was the most capable machine available without the need for a package deal. While he praised its speed and projected reliability, Jordon admitted that the 6600 would be more difficult to

²³⁴ Dept. of Supply and Property, "Procurement Request," 7; H. C. Hoyt, T-5, to R. Lazarus, T-7, November 30, 1964, "Computer Selection," Lazarus, T-7, to Whom it May Concern, Dec. 2, 1964, "Computer Selection"; "Summary of the Computer Proposal, Dec. 23, 1964," Edward A. Voorhees Papers, Box 35, Folder 3, Series 27, Selection Studies, LANL Archives.

²³⁵ Tom Jordon, T-1, to Roger Lazarus, December 1, 1964, "Computer Selection," in Lazarus, T-7, to Whom it May Concern, Dec. 2, 1964, "Computer Selection."

program and optimize than the 360/70, due to CDC's reliance on "artificial and inflexible" organization techniques for its performance (pointing to the computer's small instruction set and use of peripheral-processing units to increase CPU speeds), rather than advanced component design. The IBM series, Jordon continued, would offer greater compatibility with existing Lab equipment, and the systems themselves were more flexible than CDC's, but trailed in deliverable performance. For Jordon, the initial penalties of the 6600, its incompatibility with existing Lab equipment, and the greater difficulty of programming for its unusual architecture, were worth the flexibility it brought for future computer selection.²³⁶

B. Meixner, a programmer in weapons research from T-2, focused on the cost of re-programming codes for a new machine, in addition to the time required to bring a new computer up to full capacity. Examining the packages from both suppliers, Meixner speculated that the additional power of the 6600 might not be in full use before its 6800 replacement arrived, due to the required optimizing time, thus making CDC's initial power advantage less significant. IBM's proposal of a gradual increase in computing power, Meixner argued, also appeared to be "more in line with the projected needs of the laboratory." In addition, the 360 series offered technical refinements the 6600 lacked, such as an interrupt system, which allowed the IBM to recover automatically from a fault, and the IBM offered a level of compatibility with existing Lab equipment. Remaining an IBM shop would also reduce the number of in-house engineers and varieties of

²³⁶ Tom Jordon, "Computer Selection"; "Summary of the Computer Proposal, Dec. 23, 1964," Edward A. Voorhees Papers, Box 35, Folder 3, Series 27, Selection Studies, LANL Archives; Dept. of Supply and Property, "Procurement Request," 7.

equipment needed to support machines of two different manufacturers. Because of the advertised forward compatibility in the S/360 line, the time spent optimizing code for the 360/70 could, according to Meixner, ensure that its larger replacement entered into full production capacity quickly and efficiently when it arrived. Recommending purchase of the IBM proposal, Meixner concluded with the suggestion that, if his recommendation could not be justified based upon the current selection procedure, the criteria should be rendered more stringent, as additional selection time was preferable to purchasing the wrong machine.²³⁷

Meixner was not alone in his emphasis upon ease of programming and code optimization over potential short-term performance advantages. David F. Woods of T-6, who supervised the programming of problems in weapons design, recommended the IBM proposal for similar reasons. In his view, Woods believed the IBM machine to be "much more pleasant to deal with," a contention that supporters of the CDC proposal did not dispute. Woods contended that bringing either machine up to 20% production capacity within three months of delivery was "impossibly optimistic," citing, like Meixner, that the CDC machine's performance advantage was less significant when placed in the perspective of the time needed to program and optimize a completely unfamiliar machine. According to Woods, optimizing a single one of the Lab's more complex codes for either machine would require more than 100 hours for several expert programmers. The easier programming environment of the IBM, which included a larger vocabulary and a larger memory (131K versus 65K for the CDC) for the money, combined with a

²³⁷ B. Meixner, T-2, to R. Lazarus, T-7, November 30, 1964, "Advisory Committee Member Report on Computer Selection," in Lazarus, T-7, to Whom it May Concern, Dec. 2, 1964, "Computer Selection."

shorter projected interim period to push Woods toward the IBM. He considered the machine construction timeline of the better-established IBM to be more realistic than CDC's. Woods also cited the need for only one set of maintenance facilities if the Lab remained with IBM, if the proposals were otherwise equal, to be sufficient to justify his decision.²³⁸

Susco's estimation of the 6600 effectively echoed Meixner, although he was not as pessimistic about the time and effort needed to bring the computer to production capacity. Susco simply doubted that the Lab required the power the CDC machine offered as early as it did. Based on IBM's technical presentation, Susco found IBM's power estimates of the 360/91 to be unlikely, possibly by a factor of two, but his conservative estimates of the Lab's future computing needs meant that the gradual increases in power found in the IBM proposal, even if halved, were more than adequate. The S/360's larger instruction set than the CDC machine, and the easier transition in the Lab computing environment, were Susco's primary reasons for choosing the IBM path.²³⁹

William Wood, who agreed with Susco in his proposal to Bengt Carlson on a permanent computer study group, with members from outside T Division, diverged in his computer selection. Wood coincided with Tom Jordon's recommendation that the Lab opt for the 6600 without the 6800 option, commenting that it was a mistake to try to see "too far into the future" with supercomputing. Wood acknowledged the less attractive features of the 6600, such as its limited instruction set and difficult programming

²³⁸ David F. Woods, T-6, to Roger B. Lazarus, Chairman, Computer Selection Advisory Panel, December 1, 1964, "Advice on Computer Selection," in Lazarus, T-7, to Whom it May Concern, Dec. 2, 1964, "Computer Selection."

²³⁹ Susco, "Computer Selection."

environment, stating that they were understandable reasons for other committee members to decide against it. Indeed, Wood described the CDC machine as a "tolerable solution," and indicated that the Burroughs machine would have been his first choice, had it still been in the offing, due to its predicted superior programming environment and modular expandability. Stating that he was distrustful of predictions that computing requirements would increase geometrically, Wood believed that the IBM upgrade path would likely be adequate for the Lab's future computing needs, but he opted for the 6600 because it allowed greater flexibility in the future than a locked-in upgrade path.²⁴⁰

In his summary report, Roger Lazarus reiterated the difficulty of judging the IBM and CDC offers against one another, but did not provide elaborate detail explaining his choice of the 6600. Instead, Lazarus provided a chronology of the bid process, and stipulated that both IBM and CDC had features in their favor: a larger computing capacity increase in the near-term for the 6600, and an earlier delivery time, lower price, and more attractive system design for the IBM. The CDC, although not as "pleasant" to work with, was his choice, but only by a "very slim margin." Lazarus was not the deciding vote, as the IBM proposal already carried the majority of the committee. The following February, the Los Alamos Department of Supply and Property issued the committee recommendation to the Atomic Energy Commission, which finalized the official contract with IBM. Not unified in choice or in its priority of the given selection criteria (where the scope and detail of the criteria were deemed adequate at all), the

²⁴⁰ W. W. Wood, GMX-10, to Roger B. Lazarus, T-7, Nov. 25, 1964, "Computer Advisory Committee, September through November 1964."

advisory committee represented several overlapping interpretations of how the available computer options fit within differing visions of computing at the Lab.²⁴¹

Two related points of contention predominated the decision-making process: the ease of integrating and optimizing a new system for production, and the predicted rate of increase in computing demand at the Lab. B. Meixner and David Woods were the most vocal in their emphasis on ease of programming, contending that nominally more powerful computing hardware mattered little when it could not be put to use. Easing the transition of codes and facilities, both in the near- and long-term, was of greater importance than immediate gains in performance. Susco and William Wood argued similar points in their skepticism of the predicted rate of increasing demand, but arrived at differing conclusions. While Hoyt contended that the IBM proposal offered the most performance in the long-term, Susco doubted IBM's performance estimates, but found the estimates unimportant, believing that they exceeded the demand increases he envisioned. The anticipated ease of transition found with IBM was the most salient consideration from Susco's perspective.²⁴²

William Wood agreed that the computing-demand increases the selection committee predicted were likely overblown, but opted for the 6600 for the same reasons as Tom Jordon, warning that planning too far into the future of supercomputing would reduce the Lab's flexibility. Although, Wood made a point in his memo to Roger Lazarus to stress the importance of programming ease, citing that it was not [underline in

²⁴¹ Lazarus, T-7, to Whom it May Concern, Dec. 2, 1964, "Computer Selection"; Dept. of Supply and Property, "Procurement Request," 1-2.

²⁴² Susco, "Computer Selection."

original] the least of all the desirable attributes of computing machinery. He likewise expressed displeasure at the inadequacy of CDC's documentation and of the P-series tests in general, but believed the initial inconvenience of the 6600 was worth the freedom of selection in the future.²⁴³

The decisions and divisions within the advisory committee reflected the variety of issues new to Los Alamos' computing experts since the arrival of Stretch. IBM's changed position within scientific computing left LASL to navigate a more complicated range of technical choices than had been anticipated a few years before, when the SPARC Project was slated to provide a clearer path to Stretch's successor. T Division chose IBM as a collaborative partner in the 1950s because of that company's willingness to develop a computer in tight conjunction with LASL. The capriciousness of IBM management after Stretch, and AEC funding cuts, left Los Alamos as simply a customer (albeit an important one) in the computing market of the mid-1960s, rather than as a development partner, as T Division's expert personnel had initially hoped to repeat on LASL's next supercomputer acquisition.

Choices of vendors and technical approaches were all the more complicated following the repeated upheavals to LASL's mission, including its research goals and methods, since the Eisenhower moratorium and the passage of the Limited Test-Ban Treaty. Computing was, without question, becoming a more important resource to LASL's research community. It was less clear to LASL's computing experts what

²⁴³ W. W. Wood, GMX-10, to Roger B. Lazarus, T-7, Nov. 25, 1964, "Computer Advisory Committee, September through November 1964"; Lazarus, T-7, to Whom it May Concern, Dec. 2, 1964, "Computer Selection."

assortment of problems, and what computational demands, the Lab's Central Computing Facility would face as the decade unfolded. The committee's calculus of convenience versus performance versus flexibility left most committee members unsettled over questions that were left unanswered in the selection process. They were unsettled by what they considered to be too brief a search period, and the daunting (and what some believed to be inappropriate) search requirements T Division established beforehand, given that LASL's current and future computing needs were still far from settled after several years of rapid evolution.

The ever-increasing budgets of the 1950s, and T Division guiding the development of a computing operation primarily for its own benefit, had culminated with the Stretch collaboration. In the 1960s, with LASL's computing requirements and options under greater scrutiny from both the AEC and a much-expanded user community, T Division's control of LASL computing began to come into overt question for the first time since World War II. The narrow selection of IBM in 1965 did not end LASL's unwitting position between IBM and Control Data's battle for market share. Instead, it began a process of radically and permanently expanding LASL's computing operation, and sparked the formation of a new, dedicated computing division.

VI. Changing Plans

After selecting the IBM proposal, Los Alamos spent the remainder of 1965 preparing the CCF for the installation of the first in a series of IBM machines scheduled to arrive starting in November of that year. In September, IBM modified its final

proposal, spurring a series of contract renegotiations even before IBM notified the AEC that the entire S/360 product line delivery schedule would slip—and that one of the machines promised for Los Alamos, the Model 91 J, would not be delivered at all. IBM offered to replace the 91 J with two different models, the last of which, the Model 95, would be delayed fourteen months beyond the previously proposed final machine, the Model 92. IBM likewise withdrew from its proposal a collection of high-performance parallel disks and fast bulk storage that had been required components in the original bid criteria. IBM also announced that key software packages, such as the multiprogramming system, which were among the 360's more attractive features, would not be available for a year or longer after they were promised. As Roger Lazarus later recalled of IBM's changing proposals, "It was the last straw." While IBM's offerings worsened, those of CDC improved, with new high-performance disks entering Control Data's product line in November of 1965. CDC's software offerings for the 6600 also improved, along with the machine's reputation for reliability. A number of production 6600s had entered use in the intervening year, and the computer's initial teething problems appeared to be resolved. The original selection committee, including Jack Worlton, Ed Voorhees, and Chester Kazek, along with T. Mannon, the Lab procurement auditor, rewrote the procurement request, opting for the CDC 6600 and a collection of peripherals for slightly more than the amount in the IBM contract from January.²⁴⁴

Installed in July of 1966, the 6600 provided a significant increase in LASL's computing capacity as it entered production use over the following six months.

²⁴⁴ Worlton, et al., "Reevaluation of Proposals for an Advanced Computer," 1-4, 6-12; Fitzpatrick interview with Lazars.

Optimizing the CDC machine was difficult, as predicted. Los Alamos had to embark on an expensive process of rewriting application software and production codes for the new architecture, which lacked many of the features, such as a large instruction set and significant error-correction features, that LASL users had appreciated with Stretch and the smaller IBM systems. The 6600's architecture placed more of a burden on programmers at a time when supercomputer usage began to escalate rapidly at Los Alamos. With minimal software available from Control Data, LASL programmers and users had little choice but to develop an intimate familiarity with CDC equipment and programming requirements, if they hoped to utilize the 6600's extreme performance capabilities. Livermore did much the same with its four 6600s, but with a significantly different approach to local operating-system development than Los Alamos, an issue that will appear in Chapter 4.²⁴⁵

At the beginning of the search process, T Division expected to continue a pattern of procurement that had begun with Stretch. That is, to purchase a single, large supercomputer every few years to serve as LASL's front-runner system, while augmenting the Central Computing Facility's production capacity with several smaller computers in the between years. Even the most alarmist estimates of how quickly the demand for computing power would increase through the end of the decade fit neatly within this expectation, hence the general lack of concern over IBM's proposal to provide incrementally more powerful computers over several years. By the time LASL's first CDC 6600 arrived, the situation in the CCF had changed radically since the computer

²⁴⁵ "LASL Computer Capital Equipment, June 30, 1968," Edward A. Voorhees Papers, Box 67, Folder 8, LASL Computer Inventories 1961-1973, Series 49, LANL Archives.

search began in 1964. All estimates of computing-demand increases proved to be too conservative in the post-LTBT environment. The escalation in demand was significant to the point that, while the 6600 was being optimized to run LASL's most demanding jobs, Stretch continued to run a large number of small jobs, which comprised the majority of jobs by volume. To alleviate some of the heavy demand on Stretch, Los Alamos entered negotiations to purchase a surplus Stretch from the MITRE Corporation at the end of 1966, but that purchase was never completed.²⁴⁶

While weapons work once again grew in priority at Los Alamos, weapon-delivery technologies and national nuclear strategy changed markedly. One of the primary drives in the AEC weapons labs during the 1950s had been the production of increasingly efficient and compact weapon designs, which were intended for delivery primarily by aircraft. The introduction of the ballistic missile as a delivery vehicle further accelerated the drive toward even lighter and more compact designs, contributing to increased computer usage in the design process to meet tighter tolerances and performance requirements. The MIRV (Multiple Independently Targetable Reentry Vehicles), which began development in the mid-1960s, and entered the arsenal starting in the 1970s, brought even greater design challenges.²⁴⁷

²⁴⁶ The MITRE Stretch that LASL had hoped to purchase went to Brigham Young University in Utah. After its decommissioning in 1971, BYU purchased the Los Alamos Stretch as spare parts. Portions of the MITRE Stretch, containing components from LASL's 7030, survive on display at the Computer History Museum in Mountain View, California; Dave Woods, Jack Worlton, EAV, "Second 7030," 8/5/66, Edward A. Voorhees Papers, Box 21, Folder 3, Series 22, Stretch #2, 1966, LANL Archives.

²⁴⁷ Donald MacKenzie, *Inventing Accuracy: A Historical Sociology of Nuclear Missile Guidance* (Cambridge, MA: MIT Press, 1990), 214-217; Thomas B. Cochran, *Nuclear Weapons Databook: Volume II, U.S. Nuclear Warhead Production* (Ballinger Publishing Company, 1987), 19-20; Rhodes interview with Mark.

The MIRV consisted of multiple warheads placed atop a single missile, with each warhead capable of being guided to a separate target. Although not accepted without controversy among the US armed forces, the MIRV was considered a cost-effective means of defeating potential future Anti-Ballistic Missile (ABM) systems, and as a means of adapting the arsenal to changing US nuclear strategy without returning to the rampant AEC budget increases of the 1950s. While not as powerful as a single, large warhead, smaller, independent warheads improved the likelihood of destroying the larger number of targets stipulated with the emerging "counterforce" strategy of targeting Soviet military assets. This strategy stood in contrast with the previously dominant "counter-value" strategy of targeting cities. MIRV warheads required more exacting design work than their larger counterparts, due to the smaller tolerances, greater demands on efficiency, and the more complex delivery systems.²⁴⁸

Designing and simulating ballistic-missile warheads, particularly MIRV variants, and other increasingly compact weapons packages the DoD specified, resulted in a rapidly escalating number of complex jobs at the Los Alamos CCF. Carson Mark had been correct when he argued in 1964 that the overall staff population of Los Alamos would likely not increase markedly over the remainder of the decade. What he could not have anticipated was the intensified computer usage of the late 1960s that became the norm thereafter. The average number of jobs per month at the CCF increased from 10,400 in 1964 to 21,000 in 1967, and to 32,000 in 1970, the equivalent of running one

²⁴⁸ MacKenzie, *Inventing Accuracy*, 214-217; Cochran, *Nuclear Weapons Databook*, 19-20; See Marshall William McMurrin, "Achieving Accuracy: A Legacy of Computers and Missiles," Xlibris Corporation (2008).

job per staff member at Los Alamos per day. The growing number of jobs was a significant part of the problem, but so were the increasingly complex codes being run on the computers.²⁴⁹

The code complexities that T Division experts in the mid-1960s predicted to enter use in the somewhat distant future, were being phased in more quickly than anticipated. Two-dimensional implosion calculations, for example, entered regular use beginning in 1968, approximately five to eight years earlier than expected. One-dimensional codes were becoming more complex as well. The one-dimensional code described as the "main vehicle for thermonuclear-weapon design" had been modified to allow double the number of cells in a simulation. A "cell" was the prescribed area in which the computer simulated the behavior of subatomic particles based on a statistical or deterministic model, depending on the purpose of the code. The collection of cells, in aggregate, represented the entirety of the space where a detonation or other physical phenomena were being simulated. The more cells, the finer a representation of how materials behaved inside the extreme conditions of a detonating weapon. According to Roger Lazarus, the running time of a code increased "in proportion to the three-halves power of the number of cells." The replacement for the "main vehicle" code, which entered development in the late 1960s, increased the complexity over its predecessor by a factor of ten.²⁵⁰

²⁴⁹ Voorhees and Worlton, "Computer Plans and Requirements," September 1967; R. B. Lazarus, C-DO, to Paul McConnell, ENG-DO, "CDC 7600 Computer Installation, Building SM-132," July 5, 1972, Edward A. Voorhees Papers, Box 48, Folder 6, CDC-7600 Computer #3, Series 35, LANL Archives.

²⁵⁰ Voorhees and Worlton, "Computer Plans and Requirements," September 1967; Office of CADP, "AECM-1801 Information for LASL Advanced Computer System"; Lazarus, "CDC 7600 Computer Installation, Building SM-132."

Lazarus argued that the increases in code complexity that Los Alamos was experiencing in the late 1960s would absorb the additional performance expected with the next generation of computer hardware, even if nothing else changed. While weapons-related calculations occupied approximately 85% of all the computer time at Los Alamos in the late 1960s, non-weapons-related research followed a similar trend in escalating computational complexities and overall demand. The Meson Physics Division, for example, jumped from approximately 15-20 hours of computer time per month in 1967, to over 200 in 1972. The Division's expectation in 1964 of meeting LASL's computer requirements, and even offering excess capacity, through a single, large-scale computer acquired every few years was an impossibility as early as 1966. Immediately following the first CDC 6600 installed in the CCF, Los Alamos acquired three additional 6600 supercomputers between 1967 and 1968. This added the equivalent of 23 Stretches to the CCF in only two years, when the most extreme earlier estimate was a need of only 10. The 6600s did not end the pressure to expand the computer operation, as LASL soon entered contracts to acquire multiple copies of Seymour Cray's follow-on supercomputer, the 7600, and began early negotiations with Control Data on a machine to follow the still-unreleased 7600.²⁵¹

Between 1966 and 1968, LASL's adaptation to the post-LTBT Cold War environment had spurred an unprecedented expansion of the Lab's computing resources. Budget cutbacks remained an issue throughout the AEC lab system in the 1960s, but the weapons program's regrowth as LASL's primary research focus, and its commensurate

²⁵¹ Lazarus, "CDC 7600 Computer Installation, Building SM-132"; Edward A. Voorhees, "Executive Summary, Draft #2, June 5, 1973"; "LASL Computer Capital Equipment, June 30, 1968."

need for computing capacity as a key element of its development cycle, made it easier to justify and obtain funding for supercomputers from the Joint Committee on Atomic Energy. LASL's computer-procurement personnel, such as William Spack, who joined Los Alamos in 1946 as a young army radio engineer, became proficient at arguing before Congress and securing funds for new supercomputer systems on an almost yearly basis. The rapid pace of acquisition, and the growing dependence of the weapons program on the availability of multiple, cutting-edge high-performance computers, also served to discourage LASL from changing vendors.²⁵²

Control Data became LASL's primary supercomputer vendor following the 6600 installed in 1966, largely due to CDC's sizeable performance advantage over its competitors. Although Los Alamos considered bids from other computer companies, including from IBM, before each new system was acquired, Seymour Cray's designs continued to lead the category. With a growing fleet of Control Data supercomputers, LASL's commensurate software-development investment in that company's architectures, and the heavy user demand on the CCF's computer resources, would have made changing vendors during that period extremely disruptive to LASL's production environment.²⁵³

²⁵² William Spack, oral history interview by Nicholas Lewis, Los Alamos, NM (June 25, 2014).

²⁵³ William J. Worlton, "Computer Vendor Meetings, 1965-1968," lab notebook, HPC Archive; Lewis interview with Spack; William J. Worlton, "Computer Plans and Requirements for FY-1969 for the Los Alamos Scientific Laboratory Central Computing Facility," Los Alamos, NM: Los Alamos Scientific Laboratory (September 1967).

VII. The Computing Sciences and Services Division

The rapid escalation of computing demand and computing power at Los Alamos took place concurrently with, and as a driving element behind, a growing interest from multiple parties to reorganize computer-resource management at LASL. As early as 1964, Bengt Carlson had proposed the possibility of spinning off T Division's computing operations to a separate division, suggesting Roger Lazarus as its leader. The proposal was in recognition of the growing importance of computing as a Lab-wide resource, and the strain on T Division, a science division, of maintaining and allocating that resource for a growing, increasingly diverse user base. Carlson, as leader of T-1, had been responsible for allocating computer access to users from all of LASL's divisions since 1951, although T Division had been the primary user of those resources for much of that time. As computer use intensified and expanded beyond T Division, Carlson grew concerned over the appearance of favoritism toward his own division. His efforts to avoid favoritism sometimes brought critique from within T Division from users who believed he went too far in prioritizing computer time to non-T-Division users. As the computer operation became significantly larger and served a wider range of users, the pressures behind its planning and maintenance forced Carlson to devote less and less time to his theoretical physics research. His expertise in computing had been mostly a matter of necessity for his work in physics during the early Cold War at LASL, and he hoped that spinning off the computer operation to its own division would let him return more fully to his chosen field. Lazarus was not interested in leading a division when initially

proposed by Carlson in 1964, and the idea lacked sufficient institutional traction to proceed. As a result, the proposal was temporarily shelved.²⁵⁴

The contentious selection process from 1964 to 1966 forcefully raised the issue of a separate computing division, but this time among users outside T Division. Users primarily from W (Weapon Nuclear Engineering) and GMX (Explosives) Divisions, which were respectively the second and third largest users of LASL's computing resources after T Division, were displeased with what they viewed as T Division's dominance of the selection and advisory committees, and with T Division's selection criteria appearing to emphasize cutting-edge technologies over ease of transition and use. The process appeared emblematic of T Division's traditional approach to computing, which focused on maintaining a high level of expertise in computer technologies and software development primarily suited to T Division's extreme computational needs. For many of the newer users to LASL's computing resources, T Division's approach appeared overly concerned with maintaining the Laboratory's status as a world-respected center of scientific computing, rather than finding readily available tools for the job at hand.²⁵⁵

During the rapid growth of computing demand at LASL, multiple divisions and individual groups had begun to purchase their own computers to cope with the Lab-wide computing demand. The purchases also provided local computer access for areas of LASL remote to the CCF, and allowed individual groups and project teams to develop software and features on architectures of their own choosing, rather than for the machines

²⁵⁴ "New Division Formed at LASL," *The Atom* 5, no. 5 (May, 1968); Fitzpatrick interview with Lazarus; Lewis interview with Folkner; Worlton, "Computer Plans and Requirements."

²⁵⁵ Lazarus, T-7, to Whom it May Concern, Dec. 2, 1964, "Computer Selection"; Lewis interview with Folkner; Norman Morse, oral history interview by Nicholas Lewis, Los Alamos, NM (July 30, 2018).

T Division maintained. GMX and Material Physics (MP), for example, purchased a series of small IBM mainframes and DEC and SEL (Systems Engineering Laboratories) minicomputers in the mid-1960s to provide their groups with time-sharing capability, a relatively new concept that allowed multiple users to share a single computer's resources simultaneously. T Division's technical experts, such as Thomas Jordan and Nicholas Metropolis, followed time-sharing developments closely, but were reluctant to implement the new capability on the production machines of the CCF while the technology was still immature, and while production capacity was already strained. T Division's reluctance to provide time-sharing services on LASL's supercomputers frustrated the cross-section of users who began to request that capability starting in 1964, and provided a point of contention between users in the weapons program and the Lab's computer management that persisted into the 1970s.²⁵⁶

In 1966, Norris Bradbury formed the Office of Coordinator for Automatic Data Processing (CADP) to provide a centralized office for planning and managing computer purchases at Los Alamos, as there was little coordination or formal policy on the rapidly increasing rate at which LASL divisions were acquiring their own computer systems. It was Bradbury's first attempt at unifying LASL's overall computing resources before ultimately forming C Division. With representatives from the five most computer-resource-intensive divisions at the laboratory having representatives on the CADP

²⁵⁶ "Computers, June 19, 1967," Edward A. Voorhees Papers, box 67, folder 8, LASL Computer Inventories, 1961-1973, series 49, LANL Archives; T. Jordan, to B. Carlson, "Thoughts on Time-Sharing and Real-Time Data Processing at LASL, July 27, 1965," Edward A. Voorhees papers, box 60, folder 6, Time-Sharing, 1960s, series 45, LANL Archives; "Time-Sharing, Part I, June 29, 1966," Edward A. Voorhees Papers, box 60, folder 6, Time-Sharing, 1960s, series 45, LANL Archives.

Advisory Committee, the CADP office was also an attempt at quelling the users dissatisfied with T Division's predominance over the supercomputer selection process, and control over LASL's computer planning, in general.²⁵⁷

With multiple divisions establishing separate, often non-compatible computing operations, Bradbury feared a fragmentation of computing investments at LASL. As a consequence, he, too, began to press for a dedicated computing division, hoping to unify the variety of computer operations forming at Los Alamos. In early 1968, Bradbury authorized the establishment of the Computer Sciences and Services Division, or "C" Division, which commenced operation on April 22 of that year, and consisted of about 150 employees making the transition. Roger Lazarus and Thomas Jordan were named as the division leader and alternate leader, respectively. Paul Harper and William Spack served as the first assistant division leaders, and helped to design the new division's operations and policy guidelines. The three computing-related groups of T Division, including T-1 IBM Computer Support (whose functions were divided among the first three C-Division groups), T-7 Computer Research and Development (which operated the MANIAC II and conducted other R&D efforts), and T-13 Statistical Services, were transferred to C Division. Bengt Carlson and several other members of T-1 did not join the transition, and formed the basis of the new T-1 Transport and Reactor Theory group. Nicholas Metropolis left T's Division Office (T-DO) to become a C-Division advisor.

²⁵⁷ N. E. Bradbury, DIR, to Division Leaders and Department Heads, April 15, 1966, Edward A. Voorhees Papers, box 50, folder 3, CADP Responsibilities, series 38, LANL Archives.

William (Jack) Worlton and Edward Voorhees also joined when the Office of Coordinator for Automatic Data Processing (CADP) became a part of C Division.²⁵⁸

C Division was organized into seven groups: C-1 CCF Operations and Support, C-2 CCF Systems Support, C-3 CCF Plant and Hardware Support, C-4 Scientific Applications Services, C-5 Statistical Services and Management, C-6 Numerical Methods and Consultation, and C-7 Computer Science Research (inheriting T-7's R&D functions, including the MANIAC II). With the exception of the MANIAC II, which was still located in the Administration Building, C Division was housed in the recently expanded Central Computing Facility, located in Building 132 (then still colloquially known at Los Alamos as the "Stretch Building") in TA-3, the main administrative technical area of LASL. The CCF had undergone several additions since it had been initially constructed for Stretch, providing space for the Control Data equipment and on-site service technicians, and offices for LASL's computer personnel. The MANIAC II moved from the Administration Building to the CCF in 1969.²⁵⁹

In the official announcement of C Division's formation, Roger Lazarus stated that the new division would be more service-oriented than had been possible for T Division's computer operation. To that end, C Division would offer more programming services, via C-4, C-5, and C-6, for users from all divisions and groups, in order to alleviate the need for other divisions to hire their own programmers. That would curtail the fragmentation of LASL's application development by better optimizing user codes for the

²⁵⁸ "New C-Division, April 17, 1968," Edward A. Voorhees Papers, box 50, folder 3, CADP Responsibilities, series 38, LANL Archives; Fitzpatrick interview with Lazarus; "LASL Computer Capital Equipment, June 30, 1968."

²⁵⁹ "New Division Formed at LASL"; "New C-Division, April 17, 1968."

computers of the CCF, and by offering more comprehensive support options for expert and non-expert users alike. T Division's external programming services had typically been more limited than either T or other divisions had desired. That limitation was due to the difficulty of managing financial support within a science division to offer staff services to other divisions, and due to LASL offering less financial support to computing services, in general, until the rapid escalation of computer use starting in the mid-1960s. Having a separate, service-oriented computing division did alleviate many of the funding issues that had complicated T Division's service offerings, and it reduced the complexity of computer management and use at LASL by providing a single entity responsible for all of LASL's computer resources.²⁶⁰

The formation of C Division was likely not the preferred solution for users who were previously dissatisfied with T Division's computer-resource planning and management, particularly its commitment to intensive hardware and software R&D, which, for some users, appeared to be superfluous to LASL's core research mission. The transfer of former T-Division personnel to make up the majority of C Division's management and staff made a significant change from the T-Division approach unlikely. Along with Lazarus' promise that C Division would be more service-oriented than its predecessor, he argued that another benefit of a unified computing division was the ability to bring LASL's computing engineers to bear on the variety of computing challenges across Los Alamos. The promise carried with it an indication that Lazarus intended for LASL to remain a computer-science R&D center, and perhaps expand that

²⁶⁰ "New Division Formed at LASL"; Lewis interview with Jack Worlton; Lewis interview with Folkner.

capability to suit the needs of many other divisions. Although less than ideal from the perspective of the previously dissatisfied cohort of users from the weapons program, the direct transfer of T Division personnel was unavoidable from an operational standpoint, as T Division possessed the majority of LASL's computing staff and expertise. Bradbury could not have realistically re-created the operation (or replaced the world-class experts in high-performance computing) that T Division had assembled over the course of several decades, even if such a change had been desirable to LASL's upper management, which it was not.²⁶¹

The T Division approach had many supporters at LASL, including the director. Many supporters, often veterans of the Manhattan Project, directly linked T Division's computing efforts with the successful completion of both the first atomic and thermonuclear weapons, and in maintaining US strategic superiority over the Soviet Union. Instead of a significant reorganization, the computing experts in T Division were augmented with more funding and staff dedicated to programming and other support services. For many users, including those who supported the long-standing T-Division approach, the question in 1968 was of how much C Division would continue with the computing management and R&D philosophy of its predecessor.

VIII. Chapter Conclusion

The formation of C Division was LASL's formal recognition of computing as an essential service across all science and engineering divisions, a service that could no

²⁶¹ "New Division Formed at LASL."

longer remain embedded within organizational structures that originated in World War II. As Los Alamos struggled to reinvent itself repeatedly between 1958 and the mid-1960s, when its core mission changed under unprecedented political and technical shifts in Cold War policy, one of the few certainties to emerge was that computing had become a widespread, key resource at LASL, and would remain so, regardless of how the Lab's mission evolved. T Division remained the single largest user of computer resources through the 1960s, but its dominance over LASL's computing operation no longer satisfied the computing requirements for a diversifying range of users across Los Alamos, or suited the changing meanings of computing for those users.

For T Division, computer use and the maintenance of in-depth expertise in computing technologies had been deeply intertwined necessities from its wartime creation. A contingent of newer users embedded in the weapons program did not share this experience, and came to view computing as a means to an end in the changing technical requirements of their jobs, rather than as an integral component and field of study in LASL's mission. As T Division attempted to adjust to the rapidly escalating demand on its computing resources following the contested selection between IBM and Control Data, differing views on resource management, and the differing meanings of computing within the diversifying base of users, resurrected a call for a dedicated computing division. Institutional support grew in connection with other divisions forming divergent computing operations and programming efforts, spurring the director to avoid non-compatible fiefdoms of computing developing at LASL during the 1960s. The Computing Sciences and Services Division directly addressed many of the most

pressing service issues that had come to the foreground with the adoption of Control Data as LASL's primary supercomputer vendor. The continuity in management and other personnel from T Division raised the question among critics and supporters of the traditional Los Alamos approach to computing as to what degree C Division would maintain T Division's computing philosophy as the 1970s drew near. The new division's stance on two of the most pressing technical demands in Los Alamos computing of the late 1960s through the 1970s, time-sharing operating systems and computer networking, would soon answer that question.

-Chapter 4-

Unshared Experiences: Time-Sharing and the First Intervention in Los Alamos Computing

I. Chapter Introduction

Following Norris Bradbury's formation of C Division as a unified computing resource, differing interpretations soon emerged over how LASL's computing operation ought to be developed and managed. Computing had become a critical resource to a broader range of users than previously, resulting in more users who were making demands of LASL's primary computing operation. Supporters of the traditional T-Division approach that placed computing within the sciences found themselves at odds with a contingent of users in the weapons program, spread across several divisions, who were becoming increasingly reliant upon computing due to the diminishing returns on improving the efficiency of weapons designs, particularly as all computer users at Los Alamos faced enhanced accountability pressures from the AEC. These users tended to have fewer ties to the development of LASL's computing capabilities over the previous decades.²⁶²

Newer users in the weapons program were less likely to support T and then C Division's policy that computing was a necessary field of basic research in support of LASL's mission. This difference in views became an important point of critique for users dissatisfied with the traditional Los Alamos approach to computing. In the 1960s, the

²⁶² Mark, "A Short Account of Los Alamos Theoretical Work," 6-8; Westwick, *The National Labs*, 154-155.

emerging field of supercomputer software became a focal point for LASL's local R&D, as computer networking would also become in the 1970s. In this operating and R&D environment, some of LASL's computer users in the weapons program considered the computing operation to be insufficiently responsive to their interests, exemplified with T and then C Division's resistance to time-sharing operating systems on the Lab's supercomputers. Much of the problem for users who felt underserved stemmed from the computing leadership's need to develop and maintain a computer operation capable of simultaneously supporting LASL's primary weapons mission and a variety of other non-weapons applications that had emerged from the Lab's research diversification in the 1960s. Following the unpredicted, intensified need for computing among the weapons-program users in the post-LTBT and PERT environment, C Division's upper management became increasingly resistant to technical and operational changes it perceived as possible threats to Lab computing.

The cautious C-Division approach reflected LASL's long-standing philosophy of pursuing conscientious research practices in juxtaposition to Livermore's riskier lines of research. As a consequence, C Division's considerable R&D capabilities were focused on what the division's leadership considered to be in the best interest of maintaining stable, performance-oriented computing capacity. This focus included an emphasis on developing batch systems for LASL's high-end supercomputer systems, despite growing calls for time-sharing by users in the weapons program. The computing division was not unilaterally opposed to time-sharing on smaller systems, or to new technical and operational changes, but tended to tailor technical options to what C Division's leadership

and computer-planning experts believed would provide the most near- and long-term benefit to the most LASL computer users. The sudden increase in computing demand, which exceeded available capacity in the mid-1960s through the early 1970s, reinforced the drive for caution among C Division's leaders. When a new administration at LASL in the early 1970s spurred significant organizational changes to the weapons program, and time-sharing advocates gained control of the computing operation's purse strings, the ensuing altercation between the newly dominant weapons-program computer users and C-Division management resulted in the first high-level intervention over LASL's computer-operation planning. This intervention began a process that would increasingly place C Division's operational and R&D independence at odds with the demands of the Lab's most powerful computer users, and would ultimately redefine the traditional Los Alamos approach to computing.

This chapter will follow the development of the conflict between the supporters and opponents of C Division's research and operating policies in the early 1970s, highlighting the technical and cultural factors that contributed to the divergence of views and agendas relating to Los Alamos computing. The chapter begins with a discussion of how C Division's inherited policies of practicing caution against disrupting LASL's mission, and of emphasizing local computing R&D in new technical fields, resulted in a distinctly Los Alamos approach to the new fields of supercomputer networking and operating-system development. This approach emphasized batch-oriented networking and operating-system strategies, displeasing weapons-program users who sought time-sharing capability. Weapons-program users, who were increasingly frustrated with C

Division's development choices, began to argue to LASL management that the division's R&D strategies and independence were not in the best interest of the weapons program. Having established the technical choices and issues under contention between C Division's supporters and opponents, the chapter turns to Los Alamos' cultural perspectives on its mission and methods, describing how that unique culture distinguished LASL from Livermore, and how it escalated a disagreement over operating-system development into administrative intervention.

II. A Time to Share

The concept of multiple users simultaneously sharing a single computer's resources originated in the 1950s, and gained traction in the 1960s with the aid of US military funding, advances in software and operating-system development, and the availability of computers capable of supporting a multi-user environment. A time-sharing operating system would switch a computer's runtime very quickly among the different users connected via terminals (local or remote from where the computer was located), while giving each user the impression of being the only one on the system. At a time when computers were rare and expensive, sharing system resources increased the overall amount of user access, helping to maximize the utility of valuable computer time. The concept was first put into use at Dartmouth and MIT, with each developing time-sharing systems between 1962 and 1964.²⁶³

²⁶³ Arthur L. Norberg and Judy E. O'Neill, *Transforming Computer Technology: Information Processing for the Pentagon, 1962-1986* (Baltimore, MD: Johns Hopkins University Press, 2000), 68. For a deeper discussion of the development of early time-sharing, particularly of the Dartmouth Time-Sharing System,

By 1964, Nicholas Metropolis, having returned from the University of Chicago as an advisor for LASL's T Division, and Jack Worlton and Ed Voorhees of the Los Alamos Coordinator for Automatic Data Processing (CADP) Office, were actively investigating the development of time-sharing technologies for LASL. While T-7 and other groups in T Division experimented with the time-sharing concept on DEC minicomputers, Metropolis engaged with experts and companies in the time-sharing field on the state of their projects. Exchanges between Los Alamos and computing industry experts typically took place through trade and academic publications, personal correspondence, and professional conferences. Interest in time-sharing, in general, was growing among US computing experts and centers, and Los Alamos was no different.²⁶⁴

Computing demand brought an end to "self-service" computer use at Los Alamos, with efficiency mandating the grouping of jobs into "batches" placed into a computer's job queue. Batch operations at Los Alamos in the mid-to-late 1960s typically meant that users would bring their jobs by hand to the CCF, sometimes from sites far distant from TA-3 where the CCF was located. With LASL's technical areas spread across over seventy square-miles of AEC-owned territory, time spent traveling to and from TA-3 soaked up hours in a normal workday. A job could be delivered to the CCF in the form of magnetic tapes or punched cards, which users could either punch themselves or have T-1's (later C-1's) keypunch operators punch for them. The CCF service desk would

see Chapter 4 of Joy Rankin, *A People's History of Computing in the United States* (Cambridge, MA: Harvard University Press, 2018).

²⁶⁴ N. E. Bradbury, DIR, to Division Leaders and Department Heads, April 15, 1966; T. Jordan, to B. Carlson, "Thoughts on Time-Sharing and Real-Time Data Processing at LASL, July 27, 1965"; "LASL Computer Capital Equipment, June 30, 1968"; Yost, *Making IT Work*, 154-155.

hand the job over to the computer operators, who would load the appropriate program cards or tapes, along with the cards or tape containing the job data, and place the job into the computer's queue. The computer's job scheduler would organize and execute the jobs in its queue in an order that would minimize the amount of time the CPU remained idle, often having the CPU process a second or even a third job while performing an I/O request or other non-processor-intensive task for another job. In addition to code and application-library work, a large portion of T and then C Division's software-programming effort was focused on optimizing the automated scheduling and processing of user jobs on the supercomputers.²⁶⁵

Once complete, job results were then printed (or punched onto cards to be run again if the previous job had been an intermediary step) for the user to retrieve from a series of labeled shelves located in the hallway next to the service desk in TA-3's Building 132, also known as the "Stretch Building." If an error occurred during a job run, the process would repeat, frustrating both users and the computer operators who had to contend with dissatisfied customers. Larger jobs requiring several hours, and sometimes several days, were often run in segments over multiple nights or on weekends, depending on each division's available quota of time for the week, and the priority of the job. High-priority jobs received preference. Errors on large jobs were especially costly in programming, compiling, and processing time. For smaller jobs at least, time-sharing

²⁶⁵ Job scheduling capabilities and procedures varied by the type of job being run, its resource requirements, and the capabilities of the computer itself. The CDC 6600 and 7600 supercomputers, for example, offloaded I/O tasks from the CPU to Peripheral Processing Units, which required fewer interruptions of CPU processing time for basic system calls than in previous supercomputers; "Automatic Digital Computers"; Savage, "Understanding Stretch," 4-6; Lewis interview with Christman; Lewis interview with Folkner.

appeared to be a promising path for improving computer access at LASL, while reducing waiting time, foot traffic through the CCF, and other frustrations. Time-sharing also raised the possibility of interacting with jobs as they ran.²⁶⁶

Computer users from about ten LASL groups distributed among P (Physics), CMB (Metallurgical Chemistry), GMX (Explosives), and J (Field Testing) Divisions began to press for time-sharing solutions on the CCF's supercomputers. The primary issues for most of these groups was of remoteness from the CCF, and the desire to determine if time-sharing was compatible with their workloads. The P-2 Research Reactor Experiments group, for example, submitted about 175 jobs to the CCF per month, but because the group was housed at Omega site, about a twenty-minute round-trip drive from the CCF, the group lost a significant amount of time due to travel. Most of these groups ran smaller jobs that had typically been processed on desk calculators. J-15, the Hydrodynamic Yield group, hoped that access to a time-sharing service would result in a faster turnaround for its time-sensitive jobs. But users with heavier computing demands, such as groups focusing on compute-intensive theoretical-physics research, were wary about the introduction of time-sharing on the CCF supercomputers, fearing the loss of computing capacity.²⁶⁷

Time-sharing became a frequent topic at the CCF User's Group meetings, where T- and then C-Division representatives met with members of other divisions about the

²⁶⁶ "Automatic Digital Computers"; "LASL Computer Capital Equipment, June 30, 1968"; Lewis interview with Folkner; William J. Worlton, "A Proposal for Time-Sharing Service, May, 1969," Edward A. Voorhees Papers, box 50, folder 5, 1801 Studies, 1968-1970, series 38, LANL Archives.

²⁶⁷ Worlton, "A Proposal for Time-Sharing Service, May, 1969"; D. V. Susco and P. P. Whalen, W-4, to Bengt Carlson, T-1, through H. M. Agnew, W-DO, "W-4's Projected Computer Usage Time," July 31, 1967, Edward A. Voorhees Papers, box 6, folder, 5, File 7 W-Division 1966-1969, series 9, LANL Archives.

CCF computing operation. Despite the general interest in improved computer access, introducing time-sharing on the supercomputers in LASL's classified environment was a complicated problem in the late 1960s, due to the immaturity of time-sharing technology and the demands on the existing batch operation. As a consequence, some of the first independent computer operations among LASL divisions outside T Division were made up of minicomputers and small mainframes from DEC, Systems Engineering Laboratories (SEL), and IBM, which offered limited time-sharing capabilities for jobs that did not require supercomputer capabilities. T and then C Division made provisions in their proposed budgets for Fiscal Year 1969 and onward for the purchase of time-sharing equipment, including a dedicated time-sharing computer, in order keep that option open. Despite the technical and organizational hurdles that would have to be overcome, no one in the computing operation or its management was opposed to the concept of time-sharing on non-supercomputer systems.²⁶⁸

The construction of private housing in Los Alamos and nearby White Rock in the mid-to-late 1960s, and the consequent hiring of additional computer operators, made it possible for T Division to run the CCF computers for consistently longer hours per day than previously. By 1967, the CCF production computers—which included two CDC 6600s, two IBM 7094s, and the 7030 Stretch—were operating twenty-four hours per day, seven days per week. T-1 also operated two IBM 1401 data-processing computers for

²⁶⁸ N. Metropolis, T-DO, to L. Rosen, MP-DO, "Comments on a Proposal for Time-Sharing," November 21, 1966, Edward A. Voorhees Papers, box 61, folder 2, Time-Sharing Correspondence, series 45, LANL Archives; "Computers, June 19, 1967"; "LASL Computer Capital Equipment, June 30, 1968"; Susco, and Whalen, "W-4's Projected Computer Usage Time"; E. A. Voorhees, and W. J. Worlton, "Computer Plans and Requirements for FY-1969 for the Los Alamos Scientific Laboratory Central Computing Facility, September, 1967," cabinet 8-2 (Jack Worlton), HPC Archive.

administrative purposes. Plans to relinquish the second 6600 after LASL acquired Control Data's follow-on computer to the 6600 were quickly abandoned, due to escalating demand. The size of LASL's supercomputer operation, with a minimum of four or five large production machines, was no longer a temporary anomaly, as had previously been thought, but a permanent feature. While other parts of LASL experienced cutbacks due to wavering federal budgets, computing grew apace.²⁶⁹

As a technical note, while the computers were on a twenty-four-hour schedule, the total operational time of a computer was always less than its "switched on" time, due to scheduled and unscheduled maintenance and other errors. Scheduled maintenance was always announced ahead of time for users. The total operational time, or "good time," as described in the operator's logs, of the 6600s steadily improved after their installation, reaching about 95% of their total "on" time. Meanwhile, the "good time" on Stretch and the IBM 7094s tended to decrease month to month over the late 1960s, because, as those systems aged, they required more unscheduled maintenance and overall downtime. Stretch caught fire several times toward the end of its service life, severely impacting its "good" time. With the conversion or development of codes compatible with the Control Data machines, the IBM computers were relegated to older codes and smaller jobs, with

²⁶⁹ Chambers, "Technically Sweet Los Alamos," 375-377. The disposal process of transferring AEC townsite properties (including the former technical area around Ashley Pond) to Los Alamos County took place over a five-year period, starting in 1962. The county charter, providing management autonomous from the AEC, was approved in 1968. Much of the impetus for divesting the AEC of managing the Los Alamos/White Rock townsites did not come from residents of the county, but from members of the US Congress who disliked the "socialist" appearance of the US government owning and managing a civilian community. "Machine Usage by Group and Division, Nov. 1966," Edward A. Voorhees Papers, box 66, folder 1, IBM 7030, 7094, CDC 6600, Machine Usage, 1966 - 1968, series 48, LANL Archives; "Machine Usage by Group and Division, July, 1967," Edward A. Voorhees Papers, box 66, folder 1, IBM 7030, 7094, CDC 6600, Machine Usage, 1966 - 1968, series 48, LANL Archives;

the most demanding problems going to the 6600s. Even without the fire hazard, the IBM equipments' days were numbered, due the valuable space the large machines occupied on the CCF production floor.²⁷⁰

Two additional CDC 6600s installed in 1968, during and after the transition to C Division, were also soon operating around the clock. Before the arrival of those systems, C Division continued negotiations begun under T Division to purchase CDC's follow-on supercomputer, the 7600, for a cost of \$5.3 million, including input-output equipment. Once ordered, C Division and the AEC fast-tracked the purchase of a second 7600, expected to arrive shortly after the first, in order to cope with the computationally demanding two-dimensional codes entering use. The AEC approved the lease of the fourth 6600 (LASL had purchased the other three) as an interim machine before the arrival of LASL's first 7600, with the stipulation that the leased 6600 would be transferred to Sandia after the 7600 arrived. The LASL 6600 became Sandia's first supercomputer. Like the first two 6600s, the second two were also quickly swamped under the growing volume and complexity of jobs.²⁷¹

While negotiations took place at LASL for the first and second 7600s, C Division was already in regular contact with Control Data about Seymour Cray's expected replacement for the (still-unproduced) 7600, called the 8600. Since the purchase of its first 6600, LASL had maintained close ties with Seymour Cray, whose extraordinary

²⁷⁰ "Machine Usage by Group and Division, Dec., 1967"; Voorhees, and Worlton, "Computer Plans and Requirements for FY-1969."

²⁷¹ "1801 for CDC-7600, Dec. 17, 1968," Edward A. Voorhees Papers, box 48, folder 4, Advanced Computer Information, series 34, LANL Archives; Voorhees, and Worlton, "Computer Plans and Requirements for FY-1969."

faculty for scientific-computer design, and affable personality, earned him numerous friends and admirers at Los Alamos, particularly among the computing staff. Livermore, with its own growing contingent of Control Data supercomputers, had been in negotiations with the company since 1966 for the production of a "parallel" supercomputer system, being developed by a CDC team separate from that of Seymour Cray's. T and then C Division staff followed the development of the new system, as its new architecture targeted some of the most significant performance bottlenecks found in conventional serial processing. The original Livermore contract predicted a deliverable computer in 1969, but delays plagued the system's development, even while the real-world capabilities of its radical new approach to supercomputing remained to be seen. Given Seymour Cray's established design track record as of the late 1960s, LASL's computing personnel remained largely committed to the 8600 as the Lab's preferred, next-generation supercomputer.²⁷²

The rapid expansion of the CCF's fleet of supercomputers, and the growth of the associated support staff, in the late 1960s is a particularly relevant indicator of the central role scientific computing had attained at LASL in the PERT and LTBT era. Starting in

²⁷² Details about the distinction between serial and parallel processing methods will be addressed later in the chapter. Lewis interview with Spack; Lewis interview with Folkner; W. H. Brummett, Jr., Chairman, Contract Board, to Charles J. Purcell, Staff Specialist, Control Data Corporation, Oct. 5, 1966, "Request for Proposal for a Parallel Digital Computing System," Control Data Corporation Collection, Atomic Energy Commission (AEC): Lawrence Livermore Laboratory (LLL), 1966-971, collection number 80, box 33, folder 19; C. J. Purcell, to J. E. Thornton, September 19, 1968, "Initial STAR Configuration/Livermore," CBI 80, Control Data Corporation Collection, Atomic Energy Commission (AEC): Lawrence Livermore Laboratory (LLL), 1966-971, collection number 80, box 33, folder 19, CBI; Duane Harder, C-4, to R. B. Lazarus, C-DO, September 3, 1971, "STAR News from Joe Requa (LLL) September 3, 1971," Edward A. Voorhees Papers, box 45, folder 2, CDC-STAR Evaluation, series 32, LANL Archives; N. Krisa and F. Schilling, C-2, to Distribution, September 23, 1971, "Summary of Meeting to Discuss the CDC STAR 100 Computer," Edward A. Voorhees Papers, box 45, folder 2, CDC-STAR Evaluation, series 32, LANL Archives.

1965, Los Alamos and the other AEC labs faced a new round of budget cuts, enduring into the early 1970s, with each of the labs enduring layoffs as a consequence. Livermore experienced a 9% budget cut in 1965, slightly larger than that at Los Alamos, while Oak Ridge and Argonne suffered cuts of 25% or more. In the 1960s, the opening up of reactor development, which had been exclusively the province of the AEC, to the private sector resulted in heavier losses at Oak Ridge and Argonne. The budget cuts were also a consequence of changing political and social directions, with the Vietnam War, the Apollo Program, and new social programs diverting funds from the AEC labs. The directors of each AEC lab continued to meet annually to discuss and coordinate their respective strategies for adapting to the new political and economic climate. Los Alamos and Livermore each sought new and unique research capabilities (such as LASL's Meson Physics Facility) to distinguish the weapons labs from each other, while also emphasizing their core missions, in order to set the weapons labs apart from the non-weapons AEC labs. These drives at Los Alamos and Livermore, despite the shrinking budgets, ensured that demand and funding for computing at both labs increased into the 1970s.²⁷³

As a function of demand and funding, LASL's supercomputing capacity quadrupled every two years in the late 1960s, rather than the predicted doubling found in the worst-case scenarios of the mid-1960s. The computers were also able to operate more or less continuously thanks to the availability of more personnel. Despite these changes, increases in demand, primarily from intensive users in W (Weapon Nuclear Engineering), T, and J (Field Testing) Divisions, made access a worsening proposition.

²⁷³ Westwick, *The National Labs*, 274-276.

A contingent of users from several divisions advocated time-sharing to contend with the access issue. Metropolis, along with Carlson, Voorhees, and the other experts in LASL's computing operation with whom Metropolis consulted, believed that time-sharing technology had great potential, but was not yet ready for use at Los Alamos. When the first formal user requests for time-sharing in the CCF began to cross Norris Bradbury's desk in 1966, Metropolis elaborated on the situation through an extensive, documented exchange with Louis Rosin, leader of the new Meson Physics (MP) Division, and a strong advocate of time-sharing. Metropolis argued that, as of 1966, the infancy of time-sharing offered little benefit, but posed significant risks, to LASL's programmatic-computing operation. Although writing in 1966, the issues with time-sharing remained relevant for T and then C Division over the next several years. Although many of the technical issues were addressed, or at least explored at LASL between 1968 and the early 1970s, they still played a significant role in C Division's operational planning during that time.²⁷⁴

Drawing from a survey of the time-sharing field he conducted with the aid of computing experts from Los Alamos and Sandia, Metropolis cited a list of issues with

²⁷⁴ The rapid increase in computer demand in the late 1960s from the Field Testing (J) Division, making it one of the Lab's largest computer users, represented LASL's adaptation to the importance of computing power in the design and testing of nuclear weapons following PERT and the Limited Test Ban Treaty, and the return of the weapons program as LASL's primary mission focus. While Ploughshare research after the LTBT initially underwent modification at Los Alamos and Livermore, moving from experiments with large-scale earthmoving, to a focus on smaller-scale, underground projects, such as petroleum and natural gas fracking and storage, Ploughshare research was abandoned in the late 1960s. This contributed to the return of weapons research as the primary source of funding at LASL and Livermore. The sudden requirement for computing capacity from relatively new users connected to the enlarged weapons program, such as J Division, also demonstrates why it was difficult for LASL to anticipate how computer demand would change over the 1960s and early 1970s; "1801 for CDC-7600, Dec. 17, 1968"; "LASL Computer Inventory, July 1, 1969," Edward A. Voorhees Papers, box 67, folder 8, LASL Computer Inventories 1961-1973, series 49, LANL Archives; Metropolis, T-DO, to L. Rosen, MP-DO, "Comments on a Proposal for Time-Sharing."

time-sharing significant enough to warrant waiting for the technology to mature, at least before introducing it to the production environment. One key issue was the performance drain that time-sharing-capable operating systems caused when continuously switching a computer's runtime between users. The few existing time-sharing OSES were not particularly efficient at resource utilization, particularly when compared to the highly optimized operating systems and resource-management software LASL employed on its production computers. A loss of performance might have been more acceptable, in exchange for greater computer access, were the CCF supercomputers not already at capacity, and if the future rate of change for demand had been more certain. As demand escalated in the late 1960s, Roger Lazarus, Thomas Jordan, and some of LASL's most intensive computer users became increasingly pessimistic in their views of how time-sharing could reasonably fit within Los Alamos' classified operation.²⁷⁵

Lazarus, who followed Metropolis as the head of the T-7 Computer Research and Development group before the formation of C Division, understood T Division's perspective on time-sharing and remote-access systems. Lazarus was less hopeful than Metropolis about time-sharing as a possible solution for LASL. In Lazarus' view, utilizing a supercomputer to drive a multi-user environment defeated the purpose of supercomputing. He was not alone in this view, shared by others both in and outside of C Division. D. V. Susco and Paul Whalen of the W-4 Small Weapons Theoretical Design Group, for example, made clear in a report to Bengt Carlson and their division leader (future LASL director Harold Agnew) that W-4 was "dubious" about the proposed value

²⁷⁵ Metropolis, T-DO, to L. Rosen, MP-DO, "Comments on a Proposal for Time-Sharing"; Fitzpatrick interview with Lazarus.

of time-sharing. The W-4 group believed that time-sharing would negatively impact the performance of production codes on the CCF supercomputers. While opposed to classified time-sharing, Lazarus was more open to changes in access technologies than some of his contemporaries at LASL.²⁷⁶

Other users and C Division personnel who were not as opposed to time-sharing in the CCF were still often wary of its introduction at such an early stage in the technology's development, as they indicated in the CCF User's Group meetings. Their wariness was largely due to the heightened computer demand, and concerns that CPU time devoted to uses other than number crunching could mean needlessly diverting mission-critical resources when there were none to spare. Personnel in the computing operation had a long-standing concern over jeopardizing LASL's mission in the event of technical or other issues in the CCF, and the possible, premature introduction of time-sharing effectively raised those concerns. Although Lazarus and others in C-Division management were not opposed to remote access, or to time-sharing on the minicomputers or small mainframes that were becoming a significant part of LASL's computing power, their reticence about time-sharing in the classified supercomputer environment became a focal point for critics.²⁷⁷

Although performance and resource management were the primary points of contention, Metropolis cited several other technical issues with its introduction at Los

²⁷⁶ Susco and Whalen, W-4, to Carlson, T-1, through Agnew, W-DO, "W-4's Projected Computer Usage Time."

²⁷⁷ Fitzpatrick interview with Lazarus; "LASL Computer Capital Equipment, June 30, 1968"; John Morrison, oral history interview by Nicholas Lewis, Los Alamos, NM (July 1, 2014); Lewis interview with William J. Worlton.

Alamos. One issue was the lack of a terminal infrastructure at the Laboratory, and a dearth of suppliers who offered affordable terminal hardware. At the time Metropolis was writing in 1966, T Division was investigating possible terminal suppliers, but the available options were still inadequate for LASL's scientific-computing purposes. In any case, once the appropriate equipment was found, creating a terminal network connected to the CCF still required time for securing purchasing funds, and time for installation. Security was also a time-sharing issue that Metropolis could not ignore. "Computer security" in the mid-1960s was almost entirely limited to restricting physical access to a computer and its data storage, and the prospect of sending classified information between sites electronically posed considerable concern for T and then C Division, as the computing operation had to comply with AEC requirements for the handling of sensitive data. There was also, as yet, little means of keeping one user's data truly secure from other users on time-sharing systems, a problem that required both technical and procedural resolutions before a shared OS environment could be implemented at Los Alamos.²⁷⁸

The final significant issue for Metropolis was that, even if LASL chose to offer local time-sharing capability, there was no hardware or software available for creating a supercomputer-based time-sharing system. Such a system would have to have been created from scratch, at unknown expense, and without knowing if such a system was compatible with LASL's extensive software libraries, code base, or typical workload.

²⁷⁸ Metropolis, T-DO, to L. Rosen, MP-DO, "Comments on a Proposal for Time-Sharing"; Jeffrey R. Yost, "The Origin and Early History of the Computer Security Software Products Industry," *IEEE Annals of the History of Computing* 37, no. 2 (April-June 2015): 46.

Although T Division was, as previously discussed, amenable to in-house R&D efforts, creating a time-sharing system compatible with LASL's existing computer operation would be a daunting task. Argonne, which primarily focused on unclassified research, had attempted to use a time-sharing service from General Electric (the leader in commercial time-sharing services in 1960s), but quickly abandoned it. Oak Ridge was in early discussions with a new time-sharing company, but no plan had yet to emerge on how to implement the service. Livermore's Computation Department, which had taken an early interest in the MIT time-sharing project, was in the process of writing its own time-sharing operating system for the CDC 6600, but the early results were not satisfactory for production purposes, and development was expected to continue for an indeterminate period of time. Ultimately, Livermore would not complete its first production-worthy time-sharing OS until the early 1970s. Metropolis argued that LASL's best option, in light of the deficiencies and uncertainties of time-sharing in 1966, was to wait for the field to mature, so that Los Alamos could "learn from the mistakes of others."²⁷⁹

By 1967, T Division management still considered time-sharing to be too "radical" for use on the CCF supercomputers in the near-term, but a remote-access system had clear benefits, despite the technical challenges it would pose for a secure facility. Lacking ready-made solutions available for supercomputer remote-access, LASL would have to develop its own. Carson Mark and Bengt Carlson approved a T-1 proposal for a new R&D project, a project that had multiple goals in mind, not the least of which was

²⁷⁹ Metropolis, T-DO, to L. Rosen, MP-DO, "Comments on a Proposal for Time-Sharing."

bridging some of the gap with users on the time-sharing issue. A collaboration between T-1 and P (Physics) Division (groups from which would be the initial testers), "MUX," the Multiple User eXperiment, was intended to provide LASL users with remote-access to the CCF, while also providing LASL with hands-on experience with the incipient field of supercomputer networking. T Division produced most of MUX's infrastructure locally, with T-1 staff members Don Myers as its primary designer, Tom Gardiner as its electronics fabricator, and Ron Christman as its software developer. As its moniker implied, MUX was an experiment with integrating time-sharing-like capabilities with LASL's typical workload, and was intended to be a more cautious move into the new field than Livermore's effort at developing a full-blown time-sharing OS, which had yet to yield useable results. Its specifications completed in June 1967, the MUX system first entered use in July 1968 under the auspices of the new C Division. Coming "on the air" in two phases, MUX's software was completed and modified based on user feedback.²⁸⁰

Starting modestly with eight users, the MUX system after refinement provided up to 75 remote-terminal users the ability to run batch jobs on one of LASL's CDC 6600s. Provisioning the MUX 6600 for only unclassified work sidestepped security concerns and helped to satisfy some of the growing demand for a supercomputer for uncleared users, as they were not allowed secure access to the CCF. Roger Lazarus coordinated LASL's security plans with the AEC Department of Military Application to resolve any

²⁸⁰ R. Christman, T-1, to C. Folkner, T-1, "MUX Remote Console Software, August 18, 1967," Edward A. Voorhees Papers, box 5, folder 7, T-1 Group Correspondence 1967, series 9, LANL Archives; Jeffery L. Peterson, "MUX, adieu, faithful friend," *The Atom* (October 1979): 14-16. "On the air," the expression from radio broadcasting, appears frequently in the Los Alamos computing documentation of the 1950s and 1960s when referring to a new computer or system becoming available for use. See: R. Christman, "MUX System Description, Jan. 9, 1968," Edward A. Voorhees Papers, box 5, folder 8, T-1 Group Correspondence, series 9, LANL Archives.

remaining questions about remote-terminal security before MUX entered operation.

MUX also eventually provided computer access at Los Alamos to non-LASL visitors from the US and abroad, a provision that had not previously existed. The MUX service was designed to host a mixture of remote terminal types, including typewriter- and CRT-based consoles, but most MUX users operated teletype terminals equipped with punched-card readers. These teletype terminals offered a maximum data-transmission speed of 300 bits per second. The MUX software provided a limited file-manager capability, which allowed users to submit their files to the 6600, modify and combine their files in the 6600's storage, enter their files into the batch queue, view files they had submitted or that had been produced by executing a job on the system, and save files for later use. MUX allowed users an unprecedented level of control over their jobs at LASL, while providing a similar level of convenience as a local minicomputer, but with the processing power of a supercomputer.²⁸¹

The total budget for terminals and other equipment reached \$2.5 million. MUX itself consisted of about \$50,000 worth of hand-built hardware, using 8,000 integrated circuits and two miles of twisted-pair cable. The hand-built hardware comprised the "wired logic" that provided the interface between the user terminals and one of the 6600's

²⁸¹ Non-LASL visitors were granted occasional access to the MANIAC II, but its performance was far from that of a supercomputer, and, as an R&D platform, it lacked compatibility with the software library developed for the Control Data machines. Those who borrowed time on the MANIAC II typically had a research-driven purpose behind coming to Los Alamos to use that specific machine, such as a preference for the MADCAP programming language created for the MANIAC II; Christman, T-1, to Folkner, T-1, "MUX Remote Console Software, August 18, 1967"; Roger B. Lazarus, C-DO, to Mr. Charles Winter, U.S. Atomic Energy Commission, December 5, 1968, Edward A. Voorhees Papers, box 6, folder 6, C-DO Correspondence April 1968-Dec. 1968, series 10, LANL Archives; Christman, "MUX System Description, Jan. 9, 1968"; Ronald D. Christman, "MUX, a Simple Approach to On-Line Computing," *Communications of the ACM* 15, No. 5 (May 1972): 319-329.

Peripheral Processing Units (PPU). As designed by Seymour Cray, the PPUs handled I/O and housekeeping functions on the 6600, and were capable computers in their own right, making the 6600 an early example of using "parallelism" to increase computing performance. Although programming the PPUs was difficult, as they could not be programmed directly, only via the 6600, T Division's programmers were able to reprogram the PPUs for new functions. To the 6600's CPU, MUX data sent from the reprogrammed PPU appeared to be from local users, despite traveling from remote terminals through the MUX logic hardware.²⁸²

In addition to the custom software on the PPU handling I/O, custom MUX application software on the 6600's CPU placed the remotely submitted jobs into the supercomputer's queue for processing. A product of T Division's long experience creating job-optimization software, MUX automatically organized job execution to reduce idle time on the 6600, switching between jobs being processed as needed. The application then returned completed jobs to users via the terminal interface, which could print the results, or on punched cards, depending on the user's selection. The user's job files could also be viewed and stored on the 6600's data-storage units. The MUX software, which resided on a modified version of CDC's SCOPE 3 operating system, provided a hands-on test of the file-protection schemes that would be needed for classified information, clearing the way for a secure-side remote-access system. The decision to use Control Data's native OS to host MUX directly influenced C Division's

²⁸² Voorhees and Worlton, "Computer Plans and Requirements for FY-1969"; Peterson, "MUX, adieu, faithful friend," 14-16; Christman, "MUX System Description, Jan. 9, 1968"; Ronald D. Christman, oral history interview by Nicholas Lewis, Los Alamos, NM (July 2, 2014).

OS-planning strategy, because Control Data's updates to SCOPE tended to "break" the custom LASL software. The problem convinced the division's leadership that LASL was better off pursuing its own, custom OS development efforts in the years ahead.²⁸³

MUX was well-received among members of the CCF User's Group, and garnered a loyal following over its decade in service, but it was not a time-sharing system, as users could not interact with their jobs once they were submitted via remote terminal. MUX served a pent-up demand for an unclassified, remote batch service, and typically operated at full capacity twenty-four hours per day. The heavy demand for MUX, and the resulting call for an unclassified time-sharing system, spurred Roger Lazarus to subscribe LASL to its first time-sharing service. Time-sharing services had matured significantly since Nicholas Metropolis' first examination. Other AEC laboratories were already subscribing to a variety of time-sharing services to offset the demand on their own unclassified computer operations, while avoiding the need to develop or maintain such a service locally. Provided by General Electric at a cost of about \$10 per hour of terminal time, and about \$18 per minute of CPU time, the LASL subscription service ran from LASL teletypes or leased acoustic couplers to a GE mainframe in Albuquerque. This provided unclassified LASL users and contractors with true time-sharing capabilities, and gave C Division an opportunity to examine which, and how well, LASL jobs were compatible with the time-sharing paradigm before the development of a local service.²⁸⁴

²⁸³ Christman, "MUX System Description, Jan. 9, 1968"; Ronald D. Christman, oral history interview by Nicholas Lewis, Los Alamos, NM (July 2, 2014).

²⁸⁴ Ron Christman, one of MUX's three original developers, recalled that the system was so unusually reliable that, when it did fail on rare occasions, its maintenance staff had to "re-learn" how the system worked in order to repair it; Worlton, "A Proposal for Time-Sharing Service, May, 1969"; Ronald D. Christman, "The LASL Computer Network," Los Alamos, NM: Los Alamos Scientific Laboratory, LA-

MUX and the time-sharing service met some of the unclassified user demand, but classified time-sharing advocates argued that neither MUX nor the subscription provided the classified, true time-sharing environment they had wanted. While not a classified or a time-sharing system, C Division's development of the MUX service provided solutions for many of the technical challenges standing in the way of time-sharing and other remote-access services being offered in LASL's classified setting. The success of MUX persuaded C Division's leadership to move forward with more ambitious, local networking projects, with MUX providing evidence that networking technologies had the potential to address a range of computing issues at LASL. The first issue was of improving access to the CCF, and the other was the unrelenting expansion of data storage at LASL.²⁸⁵

III. "We can do it better."

With the success of MUX in the unclassified environment, Roger Lazarus approved a plan to apply C Division's new, hands-on experience with supercomputer-networking and remote-access technologies toward multiple problems in the classified setting, including access and data storage. Lazarus' interest in utilizing networking technologies to reform CCF access was part of an ongoing reorganization of C Division's supercomputer management strategy, which included both technological and procedural changes meant to cope with LASL's ever-increasing supercomputer demand. The most

5103-MS (December, 1972), 1-3; Paul Harper, Chairman, "CCF User's Group Meeting, February 6, 1969," Edward A. Voorhees Papers, box 59, folder 1, CCF User's Group, 1969, series 43, LANL Archives.

²⁸⁵ Worlton, "A Proposal for Time-Sharing Service, May, 1969"; Christman, "The LASL Computer Network," 1-3.

significant procedural change Lazarus instituted after becoming C-Division leader was the introduction of computer-time allocation reforms starting in 1969. The reforms were directed at a time-quota system that Bengt Carlson introduced in 1965, which had provided the first time-based access controls on the CCF computers to cope with the surging demand stemming from the PERT and LTBT requirements. The 1965 quota system provided the eight major user divisions with a designated number of daily hours, Monday through Saturday, per computer. Sundays were part of a "Reserve Quota" to reapportion time lost during the week to computer faults and maintenance, and to provide a buffer for divisions that did not have an established quota, or had not been able to secure computer time during the week.²⁸⁶

The quota system was not about limiting excess computer use so much as providing the largest users with guaranteed time each day. This compared with the competition and chance inherent to a first-come, first-served basis, which caused consternation among all users, but particularly among those in need of fast turnaround times, and predictable job-run schedules. The time apportioned was based on monthly previous-use statistics for each division, with the allocation changing as the needs of each division evolved. The established quota system favored the largest users of the CCF, who had a vested interest in maintaining their given allocation from month to month. As the number of new and smaller-scale users grew dramatically in the late 1960s, a simple time-quota system that overwhelmingly favored the traditional, heavy users of the CCF

²⁸⁶ B. Carlson, T-1, to Division Leaders and Group Leaders, "Assignment of Computer Time on a Quota Basis," August 26, 1965, Edward A. Voorhees Papers, box 67, folder 7, Quota, Time Assigning and Scheduling, 1963-1965, series 48, LANL Archives.

supercomputers was no longer adequate. In fact, it encouraged users to exhaust their allocation per month, rather than free up unneeded time for others, lest they receive a smaller allocation for the next month. In response, Roger Lazarus phased in a new time-quota system meant to provide a more fair distribution of computer time amongst all users, while still guaranteeing heavy users a predictable production schedule.²⁸⁷

The process began with an experimental weighting assigned to different blocks of time on the CCF supercomputers. After making adjustments based on user feedback, which Lazarus called "gratifyingly critical," the new system created two categories of assigned time. "Category A" was guaranteed computer time, following the previous model from T Division. "Category B" time was not guaranteed, and was subject to availability over weekdays and the weekends. Users from any division could request B-category time, which was allocated based on availability, and prioritized through a new usage calculation that took into account both CPU and PPU time on the CDC 6600s, and the number of jobs having already been run for a given division over that month. Category A time for those divisions that still required a fixed allocation could be negotiated through the C-DO (Division Office), and was based on the code runs a given division expected to make in a specified period of time. The W-4 group, for example, quickly found the new system advantageous, as it was able to negotiate Category A time around a series of large, 2D code runs for a time-sensitive project, in exchange for less guaranteed time after its major project had ended.²⁸⁸

²⁸⁷ Carlson, "Assignment of Computer Time on a Quota Basis"; R. B. Lazarus, C-DO, to Division and Group Leader List, "Assignment of Computer Time," May 22, 1969, Edward A. Voorhees Papers, box 67, folder 1, Machine Usage by Groups and Division 6600, series 48, LANL Archives.

²⁸⁸ Lazarus, "Assignment of Computer Time."

Through negotiations and adjustments, the experiment became a permanent system, which met with general approval, as it provided easier and better computer access based on need, rather than a fixed allocation, while ensuring guaranteed time when required. Between 1972 and 1975, reflecting the increased concern throughout the AEC lab system with budgetary oversight, C Division phased-in a more formal cost-accounting system on top of its computer-time allocation procedure, using project and division charge codes to control access and track usage statistics. The task of creating and enforcing time-apportionment and charging algorithms was ultimately transferred to user-based committees created for that purpose in the 1970s, since leaving C Division with that responsibility placed C-Division personnel in the needlessly difficult and thankless position of arbitrating user allocation demands among all other users.²⁸⁹

As the computer-time management system evolved with increasing demand, similar reforms went into effect for the growing data-storage problem at Los Alamos. As the number of supercomputers, along with the number and size of jobs running on them, grew in the CCF, the number of associated storage devices had grown impressively as well. C Division began to impose restrictions on files in tape storage that had gone unused for extended periods of time, in order to prevent the unnecessary accumulation of disused tapes that users had either forgotten about or had abandoned. Stricter controls on the amounts of storage permitted per user, and the amount of time in which files could go

²⁸⁹ R. B. Lazarus, C-DO, to D. P. MacDougall, ADW, "CCF Usage and Scheduling Allocation," August 11, 1972, A-98-011, Nicholas Metropolis Papers, box 30, folder 2, Computer Education, LANL Archives; W. W. Wood, T-DOT, to CCFUG-APA Committee, "LTS Allocation and Dollar Charging Algorithms for Production Codes: Various Options Compared with CROS Charges," January 6, 1975, Edward A. Voorhees Papers, box 60, folder 3, CCF User's Group, 1975, series 43, LANL Archives; Norman Morse, oral history interview by Nicholas Lewis, Los Alamos, NM (July 30, 2018).

unused in the tape library, could only go so far in addressing fundamental technical limitations of the CCF's data-storage systems and procedures. The enactment of cost-accounting procedures on user storage maintained in the CCF was also phased in during the early 1970s, in conjunction with stricter cost-accounting on computer access. The new procedures improved data-storage control in conjunction with a significant technical reconfiguration of the CCF's storage systems.²⁹⁰

The primary solution to meeting LASL's growing demand for computing capacity in the late 1960s and early 1970s was the purchase of more and faster supercomputers, a strategy that only exacerbated the data-storage problem. The CCF's tape library grew from about 32,000 tapes in 1969 to over 44,000 in 1972, consistently the largest number of tapes in storage of any of the AEC labs. Despite stricter controls on disused tapes remaining in the library, only about half the tapes in the CCF library were considered "active," and most tapes were only about half-full of data. Reel-to-reel tapes were LASL's primary storage medium, and its only archival medium, in the late 1960s. The only other archival option, punched cards, were not a practical solution, due to their low data-density, slow I/O capabilities, and cumbersome physical storage and retrieval process. C Division spent the late 1960s and early 1970s seeking other storage-technology options.²⁹¹

²⁹⁰ Ronald D. Christman, "A Review of the Hydra Project," Los Alamos Scientific Laboratory, Los Alamos, NM (1974); Christman, "The LASL Computer Network."

²⁹¹ William J. Worlton, Lab Notebook (Los Alamos, NM: Los Alamos Scientific Laboratory, 4/11/71), 25, 30; "A Proposal for the Acquisition of Interrelated Large Storage Systems for the Central Computing Facility, March 7, 1969," Edward A. Voorhees Papers, box 41, folder 1, Large Storage Systems, series 28, LANL Archives.

Although the technology was introduced to Los Alamos with Stretch in 1961, hard disk drives were rarely used as data-storage units at LASL through the late 1960s, due to the costs and technical limitations involved. The Control Data 6638 hard disk drives in use on the CDC 6600 supercomputers were expensive and had small capacities that made them less than ideal for uses other than secondary system memory, or sometimes for core dumps, as pioneered with Stretch. The 6638 drives typically hosted the system OS, and user data that was in immediate use on the computer but too large to reside in core memory. The OS and user data usually competed for I/O bandwidth when occupying the same disk drive, compromising I/O performance. Disk performance issues were a moot point for the more data-intensive weapons codes of the late 1960s, which could generate data far exceeding the CDC 6638's one-million-word capacity. For example, a magnetohydrodynamics code recently put into regular use on the 6600s could generate up to twelve-million words of data, requiring a commensurate amount of storage. Core dumps from such storage-intensive jobs, and their resumption at a later time, required the use and manual swapping of many tapes, which not only sacrificed speed, but introduced more chances for operator error. Disk units in the late 1960s were dropping in price, while increasing in capacity and performance, but, as will be discussed below, there were technical roadblocks keeping them from entering regular use at LASL as a layer above tapes as a short-to-medium-term data-storage solution.²⁹²

The primary technical difficulty in finding a solution to the data-storage issue was that each supercomputer of the CCF operated independently of the others. This meant

²⁹² Worlton, Lab Notebook, 4/11/71, 25, 30; "A Proposal for the Acquisition of Interrelated Large Storage Systems for the Central Computing Facility, March 7, 1969."

that each machine required its own set of storage devices. That arrangement was costly in terms of hardware duplication, and wasteful in the production of tapes that were only partially filled, or contained files that were duplicated on other tapes elsewhere, resulting from users moving between different supercomputers. A rapidly growing archive of tapes filled up the CCF. The isolation of each supercomputer from the others also meant that introducing new layers of storage, such as an archival bulk storage system, or a layer of fast, online disk storage, would be cost-prohibitive, if those systems had to be duplicated for each supercomputer. The storage issue only grew worse after the arrival of the CDC 7600 in 1970. Offering about five times the real-world performance of the 6600, the 7600 generated data at an even faster rate. From an access perspective, the 7600's performance also raised, for some users, the possibility of driving interactive graphics from a supercomputer, posing logistical challenges for which the CCF was not yet prepared. The networking technologies that were first put into operation at Los Alamos with MUX appeared to offer solutions to the interconnected storage and access reform problems in the CCF. Although MUX provided C Division with a relatively small-scale proof-of-concept for networking technologies, adapting computer networking to LASL's requirements would be a difficult and drawn-out process, difficulties shared with other early adopters of computer networking.²⁹³

²⁹³ "A Proposal for the Acquisition of Interrelated Large Storage Systems"; D. V. Susco, W-4, to F. F. House, T-1, "Mass Storage," January 15, 1968, Edward A. Voorhees Papers, box 6, folder 5, W-Division 1966-1969, series 9, LANL Archives; Worlton, Lab Notebook, 5/1/71, 11; "Online" storage refers to files on storage media accessible to the requesting computer without operator intervention, while "offline" storage requires an operator to retrieve the needed media, and to place it on the correct interface device before the required files are available to the requesting computer.

The first reference to a networking project at LASL appeared in July of 1968, just as MUX first entered use. C Division's technical-planning groups put forward a series of network and configuration proposals between 1968 and 1970, with plans for the "Hydra" network officially announced in May of 1970. The network's name derived from the mythological Greek serpent with multiple heads, a reference to the network's centralized front-end machine being connected to multiple supercomputers of the CCF. After C Division's dissatisfaction with its modified version of Control Data's SCOPE OS used on MUX, the Hydra software plan included a custom LASL OS for its front-end machine. Developing an operating system locally reflected C Division's growing commitment to in-house software, particularly in the mission-critical production space, despite some maturation and improvement having occurred in vendor-sourced software development since the 1960s. Discussed in detail later in this chapter, the C Division leadership's position on OS development in the early 1970s formed an integral part of its path forward for Los Alamos computing, and became the nucleus of contention for those in other divisions who desired a different direction for LASL supercomputing. Even aside from its own OS, the original Hydra project called for an ambitious supercomputer-storage network, with a suite of advanced features, including a graphical interface for fast Computer-Based Terminals (CBTs). Although not as large in scope or scale as the DoD's ARPA networking project, the Hydra network had to manage many gigabits of online and offline data spread across multiple storage technologies, while keeping that management as seamless to users as possible, and adhering to AEC security requirements.²⁹⁴

²⁹⁴ Christman, "The LASL Computer Network."

To cope with the large sizes and volumes of files that would be traveling across the network, Hydra had to move data quickly, at over four-megabits per second—exceptionally fast speeds for the early 1970s. Transfers also required extreme accuracy, as lost data meant re-attempting the file transfer. No dedicated networking interconnects existed in the early 1970s that could manage the data-transfer speeds and accuracy requirements of a supercomputer. As a consequence, C Division network planners opted to adapt Control Data's channel couplers, which were designed to connect one CDC supercomputer to another, to serve as interconnects between the CCF supercomputers and Hydra's CDC 6600 front-end machine, a task for which they were not intended. The adaptation of existing technologies for new purposes characterized the development of early supercomputer networking at LASL, with multiple factors influencing C Division's improvisation choices with Hydra.²⁹⁵

The dearth of existing technical options was one part of the drive to improvise during Hydra's development. Budgetary constraints also spurred improvisation. A lack of funding to purchase duplicate data-storage hardware pressed C Division to pursue a storage network in the first place. In the same vein, C Division repurposing one of its production CDC 6600s to act as the front-end machine of Hydra's "Star" network ("Star" referred to the configuration of a front-end machine at the center of other client computers and devices) was the result of technical and budgetary constraints. LASL could not purchase a dedicated front-end computer, which would have vastly increased

²⁹⁵ Although the CDC channel couplers could handle much higher data-transfer speeds, file transfers across Hydra were limited to 4Mb/s to avoid saturating the network; "A Proposal for the Acquisition of Interrelated Large Storage Systems"; Christman, "The LASL Computer Network"; Christman, "A Review of the Hydra Project."

the up-front cost of Hydra. It also would have complicated the purchasing requests for Hydra's hardware devices that C Division had to submit to the AEC and Congress. Replacing the 6600 could wait until another fiscal year, when the overall budget requests were not as laden with other hardware purchases.²⁹⁶

To outside observers, losing the CDC 6600's production capacity while the CCF computers remained saturated might have appeared unwise, but the single 6600 represented only 1/12 to 1/15 of LASL's supercomputing power. Hydra's network designers determined that the time saved through better remote access and data retrieval would actually offset the loss of a single 6600. The savings on storage hardware, media, and associated staff costs would also offset much of the cost of replacing the lost capacity in the future, while the network would ensure that LASL's next-generation supercomputer, likely the CDC 8600, would have an efficient storage and access system in-place before its arrival. Another benefit of using a supercomputer as the network's front-end machine was the I/O and processing performance the 6600 offered. Many other "star" network projects in contemporaneous development often used minicomputers as front-end machines, which provided throughput inadequate for LASL's requirements. For instance, Livermore had begun its own "star" storage-network project, called "Octopus," several years before Hydra. Octopus suffered from numerous delays in its development (not uncommon for early networking projects), and its initial use of a DEC PDP-6 minicomputer as a front-end machine severely impacted network performance, demonstrating to C Division the value of a larger front-end machine. Livermore

²⁹⁶ Worlton, Lab Notebook, 11/29/71, 44-46.

eventually replaced the PDP-6 minicomputer with a powerful DEC PDP-10 mainframe, resolving the performance bottleneck. C Division waiting to pursue a large data-storage networking project did not completely prevent delays or missteps with Hydra, but it provided perspective on which design choices could have been most detrimental to the development of a network capable of addressing LASL's most pressing data-storage and remote-access challenges.²⁹⁷

While Hydra itself did not use a minicomputer as its front end, minicomputers became the Swiss-army knives of Lab networking during the 1970s, being programmed and interconnected to fulfill a variety of networking functions at LASL. While this topic will be discussed in Chapter 5, it is worth noting here that the extensive use of minicomputers in LASL networking was the result of several technical and financial influences on Hydra's design. Minicomputers were, compared to supercomputers and mainframes, inexpensive computer systems that were adequately powerful and versatile to be programmed to fulfill networking functions that were not yet being addressed by vendors, such as computer-based terminals, security monitoring, and high-speed file handling. The motivations and methods behind Hydra's design were an example of the continuation and evolution of the characteristic Los Alamos approach to computing, with C Division intentionally developing in-house expertise in the immature but mission-critical new field of supercomputer networking.²⁹⁸

²⁹⁷ Ibid; Lewis interview with Christman; "User's Group Meeting, Oct. 8, 1970," Edward A. Voorhees Papers, box 59, folder 2, CCF Users' Group, 1970, series 43, LANL Archives; Pierre Du Bois, John Fletcher, and Marilyn Richards, "Livermore Time-Sharing System, Part I: Octopus, Chapter 4: Files," October 14, 1970, University of California, Livermore: Technical Information Department, Lawrence Radiation Laboratory, 1-3; Christman, "The LASL Computer Network."

²⁹⁸ Christman, "The LASL Computer Network"; Christman, "A Review of the Hydra Project"; Lewis interview with Christman; Lewis interview with Morrison.

Hydra's basic features and design were formalized in 1970. While the final, working version of Hydra differed from the original plan, Hydra's technical project manager summarized the guiding philosophy behind the network design and the organization of the project charged with creating it, which remained consistent. Of the eight major guiding points for Hydra, the most significant was C Division's stated belief that, "We [LASL] can do it better; in-house software is better; we need interesting work." This statement explicitly stated the rationale behind local R&D projects in decades past, and summarized the intention to use the project as both the source of a technical solution and as a means of fostering local expertise in networking technology. The Hydra project had an initial expected completion date of late 1972, an ambitious development timeframe that, along with other key elements underlying the network's design philosophy, was not a coincidence. The restructuring of data storage and computer access in the CCF with Hydra anticipated LASL's acquisition of the Control Data 7600 supercomputer (multiple copies, in fact) in the early 1970s, and the planned CDC 8600 in the mid-1970s. C Division's leadership intended for the new machines to be the centerpieces behind its reshaping of the CCF's classified operation, and its effort to solidify the Lab's computing R&D strategy. The two-part adaptive strategy was C Division's formal recognition that the intensified computer use and demand that began in the mid-1960 had become the new, permanent state of LASL computing in the 1970s, and that C Division needed to plan accordingly.²⁹⁹

²⁹⁹ Christman, "A Review of the Hydra Project."

In his summary, Hydra's technical manager, Ron Christman, also listed the eight other major guiding points behind the network and its development project, which expressed C Division's overall technical philosophy. This philosophy touched upon each major project and organizational strategy the division's leadership selected as the 1970s began. Despite resistance from groups such as TD-2 (Thermonuclear Weapons Design) and TD-4 (Small Weapons Design), which harbored a vision for LASL's computing future that eliminated nearly all local R&D, and favored the Livermore model of heavy reliance on supercomputer time-sharing, the C-Division philosophy outlined with Hydra proved to be a remarkably resilient influence on C Division's leadership over the 1970s.³⁰⁰

C Division's guiding philosophy in 1970 stated that time-sharing was not a "reasonable" use of supercomputers, and that the local "CROS" operating system (discussed below) for the CDC 7600s would be the basis for all future supercomputer operating systems at LASL. Most C-Division personnel and most CCF computer users were not opposed to time-sharing on smaller computers at Los Alamos, or to the

³⁰⁰ Ibid; Lewis interview with William J. Worlton; Lewis interview with Thompson. Buck Thompson, who, as a programmer in the TD-6 Monte Carlo group, wrote Monte Carlo codes for the weapons designers of TD-2 and TD-4 in the 1970s. According to Thompson, the programmers and software engineers of TD Division tended to sympathize with the difficulties C Division experienced with TD's weapons designers. In Thompson's view, TD weapon-design personnel were often vocally displeased with the software provided to them, and with LASL's computing capabilities, in general. In Thompson's estimation, their dissatisfaction was partly due to a lack of expertise with the limitations of computing technologies and software, and the difficulties of developing and adapting cutting-edge technologies to the advanced computing demands at LASL. The result, according to Thompson, was a growing tension in the 1970s between computing experts, regardless of their division, and non-experts who had less than realistic expectations of computer systems and software, rather than simply a tension between C Division and its users, as has sometimes been portrayed in later summaries of this period in LASL computing history. C Division, because of its central role in LASL computing, likewise became a central target for animosity from some, but far from all, computer users. Computer users who ran the most compute-intensive jobs, for example, tended to approve of C Division's performance-focused planning strategy.

provisioning of contract time-sharing services. In fact, early proposals for Hydra between 1968 and 1969 listed time-sharing OS support as a possible feature, but time-sharing was stripped away in later proposals as the batch-only directive took shape. A large cross section comprised mostly of compute-intensive users, such as those from W-4, and C-Division's upper leadership, primarily Lazarus and Thomas Jordan, remained opposed to time-sharing's introduction on the CCF supercomputers. As a consequence, C Division's operating-system strategy remained focused on providing only batch services in the classified environment, and on employing only in-house operating systems on LASL's highest-end supercomputers.³⁰¹

The focus on locally developed, batch-only operating systems began in early 1969, when C Division, via the AEC, signed a contract for its first CDC 7600 supercomputer and for a second 7600 later that same year. Both machines were to be delivered in 1970. Contract bids from vendors had demonstrated that no other supercomputer system then available matched the performance of the CDC 7600. Even so, Control Data could not guarantee that its SCOPE operating system for the 7600 would be completed in time for delivery of either machine, or that it would be compatible with the array of hardware devices, such as the CDC 7638 disk units, that LASL planned to use with the 7600 supercomputers. LASL's new class of supercomputers potentially would be lacking a ready-made operating system upon their arrival. This was not an unusual situation in the 1960s and early 1970s, as vendors catering to the high-end scientific market tended to depend upon their customers' abilities to develop much of

³⁰¹ Christman, "A Review of the Hydra Project"; Lewis interview with Thompson.

their own software. The AEC labs, one of the largest single consumers of supercomputers in the world into the 1970s, were typically willing to accept the tradeoff of receiving the latest hardware in exchange for minimal included software. This arrangement would begin to change in the 1970s, as demand for supercomputing resources grew among academic, corporate, and other government-funded institutions that were less willing or able to accept a computer system that lacked at least a basic suite of included software, such as FORTRAN compilers and operating systems.³⁰²

For LASL in 1969, vendor-sourced software options were limited, and were typically considered inferior to solutions developed locally. LASL lamented the need for computer vendors to develop their operating systems and other software offerings for as wide a range of customers as possible. As early as 1960, LASL computer experts had noticed and feared the trend of computer vendors emphasizing the economical development of their hardware products over the needs of the AEC labs. The trend had intensified and expanded to software by the late 1960s and early 1970s. Lab investments and collaborations with university and commercial product research efforts provided some assurance to LASL that technologies suited to its computing demands were still being developed outside Los Alamos, but the overall trend was a reduction of the importance of the AEC labs to the supercomputing industry. Computer vendors had less

³⁰² "7600 Contract Points," 1/3/1969, Edward A. Voorhees Papers, box 42, folder 2, CDC-7600 #1, Working Papers-2, LANL Archives; "7600 #2 Negotiations," 9/24/1969, Edward A. Voorhees Papers, box 43, folder 3, 7600 #2, Working Papers-2, series 31, LANL Archives; Marge McCormick, ed., "Computing Requirements of the ERDA Advanced Computing Laboratories, July, 1977," Edward A. Voorhees Papers, box 1, folder 17, Historical Articles, series 2, LANL Archives, 2; MacKenzie, "The Influence of the Los Alamos and Livermore National Laboratories on the Development of Supercomputing," 191.

and less financial inducement to tailor their products specifically to a single customer's demands, further emphasizing the need for local R&D at Los Alamos.³⁰³

The classified nature of much of the AEC weapons labs' production software also played a role with vendor software failing to meet LASL's requirements, as commercial vendors were often excluded from taking part in, or having detailed knowledge of, significant areas of weapons lab software development. The weapons labs usually provided example codes, stripped of their classified components, to vendors who needed to know something of the computational work for which their products were intended, but the examples, when available, were not always adequate or truly representative. Security restrictions often negatively influenced a vendor's ability to meet the software and hardware requirements of LASL and the other labs. Chapter 5 will demonstrate this issue by delving into how security concerns over vendor access to classified software resulted in lasting, negative consequences for Livermore's computing operation, and for its most important supercomputer vendor.³⁰⁴

Adapting Control Data's SCOPE OS for the MUX networking system had left C Division's software developers unimpressed with the suitability of Control Data's operating systems for LASL's computing environment. CDC's delays with the 7600's software package further encouraged C Division's leadership to develop a custom solution, particularly if it meant having an OS written specifically for LASL's workload, and cultivating local experience with cutting-edge, production-OS development. An in-

³⁰³ "User's Group Meeting, July 9, 1970," Edward A. Voorhees Papers, box 59, folder 2, CCF Users' Group, 1970, series 43, LANL Archives; MacKenzie, "The Influence of the Los Alamos and Livermore National Laboratories on the Development of Supercomputing," 191.

³⁰⁴ Lewis interview with Morrison; Lewis interview with Folkner.

house project at the forefront of supercomputer operating-system R&D would also help attract and retain new staff needed for maintaining LASL's increasingly sophisticated software environment. Future division leader John Morrison later stated that division management believed the only practical solution to the OS problem for the 7600 was "to write our own system from scratch, tailoring it to the needs of the majority of the Laboratory users, with unnecessary frills eliminated or minimized." Livermore, which had taken delivery of CDC 7600 serial number 1 in 1969, reacted similarly to Control Data's software situation for the 7600, and opted to develop its own OS. Unlike LASL, Livermore, after experimenting with time-sharing operating systems on the 6600 since the mid-1960s, committed to a time-sharing system for the 7600.³⁰⁵

Named after a ridge near Los Alamos, the Chili Ridge Operating System, better known as CROS, entered development at LASL in 1969, and began early testing in mid 1970, after the two CDC 7600s had arrived that summer. Control Data had managed to implement a somewhat functional version of SCOPE for the 7600s at the time of their arrival in the CCF, but C Division began transitioning to CROS as quickly as possible, in order to take advantage of the 6600 codes that had been adapted for CROS on the new hardware. While the first 7600, dubbed Machine #13 in the CCF, had significant reliability issues in the first months after delivery, Machine #14 passed its performance evaluation and entered its testing phase on three major codes in a matter of weeks. Optimizing the codes on the 7600 with an early version of CROS, and developing a

³⁰⁵ "7600 Computer Now in Operation at Livermore Lab," *Magnet* (April, 1969), 3; John A. Brockmeyer, and A. Gene Dornhoff, "Computer Interfacing at LASL," Los Alamos Scientific Laboratory, LA-UR-76-2502 (1976), 1-2; Lewis interview with Morrison; McCormick, ed., "Computing Requirements of the ERDA Advanced Computing Laboratories, July, 1977," 3.

software solution to allow the 7600s to read 6600-formatted tapes, took longer. Some users had hoped that CROS would be adapted for the 6600s, but Charlie Folkner, leader of the C-2 CCF Systems Support group, stated that there was not sufficient staff time for such a conversion. The 6600s would continue to run CDC-sourced operating systems for the remainder of their service at LASL.³⁰⁶

CROS gradually entered regular production use starting at the end of 1970, with numerous updates and bug fixes over the next year improving the performance and reliability of the codes running on the OS. The reliability of new CDC tape and disk units became a "chronic" issue for both the 7600s and the 6600s at LASL in the early 1970s. C Division had begun to phase-in a collection of new storage devices concurrently with the introduction of the 7600s. In addition to CDC tape drives, the new storage devices also included Control Data hard disk units, and the IBM 1360 "Photostore," a massive archival system that stored data on photographic film chips. The Photostore, one of only six IBM constructed, provided about one-terabit of storage in total. The improving price/performance/capacity calculus of hard drives in the early 1970s finally favored the use of disk units as supercomputer short-term storage, rather than only as extended memory. The new storage systems would not be networked until the Hydra storage network neared completion. C-2 produced software interface fixes that reduced the tape error rates, and resolved compatibility issues when exporting graphical data from CROS to tape. About 7,000 tapes were mounted on each 7600 per month, and

³⁰⁶ Charlene Douglass, Tim Locke, and Mary Maestas, eds., "Chili Ridge Operating System Operator's Information Manual," Los Alamos Scientific Laboratory, Los Alamos, NM (undated), 37; "User's Group Meeting, July 9, 1970"; "User's Group Meeting, Sept. 3, 1970," Edward A. Voorhees Papers, box 59, folder 2, CCF Users' Group, 1970, series 43, LANL Archives.

cyclic problems with reading and writing accuracy on individual tape units persisted through 1972. Unrecoverable disk errors also occurred regularly on the new CDC 7638 disk units, partly due to flaws in some of the disk platters, and the dearth of available software monitoring tools to diagnose trouble spots. Updated disk drivers written for CROS in 1972 provided a completely revised error reporting system, and helped to resolve some of the most significant disk issues.³⁰⁷

IV. Reorganization

By 1973, the two 7600s at Los Alamos were working well under CROS, with both computers reporting zero idle time in February. Updates to the OS had improved system stability, with the number of software crashes on each machine dropping from about sixteen per month in February to about three in April. The updates also resolved most of the storage-device issues of the previous two years, although the number of tape and disk errors still occasionally fluctuated. The storage issues with the 7600s were largely device-side, and Control Data continued to work on the problem. CROS had proven to be highly efficient for LASL's code base, which was no surprise, because the

³⁰⁷ "User's Group Meeting, Dec. 14, 1972," Edward A. Voorhees Papers, box 59, folder 5, CCF Users' Group, Aug. 72 - Dec. 72, series 43, LANL Archives. While the Control Data 7600s were entering full production use at LASL, the venerable IBM 7030 Stretch underwent decommissioning after ten years of service. Stretch received a retirement party (which insiders in C Division called a "wake") on June 21, 1971, with veterans of the original Stretch Project, such as LASL's Roger Lazarus, Robert Frank, and Bengt Carlson, and IBM's Werner Buchholz (author of the book describing Stretch's design) and Steve Dunwell, in attendance. Stretch was the last IBM supercomputer at Los Alamos until the Roadrunner cluster computer was completed in 2008. The decommissioning of Stretch, while bittersweet for its designers and ardent users, was necessary for freeing up space on the computer floor of the CCF for the new supercomputers LASL planned to install within the next several years. Maintaining and operating the aging Stretch had also become increasingly difficult and expensive, making it hard for C Division to justify the very limited production capacity it offered by the standards of the early 1970s. For details on Stretch's decommissioning, see: "Stretch, the Super Computer, is Gone," *The Atom* 8, no. 6 (July-August, 1971): 23-26.

OS was developed specifically for the local code requirements. The batch OS quickly earned a loyal following among many CCF computer users, particularly those who depended upon CROS' ability to provide fast turnarounds on jobs that would have taken far longer on the 6600s, if the older computers could have handled them at all. In early 1973, Machines #13 and #14 were together running over 21,000 jobs per month. A point system tracked and charged-back supercomputer usage by division, and revealed the changing demographics of computer usage at Los Alamos in the early 1970s.³⁰⁸

T Division, which had been the largest single user of LASL's computing resources since the Manhattan Project, had slipped far below TD (Theoretical Design) Division, the new largest user, and was in an approximate tie with J Division as the second largest user. A half-dozen other divisions were also making far heavier use of the CCF than only a few years before. The change for T Division was a consequence of four of its compute-heavy, weapons-related groups, and the W-4 Small Weapons Theoretical Design group, being transferred to TD after its formation in 1971. Robert Thorn, formerly the alternate division leader of T Division, became the first leader of TD. Weapons Nuclear Engineering (W) Division was dissolved altogether in 1972, and its groups were consolidated with parts of GMX to form the Design Engineering (WX) Division.³⁰⁹

³⁰⁸ "User's Group Meeting, February 15, 1973," Edward A. Voorhees Papers, box 59, folder 6, CCF Users' Group, Jan. 73 - June 73, series 43, LANL Archives; "User's Group Meeting, April 12, 1973," Edward A. Voorhees Papers, box 59, folder 6, CCF Users' Group, Jan. 73 - June 73, series 43, LANL Archives; MacKenzie, "The Influence of the Los Alamos and Livermore National Laboratories on the Development of Supercomputing," 186.

³⁰⁹ "User's Group Meeting, April 12, 1973"; "Los Alamos Scientific Laboratory Organization, 1971," Edward A. Voorhees Papers, box 1, folder 2, Organization Charts, series 1, LANL Archives; "Los Alamos Scientific Laboratory Organization, 1972," Edward A. Voorhees Papers, box 1, folder 2, Organization Charts, series 1, LANL Archives.

The reorganization of numerous divisions at LASL followed the 1970 retirement of Norris Bradbury, who had largely resisted making significant changes to the operational structure of Los Alamos after he accepted what he had believed to be his "temporary" role as director in 1945. His successor as director, Harold Agnew, physicist, fellow Manhattan Project veteran, and former leader of W Division, made a variety of administrative and organizational changes at Los Alamos. Agnew created and eliminated several divisions, while consolidating groups based on their research focus. He also placed new and existing divisions under the purview of new associate director positions. Associate directors oversaw multiple divisions grouped by mission focus, and report to the director. T Division was placed under the new Associate Director for Research, along with the Physics (P) and Meson Physics (MP) Divisions. This separated T Division from the other major computer users, TD, WX, and J Divisions, which were placed under the Associate Director for Weapons. The result was that T Division, after having its weapons-related groups transferred to TD Division, and being assigned to a new associate directorate, was no longer part of the weapons program, severely curtailing its financial and administrative influence over the computing operation.³¹⁰

Agnew's reorganization was, in large part, a response to a growing level of oversight from the AEC's new Division of Plans and Reports, which focused on increasing the accountability, while reducing the flexibility, of its laboratories. The AEC was, itself, experiencing heightened scrutiny from the US Congress. Starting in the late

³¹⁰ "Los Alamos Scientific Laboratory Organization, 1971"; "Los Alamos Scientific Laboratory Organization, 1973"; "Los Alamos Scientific Laboratory Organization, 1973," Edward A. Voorhees Papers, box 1, folder 2, Organization Charts, series 1, LANL Archives.

1960s, the commission increasingly emphasized evaluating each of its labs based on their effectiveness in adhering to, and fulfilling, their primary missions. An important consequence of these changes was that the AEC labs had less freedom to reinvent themselves and diversify their research directions in response to the nation's prevailing political winds, as they had done repeatedly in the past. The labs were more heavily restricted to research bearing on their primary missions, leaving them more vulnerable to outside interference and program cuts by AEC management. New AEC field offices began to conduct closer surveillance of the labs, and to review the products of lab research with greater scrutiny than ever before.³¹¹

The Agnew reorganization gave Los Alamos divisions the appearance of greater cohesion in their directions of research, with less research overlap between divisions. The organizational changes also, inadvertently, altered the power dynamics among the users of LASL's computing resources. As the weapons program faced greater scrutiny from the AEC, the pressure to adhere to deadlines, to fulfill project goals, and to justify the existence of research programs, added social tension between supercomputer users who had differing computing needs. Users who wanted to maintain the batch-processing emphasis in the CCF were increasingly at odds with users who desired the ability to interact with jobs in progress, even at the expense of the computer performance that batch users wished to preserve. Users with different computing needs had typically been,

³¹¹ Westwick, *The National Labs*, 272-273.

before the reorganization, spread out in groups among multiple divisions, meaning that groups with differing computing needs had often occupied the same division.³¹²

After the reorganization, groups with similar research programs and computing demands tended to be concentrated together in the same divisions, allowing, to a greater degree, users with the same computing demands to dominate the computing agenda of a single division. As a consequence, between 1971 and 1973, C Division served computer-using divisions that were, internally, increasingly unified in their demands upon the CCF, and in their objections to computing strategies that they did not perceive as catering to their interests. With 80% of the C Division's funding came from LASL's weapons program, the reorganization fundamentally altered which divisions controlled the largest shares of that funding, and, consequently, the power to influence C Division's operation.³¹³

V. Intervening in C Division

The heightened managerial demands of the AEC coincided with the growing difficulty of improving upon existing nuclear weapons designs in the 1970s, as greater efficiencies in yield-to-weight ratio and other parameters became to harder extract over time. The twin demands of added pressure from oversight and the diminishing returns on weapons design resulted in about half the number full-scale nuclear tests per year in the 1970s than in the 1960s. Computing was a critical component for demonstrating

³¹² Westwick, *The National Labs*, 272-273; "User's Group Meeting, May 10, 1973," Edward A. Voorhees Papers, box 59, folder 6, CCF Users' Group, Jan. 73 - June 73, series 43, LANL Archives.

³¹³ "User's Group Meeting, April 12, 1973"; "User's Group Meeting, May 10, 1973."

accountability, while reducing the risks and costs of the weapons program. The added importance of computing for the newly reorganized weapons program, key parts of which were increasingly pro-time-sharing, resulted in the first direct, high-level intervention in C Division's computer planning and operating strategy of the 1970s. But before elaborating on the intervention itself, it is necessary to understand how LASL's unique organizational culture and mission requirements within the AEC lab system contributed to the decisions behind C-Division strategy. Understanding the distinct cultural and mission-related influences on the Lab's computing operation clarifies why time-sharing became a controversial issue at Los Alamos, but not at Livermore, despite both weapons labs experiencing similar pressures and driving forces in the early 1970s, and highlights how the controversy at LASL could escalate to the point of administrative intervention.³¹⁴

While both LASL and Livermore were early adopters of the CDC 7600 supercomputer, the differing approaches that each lab employed to contend with the same operating-system issues underscores the fundamental differences between either lab's culture, and the nature of their respective missions. When Livermore took delivery of its first CDC 7600 a year before Los Alamos, it likewise faced the problem of the supercomputer lacking an operating system, while knowing that the Control Data OS would likely not suit that lab's requirements when it finally shipped. Much like LASL, Livermore's Computational Department began development of its own operating system for the 7600. Employing time-sharing on its supercomputers was not the contested issue at Livermore that it was at LASL. The difference was due, in part, to the jobs being run

³¹⁴ Lewis interview with William J. Worlton; Lewis interview with Morrison; MacKenzie, "The Influence of the Los Alamos and Livermore National Laboratories on the Development of Supercomputing," 186.

through Livermore's computer operation typically being better suited to time-sharing than at LASL, which depended more heavily on long, batch-oriented jobs. Livermore had purchased a CDC 1604 mainframe in the mid-1960s to serve as a platform to experiment with the creation of a time-sharing OS. Written in FORTRAN, the Livermore Time-Sharing System (LTSS) was relatively portable, and transitioned through several Control Data computer models at Livermore in the 1960s, including the 6600 supercomputer. LASL's computing advisors, such as Nicholas Metropolis, remained unimpressed with the reliability, security, and primarily the performance tradeoffs of LTSS on the CDC 6600, particularly while Los Alamos had more demand for computing time than it could supply. Livermore, with its larger budget, could afford to purchase additional computers to offset the performance lost to time-sharing, but Los Alamos could not.³¹⁵

Livermore and Los Alamos' distinct operating philosophies also played a role in the reticence of LASL's C Division toward the time-sharing model. Even as weapons design at both labs became increasingly oriented toward making incremental improvements over the 1960s and 1970s, Livermore remained committed to proposing and testing novel features and weapons-design concepts to a far greater degree than the more conservative Los Alamos. This commitment was called the "Livermore enthusiasm" at the AEC level. Los Alamos pursued design options that were more theoretically sound, and less likely to become technical dead ends. These distinct approaches were part of the original plan behind creating a second AEC weapons lab,

³¹⁵ Du Bois, et al., "Livermore Time-Sharing System," 3; Fred W. Dorr, "Computer Science and Services Division of the Los Alamos Scientific Laboratory," Los Alamos Scientific Laboratory, LA-UR-74-1566 (1974), 3, 5; Ghausi interview with Fernbach; Westwick, *The National Labs*, 318-319.

with one, Los Alamos, pursuing more theoretically grounded research, to leave the other free to conduct riskier research efforts. As a consequence, Livermore in the 1970s utilized more computing power in its design efforts than LASL, with a greater reliance on live testing to evaluate a design's feasibility. LASL's weapons program utilized more of its computing time for demonstrating the theoretical bases behind its designs, resulting in fewer uncertainties going into the testing phase. LASL's more cautious approach required less full-scale live testing for design refinement and verification, and ensured a steady supply of reliable weapon systems entering the US arsenal. Livermore's "enthusiasm" provided possible future rewards and new paths of research that might have been too risky to pursue had the US maintained only a single design approach. These design goals and methodologies translated into differing computer codes and operational strategies as the two labs evolved.³¹⁶

The relative percentages of the weapons programs at LASL and Livermore in relation to each lab's overall research effort also influenced the computing priorities of the two. Livermore's weapons program occupied about 80% of that lab's budget and mission focus, compared with about 65% at Los Alamos, which remained more diversified in its research well into the 1970s. Los Alamos had to cater to a wider variety of research applications in its computing operation, including more unclassified and non-weapons computer use, than Livermore. As a result, LASL's C Division had more

³¹⁶ Westwick, *The National Labs*, 198-201; MacKenzie, "The Influence of the Los Alamos and Livermore National Laboratories on the Development of Supercomputing," 185-187.

diverse application and security requirements than at the more weapons-focused Livermore.³¹⁷

The speed and extent with which C Division at LASL embraced new computing technologies reflected the Lab's overarching concern to avoid jeopardizing its primary mission, which included its need to support a variety of non-weapons computer use, while still maintaining capabilities in the CCF that were well-suited for weapons development. For example, C Division was willing to incorporate networked storage in the CCF, but not time-sharing, which were usually linked together, as at Livermore. This predilection of C Division was represented in the continuing drive to establish local expertise within a technical field before committing to its use in programmatic computing, and to use that expertise to tailor (or even reject) technologies for LASL's uses. Described as overly cautious by some former users and managers of LASL's CCF, the resistance of C Division's leadership to integrate time-sharing into the supercomputer facility in the early 1970s was consistent with the long-standing philosophy guiding most any mission-related science division at Los Alamos. The problem for C Division's leadership in 1973 was that many of the newly dominant computer users in the weapons program did not view computing as a *peer* in the sciences, but as merely a *service* to the sciences, and were willing to intervene in C Division's managerial structure to obtain the computing tools they desired.

Many of the long-standing computer users at Los Alamos who had been part of LASL's direct role in the development of scientific-computing technologies and methods,

³¹⁷ Westwick, *The National Labs*, 198-201.

conceived of computing as both a tool in their daily work, and as a field of scientific research in its own right. They had experienced directly the interplay between the advancement of LASL's computing capabilities through the development of local, or locally modified solutions, and the ability to meet the escalating computational requirements in nuclear physics and engineering. C Division's computer-operation planning not only aligned with this view, but formed as a consequence of it, with its adherents tending to share C Division management's caution. That is not to say those users and C Division management always agreed upon the extent or direction in which the computing-operation planning contended with the application of new technical options. Disagreements occurred, and often found expression in user group meetings and formal and informal correspondence between users and C Division. Despite the disagreements, the appropriateness of C Division's operating as a mixed service and science division was not in question. In contrast, computer users who had less first-hand experience with the concurrent, intertwined development of the Los Alamos weapons program and the Lab's local computing expertise, particularly those who engaged only reluctantly with computing technologies in order to meet the added accountability and design-refinement pressures, tended to be less charitable in their impressions of C Division's caution and relative independence in its planning strategies. In this view, C Division operating as both a service and as a science division appeared to be more of a hindrance toward obtaining new tools, than a means of preventing disruptions to LASL's computing operation, or in designing its future.³¹⁸

³¹⁸ Lewis interview with Morrison.

This dichotomy of users and views on computing at Los Alamos was not new in the 1970s, but the reorganization of LASL's weapons program meant that newer, non-expert users gained significantly more financial and administrative influence between 1971 and 1973 with the formation of TD and WS Divisions. Neither division was uniform in its views on computing, but their dominant groups and members of management were more concentrated against C Division's existing approach than had been the case in their predecessor divisions. After becoming leader of the new TD Division, Robert Thorn, an ardent supporter of time-sharing at LASL, had a far greater ability to bring administrative and financial pressure onto C Division management in support of TD Division's computer users. A missed opportunity at a technical compromise finally brought this pressure to bear.³¹⁹

Just as CROS matured into a reliable workhorse OS at LASL in the early 1970s, the Livermore Time-Sharing System (LTSS) on the CDC 7600 matured at Livermore. Users in differing camps debated in CCF user group meetings the merits of each OS. Members of C Division management and CROS-supporting users who were not flatly opposed to time-sharing in the CCF were still reticent to commit institutional support for an additional operating system, or to undertake the estimated "100-man-years" needed to convert the LASL codes to LTSS, particularly after they had just been converted or rewritten for CROS. Either action, independent of any other cost of introducing secure time-sharing at LASL, was not in the existing budget. Petitions to bring LTSS into use on at least one 7600 at Los Alamos were denied by Roger Lazarus, citing the efficiency

³¹⁹ Lewis interview with Morrison; Thomas R. Hill, oral history interview by Nicholas Lewis, Los Alamos, NM (April 25, 2019).

of CROS, and echoing the detriments of LTSS that other C-Division managers had already raised in user group meetings. Rumors suggest that a small group of users clandestinely installed LTSS onto a 7600 in the CCF during a late-night shift as a demonstration, but the consequences of their actions remain unknown.³²⁰

The pent-up demand for time-sharing in the CCF, and the dissatisfaction with C Division's resistance, reached their peak in the early summer of 1973, when the AEC denied a LANL request for funds to purchase a "medium-sized" computer for the CCF. Among other uses, the medium-scale computer would have provided an opportunity to introduce local, classified time-sharing to the CCF, without disrupting service on the supercomputers. Similar to MUX and remote-access, the new system could have served as a stepping stone to persuade C Division's leadership to reconsider its stance on time-sharing on either the 7600s, or on the next-generation supercomputer system LASL was in the process of procuring.³²¹

After the AEC found inadequate justification for the medium-sized computer and refused to fund its purchase, users demanding time-sharing in the CCF, with the support of TD leadership, petitioned Harold Agnew to intervene in the impasse with Roger Lazarus. At the behest of users representing the majority of C Division's operating funds, Agnew removed Lazarus and one of his alternate division leaders (an "alternate" was effectively the division leader's deputy), Thomas Jordan, from their positions. Jordan

³²⁰ "User's Group Meeting, Sept. 14, 1972," Edward A. Voorhees Papers, box 59, folder 5, CCF Users' Group, Aug. 72 - Dec. 72, series 43, LANL Archives; "User's Group Meeting, Dec. 14, 1972"; Lewis interview with Thompson. The computing operation at Los Alamos often employed "man years" as an estimate of how much staff time was required for a given task or project, depicted in the time needed for a single individual to finish that project, if unaided.

³²¹ "User's Group Meeting, April 12, 1973"; "User's Group Meeting, May 10, 1973," Edward A. Voorhees Papers, box 59, folder 6, CCF Users' Group, Jan. 1973 - June 1973, series 43, LANL Archives.

stepped down from C-Division leadership altogether. Lazarus temporarily became an alternate leader of C Division while his replacement, Frank McGirt, transitioned from his previous role as alternate leader of the C-4 Scientific Applications Services group to be the new division leader. McGirt had relatively little managerial experience, but Agnew appointed him to the position because he had been supportive of the CCF users who had demanded time-sharing capability. The appointment as C-Division leader was for one year, after which Agnew would determine if McGirt would be replaced, or allowed to remain as division leader.³²²

After his ouster, Lazarus commented that scientific computing had gone in directions that he could not follow, and he returned to physics after ending his run as alternate division leader. Lazarus later commented that much of his opposition to time-sharing was it being "people efficient, not machine efficient." This stance, having served LASL and its computing operation well during periods of scarcity in computing power, became a rallying point for those who opposed C Division's operating strategies and independence once the economic and political dynamics at Los Alamos changed in the early 1970s. Frank McGirt quickly announced that LASL would begin work with Livermore to import LTSS for use on one of the CDC 7600s in the CCF. The transition

³²² "User's Group Meeting, May 10, 1973"; Harry McGavran, oral history interview by Nicholas Lewis, White Rock, New Mexico (July 17, 2014); Lewis interview with William J. Worlton; Lewis interview with Morrison; J. E. Zelenka to B. T. Jones, "LANL Visit to Headquarters - May 30, 1974," CBI 80, Control Data Corporation Records, Engineering, Research and Development, 1957-1986, box 36, folder 4, Atomic Energy Commission (AEC) - Los Alamos Scientific Laboratory (LASL), 1970-1974, CBI; "Computer Science and Service Division, 1974," Edward A. Voorhees Papers, box 1, folder 1, Organization Charts - C-Division, series 1, LANL Archives; "C-4 Group Leader," Edward A. Voorhees Papers, box 1, folder 1, Organization Charts - C-Division, series 1, LANL Archives.

could not be instantaneous, as the LASL codes and software libraries to be used on LTSS would have to be converted.³²³

The CCF also had to prepare for the introduction of a new OS through the training of users and operators, and the integration of LTSS support into the nearly complete Hydra storage network. The process of modifying LTSS and existing LASL software and services for use with one another continued into 1975, during which time C Division also incorporated Control Data Corporation's new KRONOS operating system onto the remaining production CDC 6600 in the CCF, and the two CDC Cyber 73 computers that LASL had purchased for smaller-scale and legacy open and classified computing jobs. The Cyber 73 computers were based on the 6600 architecture, and so offered compatibility with 6600 software, but in a more cost-effective package. KRONOS featured both batch and time-sharing capability, fulfilling C Division's new requirement from upper Lab management to expand its time-sharing offerings, but the vendor-sourced OS required modification from C Division to suit LASL's computing environment. The Hydra CDC 6600 front-end machine retained its custom "Hydra" OS.³²⁴

VI. Chapter Conclusion

Roger Lazarus' protectionism toward computer performance at the cost of user time and effort was the norm at Los Alamos and elsewhere when general-purpose

³²³ Fitzpatrick interview with Lazarus; "CCF CDC 7600 LTSS Newsletter, May 19, 1975," A-93-015 Division Office Papers (1975-1977), box 1, folder 1, March-June, 1975, Cray-1, LTSS, LANL Archives; "Second 7600 Running LTSS, Sept. 15, 1975," Edward A. Voorhees Papers, box 58, folder 1, C-Division Memos, 1974-1975, series 42, LANL Archives.

³²⁴ Dorr, "Computer Science and Services Division of the Los Alamos Scientific Laboratory," 1-3.

computers were far more rare than they would later become. The unpredicted escalation of computing demand at Los Alamos in the 1960s, which left no computer cycles to spare, reinforced the "machine efficient" strategy within C Division's early leadership. LASL's characteristic, conservative research philosophy seemed to support this strategy, as it guarded against the more "radical" approaches undertaken at Livermore. The circumstances that reinforced C Division's commitment to the "machine-efficient" concept, and its commensurate opposition to time-sharing on the Lab's supercomputers, also fostered a reluctance to reconsider C Division's OS strategy in light of new organizational pressures, and amid the rapidly changing place of computing in the AEC weapons complex. As a consequence, C Division's OS strategy, and its already contested existence as a mixed service and science division, came under top-level scrutiny at LASL for the first time, as opponents of C Division's approach to computing were able to utilize the OS controversy as a driving point to push the Lab Director for an intervention in C-Division policy.

The intervention in C Division's management and operations planning in 1973 was the culmination of a process that began in the mid-1960s, when the technical and managerial pressures on LASL's weapons program intensified, following a renewed emphasis on weapons as the Lab's primary source of funding. Computing, as a consequence of nuclear testing restrictions and the increasing difficulties of extracting greater efficiencies from weapon designs, assumed a more central role in LASL's weapons program. Heightened administrative demands from the AEC only exacerbated the importance of computer use in maintaining and justifying LASL's mission. These

changes coincided with, and helped to foster, a shift in the composition of the CCF's user base toward users who had fewer ties with the traditional Los Alamos approach to computing. Many of these newer users did not experience the concurrent development of weapons design and an in-house computing R&D effort at Los Alamos. Instead, access to supercomputing was more of a given, and an occupational necessity, for scientists and engineers in the weapons program than in decades past. These users were typically under pressure from AEC and LASL management to meet accountability requirements and deadlines, offering little incentive to consider C Division's operational goals, particularly when C Division's upper leadership appeared intractable in its views on LASL's computing requirements.

In the early 1970s, changes in the user community began to tilt the economic and administrative balance within the LASL weapons program began away from the previous, traditionally dominant users of Los Alamos computing. Despite its support base shifting, C Division's management did not significantly revisit its long-standing operating strategy to contend with the AEC pressures being placed upon weapons-program users, or the pressures those users were, in turn, placing on C Division. C Division not recognizing the importance of rapidly changing organizational pressures at LASL, and not taking sufficient conciliatory action to ameliorate user demands, provided an immediate, identifiable point of frustration for the newly dominant computer-using weapons divisions. As a consequence, time-sharing served as an example of C Division being out of step with its user base, and became a rallying point for those who did not perceive computing to be an equal among the sciences, and wished to curtail C Division's

independence. The intervention in C-Division management brought time-sharing to LASL supercomputing, but the overall structure and operational philosophy of the division remained intact. C Division's attempt to reassert its operational and planning independence after the introduction of a new class of supercomputer in the mid-1970s would result in a complete reshaping of LASL computing.

-Chapter 5-

The End of the Beginning:

The Cray-1, DEMOS, and the Reorganization of C Division

I. Chapter Introduction

The change in C-Division's leadership in 1973 altered the division's long-term planning, and opened the door to time-sharing on the supercomputers of the Central Computing Facility. Converting LASL codes and undertaking the variety of other steps needed to utilize LTSS on the CDC 7600s of the CCF strained C Division's software-development and support and complicated the rollout of new supercomputer systems at LASL (as this chapter will discuss). C Division's traditional research-intensive approach to computing continued following Lazarus' departure, and even expanded in the late 1970s. Local R&D in C Division enhanced LASL's computing capabilities between the mid and late 1970s, especially with key advancements in supercomputer networking and simulation visualization. While C Division's custom networking made Los Alamos a leader in that field, the division's attempt to maintain its leadership in the field of supercomputer operating-system development in the late 1970s intensified upper-level scrutiny of, and animosity toward, C Division's approach to computing.

Discord within segments of the weapons program over C Division's status and operation as a mixed science and service division did not vanish in the years after Roger Lazarus stepped down. Rather, the introduction of time-sharing and supercomputer networking in the CCF had eased some of the tension over the division's operating strategies, as a variety of beneficial new services and capabilities quickly became

available to LASL computer users. C Division, following from its conscientious approach to new technologies, successfully avoided a disastrous supercomputer acquisition similar to Livermore's acquisition of the STAR-100 (to be discussed in this chapter), in a move that also eased some of the tension over its usual caution. After several years of somewhat improved, or at least less-poor user relations, C Division's attempt to establish LASL at the forefront of advanced supercomputer operating-system design with another batch-only OS reinvigorated and intensified animosity with the users who had most strongly championed for time-sharing in the CCF, and most fervently believed that C Division planning and R&D should happen under more strictly user-defined parameters. While the reorganizations of LASL's weapons program in the early 1970s altered the Lab's power dynamics in favor of users demanding both time-sharing and an intervention in C Division's planning strategy, a weapons-program reorganization in the late 1970s brought some of the strongest opponents of the C Division approach to computing into the Lab's highest offices. The resulting altercation, fueled by a difficulty within C Division to adapt to prevailing political dynamics, effectively reinvented C Division, eliminating the "sciences" in the Computing Sciences and Services Division, and ending the traditional Los Alamos approach to computing that had evolved decades before.

II. Control Data and the Case of the Wrong Computer

After relinquishing his position as division leader in favor of Frank McGirt, Roger Lazarus remained for sixteen months as an alternate C-Division leader until November of

1974. During that time, contract negotiations between C Division and the Control Data Corporation uncovered an impending crisis over the machine's architecture. Lazarus had been closely involved with LASL supercomputer acquisitions since the collaboration on the IBM 7030 Stretch, and, as division leader, had been the point figure during the purchase of each new supercomputer. As mentioned in previous chapters, around the same time C Division acquired a new supercomputer, it usually began searching for its *next* supercomputer system, owing to the long lead times between the proposal of a new design, the evaluation time needed to assess rival computers' merits, and the ultimate delivery of a working system. Since the mid-1960s, the staff in charge of computer acquisition in the Coordinator for Automatic Data Processing (CADP) office and C-DO (Division Office) preferred long lead times, and meticulous evaluations. Ensuring that a computer lived up to the promises of its designers, the preemptive search for future supercomputers also provided Los Alamos with time to prepare its computing center, code base, and users for the necessary changes.³²⁵

Since the founding of C Division, division management and representatives of the CADP office had maintained close ties with computer and other technology vendors to remain apprised of potential new systems or other developments relevant to LASL's computing needs. Maintaining close contact with vendors, and remaining up-to-date on their technical offerings, were important functions of C Division. The division took those functions quite seriously following the difficulties LASL encountered in the mid-1960s with the rapid sea change in the high-performance computing industry, which caught the

³²⁵ "AEC 1801," Edward A. Voorhees Papers, box 50, folder 1, CADP Organization, series 38, LANL Archives.

Lab off guard with the elimination of its preferred upgrade path from the IBM 7030 Stretch. Despite C Division's conscientiousness, the supercomputer market of the early 1970s was even more difficult to navigate than that of the mid 1960s.³²⁶

The Control Data Corporation 6600 supercomputer first introduced Los Alamos to Seymour Cray, the computer's principal designer. The CDC 6600 and 7600 became the predominant supercomputers at Los Alamos and Livermore due to the dearth of competitors offering commercial machines with comparable performance, and, primarily, the large existing investments of the labs in software and compatible hardware for the CDC ecosystem. Between the mid-1960s and the early 1970s, LASL and Livermore together purchased nineteen Control Data supercomputers (along with massive purchases in storage, expanded memory, and other hardware devices). The result was a close working relationship between Control Data and the two weapons labs. For LASL's computing personnel, Seymour Cray was Control Data's most important representative.³²⁷

Seymour Cray often appeared in person to supervise the delivery of a supercomputer and its installation, particularly with the delivery of early serial numbers of each of his latest supercomputer designs. According to LASL interviews, a case of Leinenkugel's brand beer from Cray's hometown (and the site of his design studio) of Chippewa Falls, Wisconsin, also accompanied each supercomputer. Cray's formidable technical knowledge and disarming personality, and the fact that he was an engineer rather than a "manager" type, earned the supercomputer designer a great deal of respect

³²⁶ Dorr, "Computer Science and Services Division," 2; Fred W. Dorr, "Computer Science and Services Division Activities and Plans, April, 1978," Los Alamos Scientific Laboratory, LA-7093 (1978), 17.

³²⁷ Dorr, "Computer Science and Services Division," 2; Dorr, "Computer Science and Services Division Activities and Plans, April, 1978," 17.

and admiration at Los Alamos. When LASL began searching for its follow-on machine to the CDC 7600, Cray's 8600 was an early favorite within C Division. This was partly based on the reputations of Cray's previous designs, the 1604, 6600, and 7600. The 8600, as the design proposal existed in April 1971, also appeared well fitted for LASL's needs. Based on the CDC 7600 platform, the 8600 would likely have been an easier software transition from the Control Data computers already in use at LASL (the reasons why none of the supercomputers under consideration were going to be a truly "easy" transition will be explained below). Cray's new computer also offered approximately four times the performance on the most common job types at Los Alamos compared with its putative rivals. C-Division personnel conducted four site visits to Control Data in Minneapolis in 1971 alone, each time receiving regular updates on Cray's 8600, and on another CDC supercomputer project then in progress. The LASL group also met that year with Texas Instruments, IBM, and Burroughs. The TI Advanced Scientific Computer (ASC), at least on paper, offered the strongest competition for Control Data.³²⁸

The CDC 8600 architectural concept was effectively four 7600s operating in unison and sharing a single memory. Cray's 8600 team opted for a multiprocessor design because component performance had not increased significantly since the 7600, meaning that another single-processor system, like the 6600 or 7600, would not yield the desired processing power. Improvements in computing hardware performance were (and still are) typically achieved in two ways: improving component technology (such as lower

³²⁸ John Morrison, oral history interview by Nicholas Lewis, Los Alamos, NM (July 1, 2014); Charles Folkner, oral history interview by Nicholas Lewis, Los Alamos, NM (August 13, 2015); Jack Worlton, "The Cray-1: The Machine and the Man, September 21, 2006," cabinet 9-2 (Jack Worlton), HPC Archive; William J. Worlton, Lab Notebook, "4/21/71 to 2/1/72," HPC Archive.

logic-gate delays and faster memories); and changing computer architectures (the internal organization and operational characteristics of the computer). Increasing computing performance through faster components was the easier route from a design and programming standpoint, because moving to new components made the same computer faster, while leaving its architecture and the software that ran on it largely unchanged. A LASL physicist later said, from a user standpoint, he preferred component upgrades on the same architecture, because all he had to do for his codes to work faster was "nothing."³²⁹

The component problem was a significant technical "reverse salient" for the scientific-computing industry as a whole across the late 1960s and early 1970s. Most general-purpose computers since the 1940s had featured serial architectures, meaning that they executed instructions sequentially, one after another, at the rate of only one instruction per clock cycle. The processors of these serial computers were also usually "scalar," in that they could apply an instruction on only one item of data at a time. Serial computers began to reach a point of diminishing performance returns in the 1960s, as the rate of increase in CPU performance began to slow, and failed to keep pace with the increasingly complex problems coming from science, engineering, and mathematics. Computer-based simulations increasingly suffered from the serial bottleneck. The time lost to the sequential execution of instructions on a central processor, for example, was a

³²⁹ Murray, *The Supermen*, 116-117, 129; MacKenzie, "The Influence of the Los Alamos and Livermore Laboratories," 181-182; Harold S. Stone, *High-Performance Computer Architecture* (New York: Addison-Wesley Publishing Company, 1990), 2-5, 9-13, 17-20; Sean Blanchard, oral history interview by Nicholas Lewis, Los Alamos, NM (June 19, 2017); Neil R. Lincoln, moderator, "Reminiscences of Computer Architecture and Computer Design at Control Data Corporation," OH 321. Oral history interview moderated by Neil R. Lincoln (May and September 1975), 19, 21-22. Charles Babbage Institute, University of Minnesota, Minneapolis. <http://hdl.handle.net/11299/104327>.

particular problem on large datasets that required the same instruction to be executed on each number in a dataset. While a significant amount of programming and computer-system development since the 1950s had been focused on mitigating the limitations of serial processing, including the prefetching of instructions, out-of-order execution, and other means of preventing processor idle time, serial computing was, itself, a bottleneck that grew worse with each supercomputer generation.³³⁰

Several technical approaches offered solutions to the serial bottleneck in the late 1960s. Seymour Cray at Control Data opted for the *multiprocessor* approach for the 8600, quadrupling the number of scalar CPUs working in parallel on the same problem. A competing approach was *vectorization*, in which a single processor could execute the same instruction across all, or least a large portion, of a dataset. Computer companies in the late 1960s and early 1970s sought to bring both solutions to the scientific-computing market. Both the multiprocessor and vector approaches raised major programming hurdles, requiring the development of programming aids and languages that could take advantage of the drastically different means of processing data that the new schemes required. Codes like those at Los Alamos, which were originally written for serial processing, would have to be modified for either form of parallelism, although more so for multiprocessor systems, as the programmer would have to balance the workload on each processor. A simpler, but still undesirable task with vector computers, before the

³³⁰ Nicholas Metropolis, "Limitations of Serial Computers, 10/3/90," A-98-011, Nicholas Metropolis Papers, box 36, folder 1, Parallel Processing, LANL Archives; William J. Kaufmann III, and Larry L. Smarr, *Supercomputing and the Transformation of Science* (New York: Scientific American Library, 1993), 41-42; J. T. Pinkston, "Supercomputer Trends and Needs: A Government Perspective," in Nicholas Metropolis, et al., eds, *Frontiers of Supercomputing* (University of California Press, 1986), 124-126, 135-139;

development of automatic vectorizing compilers, was the need to manually identify which parts of a code were to be sent to the vector units on the supercomputer. Not all codes were amenable to vectorization, so neither option was ideal for all situations. While Cray's 8600 would have entailed a difficult transition to a multiprocessor architecture (a challenge for which C Division's leadership believed its software R&D groups were equal), Cray estimated the 8600's performance would be roughly double that of its nearest competitor, which appealed to C Division's desire for ever-escalating processing power.³³¹

Consistent with Cray's preference to build cutting edge architectures with older reliable components, the 8600 was intended to use discrete transistors and other components wired together by hand, rather than newer integrated-circuit technology. Cray contended that integrated circuits, as of the early 1970s, would make the computer larger than soldered discrete components, a severe problem when attempting to reduce the spacing—and hence the electronic delays—between all parts of the computer. However, the tight packaging of discrete components on the 8600 test units resulted in excess temperatures, too many delicate solder points, and high failure rates that even Cray could not resolve. The supercomputer could not be packaged tightly enough to offer the desired performance *and* function reliably. Despite years of labor in the

³³¹ Nicholas Metropolis, "An Introduction to Massively Parallel Processing, 9-5-90," A-98-011, Nicholas Metropolis Papers, box 36, folder 1, Parallel Processing, LANL Archives; Elzen and MacKenzie, "The Social Limits of Speed," 48. It should be noted that the multiprocessor and vector approaches were not mutually exclusive from a technical standpoint, as many later vector computers would also employ multiple processors. It was the difficulty of making either approach to parallelism work in the early 1970s that usually encouraged developers to stick with one approach or the other until the early 1980s. When Cray began his working on his first vector-based supercomputer, he considered making it a multiprocessor system, but ultimately opted against the added complexity.

Chippewa Falls lab, Cray realized that the 8600, as it had been designed, was simply not going to work. Control Data, which had expanded and diversified rapidly following its successes in the 1960s, could not afford to begin the 8600 development project anew, as Cray had requested. Rather than wait a year or so for funding to become available to redesign the 8600, as Control Data CEO William Norris had proposed, Cray opted to leave Control Data and form his own company in 1972.³³²

During a site visit to Chippewa Falls in May of 1972, C-Division representatives, including Roger Lazarus and Jack Worlton, learned from Seymour Cray himself about the new company. Even before the company had an official name, with "Cray Research Associates" being floated as a possibility, and before the decision whether or not the company would be independent from Control Data was settled, Cray spelled out the lessons he had learned from the 8600, and how they would influence his future design work. Cray stated that he would be starting from a "blank sheet of paper," a favorite, if not completely accurate description Cray used for his design of a new machine. The new supercomputer, Cray argued, would employ integrated circuits to resolve the packaging problems that plagued the 8600. Other details (and there were many) were not yet settled. With frequent subsequent visits, and mail and phone communications, the LASL group remained in close contact with Cray and what would become Cray Research, Inc. (CRI).³³³

³³² William Spack, oral history interview by Nicholas Lewis, Los Alamos, NM (June 25, 2014); Lewis interview with Folkner; Hall, *Generation of Wealth: The Rise of Control Data*, 26, 92; Murray, *The Supermen*, 116-117, 129; Lincoln, moderator, "Reminiscences of Computer Architecture and Computer Design at Control Data Corporation," 19, 21-22.

³³³ William J. Worlton, "2/8/72 to 2/20/74," lab notebook, HPC Archive. Despite Cray's claim that he started each computer design from scratch, all of his supercomputer architectures bore significant similarities. Sidney Fernbach noted this aspect of Cray's designs, and stated that it was one reason moving

Despite C Division's interest in Cray Research's as yet unnamed supercomputer, there was no telling when that machine would be available. LASL's latest request for bids to computer vendors had called for a supercomputer system available no later than the end of 1972. C Division had hoped that the 8600 would be nearing completion by that point, but that timeframe was no longer possible. As of spring 1972, the time of Cray's announcement, only two of the bids LASL had received from computer vendors, one from TI, and the other from a Control Data project competing with the 8600, were close to meeting the desired delivery timeframe and performance specifications. In C Division's view, neither option was a desirable replacement for the 8600, which had promised a great performance leap. The 8600's multiprocessor design also appeared to offer better average performance increases across LASL's code base compared with the vector approach, which would provide performance benefits only to the vectorizable codes at Los Alamos. Scalar performance increases on the vector supercomputers, depending on the architecture, might be modest.³³⁴

The 8600 project continued at Control Data after Cray's departure, but at reduced funding and priority, aiming to cheaply resolve the system's design problems. Cray's desired redesign of the 8600 had not been feasible in light of Control Data's financial challenges in 1971 and 1972. In addition to costs stemming from an ongoing legal battle with IBM, a recession in 1970 had significantly slowed computer orders. As financial

to the Cray-1 from Cray-designed CDC supercomputers was relatively easy from a software-transition standpoint. See: Marilyn Ghausi, interview with Sidney Fernbach, February 8, 1989. Accessed October 3, 2016. <http://www.computer-history.info/Page1.dir/pages/Fernbach.html>.

³³⁴ William J. Worlton, "2/8/72 to 2/20/74."

pressures mounted, Control Data was expending large sums on its *other* supercomputer project, the STAR-100.³³⁵

The STring ARray (STAR) 100, with ERA veteran and long-time Cray assistant James Thornton as its principal designer, was, along with TI's contemporaneous Advanced Scientific Computer (ASC), the first attempt to produce a commercially viable vector supercomputer. The STAR-100 project dated back to 1964, with a request from Livermore to produce a "parallel digital computing system" for that lab by 1970. Sidney Fernbach, leader of Livermore's Computational Department, had originally wanted the AEC to fund development of the SOLOMON parallel-computer concept that Daniel Slotnick from Westinghouse proposed in 1962. The AEC insisted on a bid process open to other computer vendors, but the AEC's terms on the bid eliminated all but Control Data's proposal of the STAR-100 (IBM did not submit a serious entry). The formal request from the AEC, on behalf of Livermore, for the development of a large parallel computer, which the AEC planned to lease, was signed in 1966.³³⁶

The project explored the technical options for using parallelism, in this case, vectorization, to produce vast increases in computer performance for scientific applications, and to overcome the serial bottleneck before the problem hampered the AEC weapons labs. Livermore provided some development funding for the STAR

³³⁵ Hall, *Generation of Wealth: The Rise of Control Data*, 164; Joe Herman, and Jane Strovas, "Control Data: A History," September, 1979, CBI 80, Control Data Corporation Records, Series 24, Historical Archives and Company History, Box 2, Folder 14, CBI, 18, 33.

³³⁶ MacKenzie, "The Influence of the Los Alamos and Livermore National Laboratories," 192; Ghausi interview with Fernbach. For more information about the SOLOMON computer concept and its announcement, see: Daniel L. Slotnick, W. Carl Borck, and Robert C. McReynolds, "The SOLOMON Computer," in *AFIPS '62 (Fall) Proceedings of the December 4-6, 1962, Fall Joint Computer Conference*, 97-107.

project (\$5 million, according to Sidney Fernbach, out of the \$75 million the STAR would ultimately cost Control Data in development funds). Livermore also provided unclassified sample codes that Control Data used for optimizing the system's architecture and assessing the computer's performance for the AEC weapons labs. The code selection would later come back to haunt Livermore and Control Data Corporation.³³⁷

The CDC 6600 had been the most complex computer-design project Control Data had undertaken before work began on the STAR. As a consequence, much like the Stretch and LARC projects in the 1950s, it was difficult to accurately predict the timeline. To make matters worse, the STAR-100 relied upon a never-before-commercialized implementation of parallel computing. The STAR's design team attempted to maximize the computer's commercial appeal by including a large number of instructions catering to a variety of users. This move ballooned the STAR's instruction set to over 200 instructions, over three times that of the 6600. The added complexity and development time contributed to the STAR's significant delays and design setbacks. As a consequence, the STAR missed its original 1970 delivery estimate, and the revised 1971 estimate as well. This meant that Control Data was supporting two supercomputer

³³⁷ Charles J. Purcell, "The Control Data STAR-100—Performance Measurements," in *AFIPS '74 Proceedings of the May 6-10, 1974, National Computer Conference and Exposition*, New York: ACM (1974), 385-387; MacKenzie, "The Influence of the Los Alamos and Livermore National Laboratories," 192; D. J. Theis, "Vector Supercomputers," CBI 80, Series 7, Control Data Corporation Records, Engineering, Research and Development, 1957-1986, Box 23, Folder 9, General, 1968-1975, CBI, 1-3; Neil R. Lincoln, "Supercomputer Development—The Pleasure and the Pain," *Datamation* 23 (May, 1977): 221-222; W. H. Brummet, Jr., Chairman, AEC Contract Board, to Charles J. Purcell, Control Data Corporation, "Request for Proposal for a Parallel Computing System," October 5, 1966, CBI 80, Control Data Corporation Records, Series 7, Engineering, Research and Development, 1957-1986, Box 33, Folder 19, AEC - LLL, 66-71, CBI; Ghausi interview with Fernbach; Murray, *The Supermen*, 129.

projects, each featuring different deployments of parallelism, while both exceeded their estimated budgets and passed their delivery dates.³³⁸

The CDC 8600 on the backburner, and the uncertainty when Seymour Cray's first independent supercomputer would be available (regardless of its still-unknown capabilities), brought the STAR and the TI ASC to the forefront of LASL's computer search. With delivery estimates of late 1972 and 1973, respectively, the CDC STAR-100 and TI's ASC were the only next-generation supercomputers close to fulfilling LASL's RFP (Request for Proposals) for a new system. C Division had reservations about the suitability of the STAR's untried architecture for LASL's computing requirements. The STAR appeared more heavily geared toward vector processing than was suitable for most LASL codes, and the real-world performance of the completed machine was difficult to pin down, due to frequent changes Control Data made. Control Data's performance estimates based on Livermore sample codes were the best information available until Los Alamos could conduct its own code tests. Unanswered technical concerns were not the only stumbling block C Division faced in its increasingly constrained decision-making process.³³⁹

With the 8600 eliminated from the running, Livermore and the AEC encouraged Los Alamos to take delivery of the STAR, partly because of the AEC's investment in the project, and partly in the interest of maintaining common architectures between LASL

³³⁸ Lincoln, "Supercomputer Development," 223; Elzen and MacKenzie, "The Social Limits of Speed," 48.

³³⁹ Control Data Corporation, Advanced Studies Division, "A Proposal to the Atomic Energy Commission for the Los Alamos Scientific Laboratory for a STAR-100 Computer System, July 1971," Edward A. Voorhees Papers, box 44, folder 8, CDC-STAR Proposal, Oct. 1971, LANL Archives; MacKenzie, "The Influence of the Los Alamos and Livermore Laboratories," 195.

and Livermore. A joint purchase between LASL and Livermore might even allow the AEC to negotiate a better deal with Control Data. Part of the technical justification was that maintaining a common set of architectures between the two labs might also make software exchanges easier than if LASL and Livermore installed incompatible systems. The two weapons labs were already able to share math libraries, programming language and compiler development, and other software tools and libraries, across their existing collections of mostly Control Data supercomputers. Both labs adding a STAR might have allowed them to share the burdens of troubleshooting and software development for the new architecture. The AEC weapons labs no longer commanded the singular purchasing power they once had, due to the customer base for supercomputing expanding considerably in the 1970s. In fact, the first order for a STAR-100 came from General Motors. Combining LASL and Livermore's forces in the purchase of the STAR might have offset some of the labs' overall decline in purchasing power.³⁴⁰

C Division's leadership was disinclined to adopt a new supercomputer simply for the "sake of appearances," as Frank McGirt later described it. It was still unclear that the STAR would be suited to LASL's requirements, and purchasing such a system along with Livermore meant both labs would be taking the same gamble at the same time. CCF users largely echoed C Division's reticence. Some users were uninterested in any dramatic shift toward parallel computing, due to the disruptions that necessarily

³⁴⁰ MacKenzie, "The Influence of the Los Alamos and Livermore Laboratories," 191; F. McGirt to Staff Meeting Attendees, STAR Project Members, STAR Review Board, "C Division STAR Commitments," January 17, 1974, Edward A. Voorhees Papers, series 33, box 47, Computers and Computing, Contract Negotiations, folder 4, Contract Negotiations, LANL Archives; William J. "Jack" Worlton, interview by Nicholas Lewis, St. George, UT, November 29, 2015.

accompanied radical architectural changes. Those who were not opposed to parallelism had largely favored the 8600 as LASL's next-generation supercomputer, because it appeared to offer an easier transition to parallelism than full-blown vector processing, as the 8600 was essentially four scalar computers working in unison. The STAR, before any official performance testing was complete, appeared to be heavily biased toward vector performance, but LASL computer users were uncertain about the suitability of their code base for a rapid shift to vectorization. Their more immediate need was for a significant increase in scalar performance. The difficulty of assessing the STAR's performance, which was giving C Division reservations about the machine, also meant that LASL had few concrete justifications to outright reject the computer.³⁴¹

Sidney Fernbach remained enthusiastic about the STAR, but, unlike his enthusiasm for time-sharing, this did not transfer to Los Alamos. Part of LASL's reservations over the STAR reflected Los Alamos' wariness about the heightened involvement of AEC headquarters in lab management. The AEC's interest in greater managerial control and oversight of its labs had been escalating since the late 1960s. In the early 1970s, the AEC attempted to institute a more rigorous, "algorithmic" procedure for computer selection and auditing at its labs, along with tighter controls on hardware purchases and the operation of lab computer centers. The result was a charged debate

³⁴¹ W. F. McNaughton, to L. D. Altman, "STAR Proposal," March 10, 1970, CBI 80, Series 7, Control Data Corporation Records, Engineering, Research and Development, 1957-1986, Box 27, Folder 12, Los Alamos Scientific Laboratory, CBI; McGirt "C Division STAR Commitments"; J. Worlton, "Some Estimates of the Performance of STAR Relative to the 7600," Edward A. Voorhees Papers, box 47, folder 1, Contract Negotiations - 1, series 33, LANL Archives; Lewis interview with William J. Worlton; T. R. Suter, to Distribution, "LASL Information Session, 9/23/73," September 27, 1973, CBI 80, Control Data Corporation Records, Engineering, Research and Development, 1957-1986, box 36, folder 4, Atomic Energy Commission (AEC) - Los Alamos Scientific Laboratory (LASL), 1970-1974, CBI.

between the AEC and C Division at Los Alamos, eventually forcing the AEC to scale back some of its demands. Roger Lazarus argued that the proposed "algorithmic" procedures for selection and auditing were "unmanageable for LASL use." The proposed process did not recognize the technical or qualitative differences between business and scientific computing. The AEC plan hoped to shape purchasing and assess the operations of both types of computing, without considering the nuanced influence of LASL's scientific research. On the AEC's proposal of using an "algorithm" to make any computer selection at all, Lazarus stated that there "are many ways to fake an algorithmic selection" that would render the algorithmic approach an artificial tool for selecting the desired computer, rather than finding the best computer. The algorithmic approach ended there for Los Alamos, but the AEC's intervention in LASL supercomputer purchasing did not.³⁴²

Into the spring of 1973, LASL had yet to commit to a STAR-100 order. Roger Lazarus had approved estimates of the installation costs of a STAR in the CCF, but continuing delays with the supercomputer's shipment estimates, which had slipped into 1974, provided C Division with time to continue searching for other options. Because of the STAR's delays, Livermore took delivery of an interim CDC 7600 in June to meet its growing computing demand. LASL was considering doing the same, after having gone three years without the installation of an additional supercomputer in the CCF. While

³⁴² Harold Agnew, "Headquarters Involvement with Details of Laboratory Management," Edward A. Voorhees Papers, box 51, folder 14, Headquarters Memos, 1974, series 38, LANL Archives; Hal Butler, MP-1, to Harold Agnew, DIR-O, "The Horrors of Computer Procurement," July 2, 1974, Edward A. Voorhees Papers, box 51, folder 14, Headquarters Memos, 1974, series 38, LANL Archives; R. B. Lazarus, et al, "On the Selection of ADP Systems – A Position Statement," November, 1970, Edward A. Voorhees Papers, box 51, folder 10, Selection Criteria, series 38, LANL Archives.

contemplating strategies for offsetting the computer operation's need for more capacity, C Division remained in contact with Seymour Cray about his new computer, which revealed tantalizing hints, but concrete details about the machine and its expected delivery time frame remained unknown. C Division continued to explore the possibility of the delayed TI ASC supercomputer, but it appeared to offer few advantages over the STAR. LASL also hosted visits from Burroughs and IBM for updates on their prospective scientific-computing systems, but found little of interest in their offerings. In June 1973, in need of greater computing capacity, and with no clearly superior systems then on the market, LASL consented to a joint LASL/Livermore contract with Control Data for the purchase of two STAR-100s at about \$10 million each, including peripherals.³⁴³

Oral-history interviews and C-Division documentation describe substantial AEC pressure on Los Alamos to accept the STAR contract. Sidney Fernbach described this period in lab computing history as one of increasing bureaucratisation from AEC management, with non-computer experts becoming heavily involved in the decision-making processes. Los Alamos' computing experts and users likewise rankled under the AEC's tighter organizational and budgetary interventions in the Lab's computing operation, of which the STAR became a prominent example. Roger Lazarus, just before stepping down as C-Division leader, distributed a pointed memo to the CCF users group:

³⁴³ "Proposed CDC STAR-100 System for LASL," February 16, 1973, Edward A. Voorhees Papers, box 48, folder 7, CDC-STAR Evaluation, series 32, LANL Archives; Worlton, "2/8/72 to 2/20/74," lab notebook; R. B. Lazarus, C-DO, to CCFUG Representatives, "CDC STAR," July 5, 1973, Edward A. Voorhees Papers, box 45, folder 3, Multiple Procurement Working Papers-1, series 32, LANL Archives; "Computer Installation CDC STAR-100," March 1, 1972, Edward A. Voorhees Papers, box 45, folder 3, Multiple Procurement Working Papers-1, series 32, LANL Archives.

"Once again, the CDC STAR has been selected for LASL and LLL [Livermore] over the TI ASC-4X," confirming for users who might still be in doubt, and communicating that the decision had not been a welcome one. He clarified that the delivery was scheduled for June 30, 1974, one year out, and that C Division and its users needed to come together to "make sure that the STAR can be reasonably well used as soon as it gets here."³⁴⁴

III. Falling STAR

With the signing of the STAR contract, C Division established the STAR Advisory Board, made up of eight high-level members of C Division, including Roger Lazarus (alternate division leader), Bill Spack (assistant division leader), Charlie Folkner (leader of C-2 CCF Systems Support), Jack Worlton (alternate leader of CADP), and Ed Voorhees (CADP leader), to determine how best to evaluate the STAR and integrate it into LASL's computing environment. Frank McGirt, the new division leader, did not have high hopes for the STAR. He thought the Lab should treat it as an experimental system that, one day, might become useful for weapons design.³⁴⁵

McGirt was hesitant to divert too many resources away from other C-Division R&D efforts, like code development work on LTSS, and the completion of the Hydra

³⁴⁴ Lazarus to CCFUG Representatives, "CDC STAR," July 5, 1973"; "Computer Installation CDC STAR-100"; William J. Worlton, "1/10/73 to 1/31/75," lab notebook, HPC Archive; McGirt, "C Division STAR Commitments"; Worlton, "Some Estimates of the Performance of STAR Relative to the 7600"; Ghausi interview with Fernbach.

³⁴⁵ C. Folkner, C-DO, to F. McGirt, C-DO, "HELP! STAR Testing," September 25, 1973, Edward A. Voorhees Papers, box 46, folder 4, Contractual, series 33, LANL Archives; McGirt, "C Division STAR Commitments"; Suter, "LASL Information Session, 9/23/73"; T. R. Suter to G. N. Hanson, "LASL Visit to STAR-100 DD 10/16/73," October 16, 1973, CBI 80, Control Data Corporation Records, Engineering, Research and Development, 1957-1986, box 36, folder 4, Atomic Energy Commission (AEC) - Los Alamos Scientific Laboratory (LASL), 1970-1974, CBI; Zelenka, "LANL Visit to Headquarters - May 30, 1974."

network, to adapt LASL's workload to the STAR. In his view, it would be a mistake to trust the untried computer with too many mission-critical jobs too quickly. Control Data's representatives to LASL noted the Lab's cautious approach to the computer, but believed it to be a product of Frank McGirt wishing to appease Harold Agnew and the computer users who had recently ousted Roger Lazarus. LASL computer users were, in this line of thought, simply being obstinate. Control Data believed LASL's requests for extensive performance testing before delivery were an attempt by C Division's leaders to demonstrate that the division was falling in line.³⁴⁶

Whatever Control Data's representatives to LASL believed, CDC's delays in completing the STAR's hardware for testing left C Division worried. Livermore did not appear to share LASL's concerns, having, in August of 1973, committed to a second STAR system. Into the fall of 1973, the LASL advisory board pressed Control Data to allow C Division to run its own performance tests on the STAR architecture, because the computer was finally approaching its deliverable form. The timetable for testing was narrow; the STAR contract stated that any testing LASL wished to perform had to be submitted to Control Data by December 30 of that year, and only two months after that were permitted for the two parties to remediate the problems discovered through testing. With the STAR still undergoing changes through late 1973, time was running out. After Control Data agreed to LASL providing its own test routines, Charlie Folkner, a C-Division software expert, sent out an urgent request to Frank McGirt in the fall of 1973,

³⁴⁶ McGirt, "C Division STAR Commitments"; Zelenka, "LANL Visit to Headquarters - May 30, 1974."

requesting that the division work quickly to pull together a test procedure for the STAR. Jack Worlton delivered the division's test routines shortly before the year was out.³⁴⁷

The test results, reported to C Division in early 1974, were alarming. The Livermore test codes had been far too easy a test for the STAR. During the computer's development, Livermore could not provide Control Data with classified weapons-design codes. To satisfy security concerns, Livermore provided declassified portions of a computationally intensive hydrodynamics weapons design code, which Control Data used to determine if the STAR had met the benchmarks specified in the AEC contract. The Livermore test codes were poorly chosen, as they were representative of only a tiny percentage of the codes then in use at either Livermore or LASL, and they failed to reveal that the STAR would be unsuitable for the majority of the weapons labs' most important and demanding computing problems. The LASL test routines demonstrated that the STAR suffered from extremely poor scalar performance. The STAR-100, as mentioned earlier, was one of the two first commercially available "vector" supercomputers, but Control Data had failed to strike a balance with scalar performance, which was critical for most of the jobs in use at both labs. Favoring vector processing to an overwhelming degree, the STAR offered approximately one-quarter the scalar processing power of a CDC 7600. Some thought it was actually closer to the 6600 in scalar mode.³⁴⁸

³⁴⁷ Folkner, "HELP! STAR Testing." It was largely Charlie Folkner's work in developing test codes during the supercomputer search of the mid-1960s that determined IBM's S/360 computer lineup would not meet the performance specifications IBM had promised to LASL; Zelenka, "LANL Visit to Headquarters - May 30, 1974"; Bill Spack, C-DO, to Charlie Folkner, C-DO, "LASL's STAR Configuration, LLL's Twin STAR Arrangement, Aug. 10, 1973," Edward A. Voorhees Papers, box 46, folder 3, STAR Proposal March 1973, series 33, LANL Archives.

³⁴⁸ Worlton, "Some Estimates of the Performance of STAR Relative to the 7600"; "Effectiveness of STAR on the LASL Workload," A-93-015 C-Division, box 4, folder 5, (Oct.-Sept. 1975), LANL Archive; MacKenzie, "The Influence of the Los Alamos and Livermore Laboratories," 195; Lewis interview with

Because of its architectural design, the STAR was also imbalanced in favor of "long" vectors, giving the machine poor "short" vector performance. Datasets to be processed were first loaded into the STAR's main memory, before being moved to the processor, where the vector instructions were applied to each of the data points. The results were then returned to main memory. The time required to move datasets to memory and back, which made the STAR a "memory-to-memory" vector supercomputer, meant that the computer could only provide users with a significant performance benefit on large datasets of 100 to 1,000 numbers. Datasets larger than 1,000 numbers often exceeded the system's available storage capacity. On smaller datasets, of fewer than 100 numbers, the STAR offered little to no performance advantage over a traditional scalar computer. In fact, LASL testing determined that the STAR's short-vector performance was only about 40% that of the CDC 7600. To make matters worse, if the STAR reached a branch in the code, requiring the computer to load and process new data, the vector data in memory had to be removed and then reloaded later. This design element severely limited the computer's ability to reach its full performance on most LASL/Livermore jobs. The Livermore test code, because it contained no data-dependant branches, failed to identify this key problem with the STAR architecture. The parties involved did not fully recognize these issues with the STAR's architecture until late in the development

Morrison; Lewis interview with Folkner; Charles A. Folkner to John Zelenka, January 15, 1974, CBI 80, Control Data Control Data Corporation Records, Engineering, Research and Development, 1957-1986, box 36, folder 4, Atomic Energy Commission (AEC) - Los Alamos Scientific Laboratory (LASL), 1970-1974, CBI.

process largely because of the novelty of vector processing in the early 1970s, and the general lack of experience with the technology.³⁴⁹

Los Alamos brought its concerns to Control Data following the first battery of tests, requesting that the company complete the STAR's instruction set (which was still not fully operational in early 1974) so that a final, representative set of performance measurements could be conducted before shipment. Control Data's official response to the LASL tests was muted, so as not to risk the AEC refusing to accept the three STARs on order. The company agreed to conduct further tests, but wished to continue with the contractually stipulated Livermore codes that had failed to reveal the STAR's architectural weaknesses. Meanwhile, preparations to ship the LASL and Livermore STARs continued.³⁵⁰

Exchanges between LASL and Control Data's Albuquerque office grew heated, with Charlie Folkner in disbelief that CDC was still arguing that the STAR performed as specified. Folkner accused CDC of renegeing on its commitments.³⁵¹ With Control Data continuing its computer deliveries as planned, C Division called a meeting with TD Division, the largest user of CCF resources, to ask if those weapons-program users could

³⁴⁹ "Effectiveness of STAR on the LASL Workload"; Worlton, "Some Estimates of the Performance of STAR Relative to the 7600"; MacKenzie, "The Influence of the Los Alamos and Livermore Laboratories," 195; Frank McGirt to John Zelenka, March 19, 1974, CBI 80, Control Data Control Data Corporation Records, Engineering, Research and Development, 1957-1986, box 36, folder 4, Atomic Energy Commission (AEC) - Los Alamos Scientific Laboratory (LASL), 1970-1974, CBI; Lewis interview with Morrison.

³⁵⁰ H. N. Frazier to E. H. Russell, "Report on Trip to LASL March 26 & 27, 1974," CBI 80, Control Data Control Data Corporation Records, Engineering, Research and Development, 1957-1986, box 36, folder 4, Atomic Energy Commission (AEC) - Los Alamos Scientific Laboratory (LASL), 1970-1974, CBI; McGirt to Zelenka, March 19, 1974.

³⁵¹ Charles A. Folkner to John Zelenka, March 15, 1974, CBI 80, Control Data Control Data Corporation Records, Engineering, Research and Development, 1957-1986, box 36, folder 4, Atomic Energy Commission (AEC) - Los Alamos Scientific Laboratory (LASL), 1970-1974, CBI.

make effective use of the STAR. Jack Worlton stated that their answer was a unanimous "No!" C and TD Divisions together decided that, "It is better to get *no* machine than to get the *wrong* machine," and began a concerted effort to cancel the STAR contract.³⁵² The AEC's pressure on Los Alamos to agree to the computer in the first place, and the fact that LASL had signed the \$10 million contract, made securing permission to cancel an unlikely proposition. On top of that hurdle were the fears within Los Alamos and AEC headquarters that Control Data would take the Lab to court if it attempted to back out of the contract. C Division needed time to build its case against the STAR, which meant working to delay the computer's shipment as long as possible. The division's leadership believed that once the computer had shipped, backing out of the contract would truly have been impossible.³⁵³

In addition to requesting more time-consuming testing, C Division also requested system-performance improvements that were well beyond what Control Data could provide before the expected shipping date. Meanwhile, the division requested the pre-shipment completion of specific hardware and software features that were not scheduled to be ready until later in the year. Despite Control Data accusations that Los Alamos was deliberately delaying shipment, the adroit tactic worked, as the LASL requests

³⁵² William J. Worlton, "The Cray-1: The Machine and The Man," in *The Cray-1 Supercomputer: 30th Anniversary Event—Celebrating the Man and the Machine*, lecture, Mountain View, CA: Computer History Museum (September 21, 2006).

³⁵³ Frazier to Russell, "Report on Trip to LASL March 26 & 27, 1974"; Lewis interview with William J. Worlton.

themselves needed to go through a lengthy negotiation process that delayed the AEC's final approval for delivery.³⁵⁴

While C Division stalled Control Data, Jack Worlton brought the LASL test results to Harold Agnew, hoping that the director could intervene with the AEC. Worlton's presentation convinced the director that the STAR would significantly harm LASL's computing operation, and, as a result, impact the Lab's mission for years to come, while users and LASL software developers struggled needlessly to adapt the Los Alamos workload to the underperforming STAR architecture. Agnew took the issue up the chain of command to AEC headquarters, where Agnew, Worlton, and Frank McGirt presented their case on April 30, 1974, only two months before the STAR's contracted delivery date. The AEC was in the midst of its own organizational shakeup, and would be abolished by the US Congress later that year. Amid the turmoil, and the strength of LASL's presentation, the AEC relented, and allowed Los Alamos to renegotiate the contract with Control Data. Because LASL *had* signed the original contract, and because the STAR technically met the contract benchmark requirements (based on poorly chosen test codes), the company was not obligated to cancel the agreement. However, Control Data did not wish to add to the already growing negative publicity of the much-delayed, over-budget, and underperforming STAR. Nor was it eager to begin another legal battle, after the lengthy (although successful) antitrust case against IBM had recently concluded. The company agreed to renegotiate with Los Alamos, meeting repeatedly in May 1974.

³⁵⁴ "Assume the STAR Contract is Canceled," April 30, 1974, Edward A. Voorhees Papers, box 47, folder 4, Contract Negotiations, Series 33, LASL Archives; William J. Worlton, "1/10/73 to 1/31/75"; Lewis interview with William J. Worlton.

Rather than accept the STAR, LASL used the \$10 million in the contract to order two CDC 7600s and a smaller Cyber 73 to quickly meet the pent-up computing demand in the CCF. Delivered by early fall 1974, the three systems nearly doubled LASL's existing computing capacity, and truly doubled the number of time-sharing computers available to classified users.³⁵⁵

Sidney Fernbach, one of the STAR's staunchest supporters since the 1960s, conceded in a meeting with Control Data that the Los Alamos test findings should be used as the metric by which CDC and Livermore worked to address the computer's architectural issues. The additional work pushed back the delivery of Livermore's two STARs until September of 1974 and January of 1975, respectively. Fernbach later commented that, had LASL not beat him to it, he would have tried to back out of the Livermore contract with Control Data. Having just released Los Alamos from its STAR, the company was not willing to cancel Livermore's order of two STARs. Upon arrival, the STARs were compatible with only about a dozen Livermore codes out of over two hundred. Fernbach later argued that the conversion effort was instructive for Livermore, providing a testing ground for vectorizing codes that many experts considered impossible to vectorize. Conversely, most users and Livermore management considered what became a seven-year-long attempt to bring the STAR-100 into regular production use to have been a fruitless exercise. Fernbach thought that the Livermore STAR controversy

³⁵⁵ Zelenka, "LANL Visit to Headquarters - May 30, 1974"; Lewis interview with William J. Worlton; Westwick, *The National Labs*, 269; Lewis interview with Folkner; Dorr, "Computer Science and Services Division," 2.

led to his dismissal as division leader in the early 1980s, when changing political winds at the lab sought a scapegoat for what users perceived as a boondoggle.³⁵⁶

The STAR-100 incident illustrated how several prominent forces, including technical bottlenecks in supercomputing, the declining prominence of the AEC labs in the supercomputing market, and the heightened managerial and operational oversight of AEC headquarters, interacted with the AEC weapons labs in the early 1970s. The incident also illustrated how Los Alamos and Livermore reacted differently to those forces. While both labs sought solutions to technical bottlenecks, and both experienced greater oversight (and outright interference) from the AEC, the reaction of C Division reflected LASL's characteristic caution toward disrupting the Lab's primary mission. It also reflected C-Division's distinctive reticence toward overreliance on outside computing expertise and untested technical systems. In the case of the STAR-100, the desire of LASL computer users to avoid significant disruptions in their existing workload, and C Division's traditional wariness of unproven technologies being introduced too quickly into the production environment, unified LASL computer users and experts who had otherwise been divided on the issue of time-sharing. The encroachment of AEC-level oversight served as a galvanizing force for C Division and its computer users in their resistance to a technical solution that grew out of, and better reflected, Livermore's more-risky operational approach.³⁵⁷

³⁵⁶ T. R. Suter to Distribution, "S. Fernbach Visit to STAR-100 DD 5/7/74," CBI 80, Control Data Control Data Corporation Records, Engineering, Research and Development, 1957-1986, box 36, folder 4, Atomic Energy Commission (AEC) - Los Alamos Scientific Laboratory (LASL), 1970-1974, CBI; Ghausi interview with Fernbach; MacKenzie, "The Influence of the Los Alamos and Livermore Laboratories," 192-193.

³⁵⁷ Ghausi interview with Fernbach.

For Control Data, the STAR-100, of which only five were produced, with two of those remaining at CDC, and the 8600, which was ultimately canceled altogether, solidified CEO William Norris' view that the future of computing was not in pushing the limits of supercomputer performance, but in applying the existing capabilities of computers toward services. The transfer of IBM's lucrative Service Bureau Corporation to Control Data as part of the antitrust settlement helped further that ambition. While the company developed follow-on supercomputers based on the STAR, its hesitation over whether it should remain in the supercomputing market between 1974 and 1976, and its decision to build upon its existing investment in the STAR architecture, ultimately cost Control Data its primacy in the supercomputing field. For Los Alamos, canceling the STAR contract had significant benefits for the Lab's computing operation, with one of the most important being the financial, staff-time, and other resources not expended on adapting the Lab's workload to a troubled supercomputer architecture. Los Alamos failed to acquire a next-generation supercomputer, but adding two CDC 7600s and a Cyber 73 provided immediate, much-needed production capacity. The STARs at Livermore, which were far from production-ready when shipped, never truly reached their full capacity before being decommissioned seven years later, and would have exacerbated the growing computer-capacity deficit at Los Alamos. Possibly most significant for LASL's long-term computing operation, avoiding the STAR left Los Alamos free as an organization to pursue a supercomputing icon.³⁵⁸

³⁵⁸ Elzen and MacKenzie, "The Social Limits of Speed," 48, 51-52; Murray, *The Supermen*, 131-132. For a more in-depth discussion of Control Data and its computing services, see: Jeffrey R. Yost, *Making IT Work: A History of the Computer Services Industry* (Cambridge, MA: The MIT Press, 2017).

IV. Suddenly Seymour

In October of 1973, Jack Worlton, Roger Lazarus, and Charlie Falkner, following a meeting with Control Data about the STAR-100, met with Seymour Cray, his chief engineer and company co-founder, Lester Davis, and other employees of Cray Research, Inc. in Chippewa Falls, Wisconsin for an update about his new supercomputer. The meeting happened in Cray's new lab, only a few hundred yards from the Control Data facility that Cray built in the 1960s as his design studio for the CDC 6600. For the as-yet unnamed new system, Cray had fundamentally reconsidered his approach to the CDC 8600 design. The Los Alamos group had already heard in 1972 about Cray and his design team's move to integrated circuits. The update in late 1973 was even more significant, as it indicated that Cray had learned not only from the packaging mistakes with the 8600, but also from the architectural mistakes of the STAR-100.³⁵⁹

Unlike the multiprocessor 8600, the new computer opted for the greater simplicity of a single processor. Cray had also backed off from the overly ambitious eight-nanosecond clock speed of the 8600, for a more sensible, but still very fast twelve nanoseconds. The new system, like the STAR and TI ASC, would be a vector supercomputer, to contend with the serial bottleneck. Cray's would be a "short" vector machine that could reach full vector performance on smaller datasets, unlike the "long" vector-oriented STAR. Cray accomplished this feat by incorporating vector registers into the CPU, making it a "register-to-register" vector supercomputer, rather than a "memory-to-memory" system like the STAR. Registers were small, but extremely fast memory

³⁵⁹ William J. Worlton, "2/8/72 to 2/20/74"; Murray, *The Supermen*, 130.

partitions residing directly on the processor. The processor executed vector instructions on the datasets residing in the registers, without having to move that data from or to the much slower and more distant main memory. Cray's team also recognized the STAR's problem with scalar performance. The new Cray computer, instead, balanced scalar and vector processing, resulting in an estimated speed improvement of four to five times the CDC 7600's vector performance and about three times its scalar performance. The result was a supercomputer that offered a significant improvement on scalar jobs over preceding supercomputers, while offering a workable means of contending with the serial bottleneck.³⁶⁰

The Cray approach provided an easier transition to parallel supercomputing for traditional supercomputer customers like LASL. Because the new computer handled scalar jobs as well as it handled vectorized applications, users would experience significant performance improvements on their conventional scalar programs while the development of programming aids and methods eased the transition to vectorized applications over time. The LASL group left the meeting intrigued over the possibilities the Cray computer offered Los Alamos. In early April of 1974, while preparing LASL's case to the AEC against the STAR contract, Worlton met again with Cray Research, which had grown to twenty-one employees. In the meeting, Worlton learned that the new supercomputer's development work had progressed quickly. About 90% of the computer's board design was complete. In all, about 500 flat, densely packed integrated-

³⁶⁰ Worlton, "1/10/73 to 1/31/75"; Richard M. Russell, "The Cray-1 Computer System," in *Communications of the ACM* 21, no. 1 (January 1978), 63-64, 67-68; Murray, *The Supermen*, 133; Jack Worlton, "The Cray-1 Computer," A-93-015, Records of C-DO, box 3, Cray-1, folder 9, FEDSIM Report, LANL Archives.

circuit (IC) modules, connecting a grand total of 144,000 IC chips, made up the computer processor. Cray estimated the completion of all boards and wiring by 1975, before a third- or fourth-quarter 1975 shipment date. A drawing Jack Worlton made in his lab notebook during the meeting also depicted what would become the signature cylindrical "C" shape of the Cray-1, complete with a bench-height power supply at its base.³⁶¹

In January of 1975, Seymour Cray gave a lecture at Los Alamos officially announcing the Cray-1. By the end of Cray's talk, C-Division's leadership was reportedly convinced that the Lab should purchase the new computer. Seymour Cray, anticipating the response his claim would elicit, stated to the Los Alamos audience that Livermore would likely receive the first Cray-1. Cray's intimation that LASL's rival might be the first recipient of his computer prodded C Division into action. That same month, Jack Worlton, the new C Division leader Harold Agnew appointed after Frank McGirt stepped down due to illness, met with the new Energy Research and Development Administration (ERDA) Division of Military Application (DMA) to ask about the feasibility of Los Alamos acquiring the first Cray-1.³⁶²

The 1974 Congressional Energy Reorganization Act had abolished the Atomic Energy Commission in response to concerns over its dual role of both promoting and regulating atomic energy. Those primary functions of the AEC were divided between two new agencies, the Nuclear Regulatory Commission and ERDA, which inherited the

³⁶¹ Worlton, "1/10/73 to 1/31/75"; Murray, *The Supermen*, 133-134; W. J. Worlton, C-DO, to Seymour Cray, August 27, 1974, Records of C-DO, box 1, (1970-1975) Cray-1, LANL Archives.

³⁶² Worlton, "The Cray-1"; "A Proposal for the Acquisition of a Cray-1 Computer for the Los Alamos Scientific Laboratory," Records of C-DO, box 1, folder 3, (Feb. - March 1975) Cray-1; J. Worlton, C-DO, to CCFUG Committee on Planning, "Discussions with DMA on Cray-1 Acquisition," February 3, 1975, Records of C-DO, box 1, (1970-1975) Cray-1, folder 4, LANL Archives.

AEC's lab system. ERDA, itself, was short-lived, being replaced by the Department of Energy in 1977. ERDA and the DOE better reflected the diversification of research beyond atomic energy that had occurred in the labs since the AEC was founded in 1947, and reflected growing economic and national-security concerns in the US after the energy crisis of 1973.³⁶³

For Los Alamos, the transition to ERDA oversight would pose new administrative challenges. Among the first challenges for LASL was the heightened oversight ERDA inherited from the AEC regarding lab supercomputer purchasing, and the new stipulations ERDA managers put in place after the STAR fiasco. Stu Milam of the Program Analysis and Budget office of the Division of Military Application (DMA), the ERDA division that housed the lab system, was "enthusiastic" about the DMA acquiring the new Cray computer, and believed that LASL had a "leg up" on Livermore. With Milam's support, the DMA approved Jack Worlton's proposal to install the first Cray-1 at Los Alamos. Sidney Fernbach likewise hoped to acquire the first Cray-1 for Livermore, but, despite Seymour Cray's goading at Los Alamos, the DMA refused his proposal, because of Livermore's contract for two STAR-100s. Los Alamos, after backing out of its STAR contract, did not possess a state-of-the-art supercomputer, and could successfully argue to the Congressional Joint Committee on Atomic Energy (JCAE) its need for one when the Cray-1 was announced. A request for proposals (RFP) among computer vendors had to take place, per ERDA requirements, before a contract could be signed for a specific system, but, given the lack of rivals that could match the Cray-1's

³⁶³ Westwick, *The National Labs*, 269-270, 311.

capabilities, an RFP was almost a formality. All the same, two significant problems stood in the way of LASL's purchase of the Cray-1.³⁶⁴

The first problem was financial. Los Alamos would be in the middle of its fiscal year when the Cray-1 became ready to ship, and the complete \$7.5 million purchase price of the new computer would not be in the budget until fiscal year 1977. Los Alamos did have \$2 million in the FY 1976 budget, which the DMA suggested might be used as partial payment until FY 1977 made the full purchase amount available. If Cray Research would not accept a partial payment, other arrangements would have to be made. The other problem was ERDA oversight. While the DMA and ERDA management were supportive of one of the ERDA labs being the first to acquire the Cray-1, ERDA stipulated that its labs could not purchase an unevaluated supercomputer, which included the Cray-1. Cutting-edge supercomputers had previously been exempted from the AEC's typical pre-purchase evaluation requirements. The AEC and Livermore's involvement with the STAR's development and its consequent architectural issues were a source of chagrin and congressional pushback at ERDA, and spurred the policy change. Like Los Alamos, Cray Research also had a financial problem. The startup company needed to begin sales, and not for partial payments, to remain afloat, and could not wait long for a firm purchase commitment.³⁶⁵

Jack Worlton credited Seymour Cray with proposing to install the computer at Los Alamos for a free, six-month evaluation. After this trial, the computer could be

³⁶⁴ Worlton, "The Cray-1"; Ghausi interview with Fernbach;

³⁶⁵ J. Worlton, C-DO, "Discussions with DMA on Cray-1 Acquisition"; William J. Worlton, "March 1975 to April 1976," lab notebook, HPC Archive; Worlton, "The Cray-1."

returned if it proved unsatisfactory. The proposed evaluation had the benefit of solving a great many problems. For Los Alamos, it would provide time to find the funds to purchase the Cray-1, and it would satisfy the ERDA stipulation that a supercomputer undergo evaluation before the signing of a purchase contract. LASL could undertake the ERDA-required vendor bids while the evaluation took place. For Cray Research, the evaluation would secure the installation of its first machine, with Cray betting that Los Alamos would find the system satisfactory. It would also provide the company's new computer with a proper evaluation, encouraging customers who might otherwise be wary of a company that still had not sold a single computer. Worlton recorded in his lab notebook that the evaluation had to be sufficiently difficult, so that ERDA and the JCAE would not see through the "ploy" of getting the Cray-1 in the door before funding had been secured or an evaluation completed.³⁶⁶

The evaluation contract, signed between Cray Research and the Regents of the University of California (which administered LASL under contract with ERDA) on March 3, 1976, stipulated a six-month evaluation between April and September 1976. The required LASL Request for Proposals to allow vendors to compete with the Cray-1 went out in January of 1976, running until that October. ERDA stipulated in the contract with Cray Research that Los Alamos conduct no production work on the computer during the evaluation period. Another clause gave LASL the option to purchase the computer, rather than return it, at the end of the evaluation. The computer arrived in early April, and was installed under the watchful eye of Seymour Cray himself, who was noted to be

³⁶⁶ William J. Worlton, "March 1975 to April 1976," lab notebook, HPC Archive; Worlton, "The Cray-1."

polite but firm with the unloading crew to be careful with their charges. The Cray-1 was installed in Building 132's machine room 290, the room that originally housed the IBM 7030 Stretch, which had been retired and removed in 1971. Building 132, also known as the Central Computing Facility (CCF), had undergone five significant additions between 1964 and 1971, increasing its internal volume by about seven times in the process. The multiple expansions provided floorspace for the Control Data machines and ancillary equipment, and additional office space for C Division and vendor support personnel, to cope with the surging demand for computing capacity. Positively diminutive compared with the IBM computers of the 1960s, and significantly more compact than the Cray-designed Control Data systems that preceded it, the Cray-1 fit easily in the cavernous CCF machine room, and was ready to begin its evaluation on April 5.³⁶⁷

Among historians of computing, and at the former ERDA labs themselves, the evaluation is typically regarded as the beginning of the ERDA labs' transition to computers from Cray Research, but the transition actually began over a year earlier, when LASL, Livermore, and several other ERDA labs began negotiating and forming their operating-system strategies for the Cray-1, long before Serial No. 1 arrived on the Los Alamos mesa. Before Seymour Cray delivered the January 1975 talk at Los Alamos, Jack Worlton sent a letter to Cray, himself, in August of 1974 regarding operating-system strategies for the new computer. Worlton, speaking independently of LASL and C

³⁶⁷ "Loan Agreement ZA-89666, 2/25/76," Records of C-DO, box 2, Cray-1, folder 9, LANL Archives; Normand, et al., "First Record of Single-Event Upset on Ground, Cray-1 Computer at Los Alamos in 1976," 3114-3115; "Letter to Vendors," January 20, 1976, Records of C-DO, box 3, folder 12, Class VI Computer Installation at LASL, LANL Archives; "ENG-SK-7402 CCF BLDG. SM-132," February 1974, Book Case 1-1, Buildings folder, HPC Archive.

Division, asked that Cray take seriously the growing issue of software compatibility when organizations selected new supercomputers for their users. Worlton noted that supercomputer customers, a growing and diversifying demographic in government, academia, and industry, were increasingly sensitive to software portability. Already planning ahead to the possibility of a LASL proposal to purchase Cray's new computer, Worlton suggested that Cray Research consider supporting Control Data's KRONOS batch/time-sharing operating system. KRONOS had, at least among supercomputer operating systems, a relatively broad user base, including the ERDA nuclear-weapons labs. While the ERDA labs tended to use custom OSES on their top-tier supercomputers, Worlton understood and acknowledged that it would make little sense for Cray Research to bend its software strategy around an OS with a non-existent commercial presence, like CROS or LTSS.³⁶⁸

The personal appeal to Seymour Cray did not result in CRI pursuing KRONOS compatibility, but CRI's operating-system and overall software strategy would ultimately take shape in response to the broadening supercomputer user base beyond the ERDA lab system. The labs would, ultimately, develop custom operating systems for the Cray-1, but Worlton's letter marked the beginning of an intense debate at LASL over its OS strategy for the Cray-1, a system that would not be available for production use for nearly two years. The intensity of the LASL OS debate reflected the excitement that the Cray-1 generated in the Los Alamos computing community. Some of the stir was due to the fact that Seymour Cray, already a much-beloved figure at LASL, appeared to be rescuing the

³⁶⁸ W. J. Worlton, C-DO, to Seymour Cray, August 27, 1974, Records of C-DO, box 1, folder 4, (1970-1975) Cray-1, LANL Archives.

Lab. Many computer users and experts at LASL feared that, as the computer market expanded, the US nuclear-weapons labs might become less economically viable sources of revenue for computer vendors. As a consequence, supercomputers entering the market might be poorly suited to LASL's needs, at least without effectively abandoning simulation codes that had been calibrated through numerous, expensive live tests. The Cray-1, with its excellent balance between scalar and vector processing, appeared to be the lifeline many LASL users were seeking.

The Cray-1 had a lot to offer most users, including those who simply wanted faster scalar processing. The OS debate intensified in the early spring of 1975, primarily between C and TD Divisions, with previous animosities surrounding the batch vs. time-sharing issue reasserting themselves, as computer-operations staff and users from various user divisions vied to determine the shape of LASL's operating-system strategy. Representatives made their respective cases for specific lines of OS development. Hoping to establish a formal dialog for the decision-making process, rather than allow the informal discussions already taking place to become heated and divisive, the C Division Office produced guidelines for evaluating the key concerns when selecting an OS for the Cray-1. Jack Worlton wrote an example argument, following the new guidelines, where he presented his case for basing the Cray-1's OS on CROS, citing that a CROS-based operating system would preserve the maximum amount of existing LASL software and programming investment for both C and TD Divisions. Worlton's argument was based on the fact that the bulk of LASL's application codes were developed for CROS, which represented 80% of the jobs being run in the CCF, versus 15% and 5% for LTSS and

KRONOS, respectively. Although TD Division had been the primary champion behind introducing LTSS to Los Alamos, the vast majority of its computer usage was on systems running CROS. Worlton's analysis provided a point-for-point breakdown of each operating system's merits and detriments, encouraging others to do the same.³⁶⁹

The common theme among most of the Cray-1 OS debates was of software portability, the ability to move operating systems (and the applications that ran on them), between computer systems with relative ease. "Relative" being the keyword, as the move between architectures always brings unforeseen difficulties. Worlton's argument in favor of developing CROS for the Cray-1 dealt with software portability in two ways. The first was, with LASL's great experience with that operating system, and its being the most heavily used OS at the Lab, transitioning CROS to the Cray would minimize the redevelopment effort for most of LASL's code base. Three of LASL's four CDC 7600s were dedicated to CROS, while a single 7600 ran LTSS, and the two remaining 6600s and one Cyber 73 were running KRONOS. The CROS and KRONOS machines were typically running at saturation each week, while the LTSS machine, dedicated almost entirely to TD Division's use, was the only CCF computer that experienced regular periods of idle time. The second Cyber 73 in the CCF had replaced the CDC 6600 as the Hydra network's front-end computer, and ran the custom Hydra OS. (The rollout of the Hydra network will be discussed later in this chapter.)³⁷⁰

³⁶⁹ J. Worlton, to Distribution, "Guidelines for the Cray-1 Operating System Decision," Records of C-DO, box 1, folder 1, (March-June, 1975) Cray-1, LTSS, LANL Archives.

³⁷⁰ Worlton, "Guidelines for the Cray-1 Operating System Decision"; "The Computer Science and Services Division Annual Report, January—December 1975," Los Alamos Scientific Laboratory, Los Alamos, NM, LA-6228-PR (March, 1976), 5.

The second part of Jack Worlton's portability argument was aimed at the distant future. By establishing a single OS that would move from the Cray-1 to its successors, LASL would expend far less software redesign effort in the long term. Robert Thorn, the leader of TD Division, and the leading figure in C Division's reorganization in 1973 over the time-sharing debate, argued effectively the opposite of Worlton's viewpoint. Thorn's stance was unsurprising, given his strong preference for time-sharing operating systems, despite the batch-only CROS being the most heavily used OS across LASL as well as in TD Division specifically. Thorn's argument was in favor of a joint OS-development project with Livermore and Cray Research. According to Thorn, CRI indicated that it would not develop an OS independently, but was willing to collaborate on an operating system with its first customer. Thorn contended that, considering LASL and Livermore would both likely acquire early Cray-1 computers, a collaboration among the two labs and the computer's vendor would do the most to prevent a duplication of effort. Settling on a single OS between LASL and Livermore would, as a consequence, ease the transfer of codes and other software between the two labs.³⁷¹

C Division's C-4 CCF Software Development and Support group broke with both Worlton and Thorn, arguing in favor of an entirely new operating system. C-4's experts believed that neither CROS nor LTSS could take full advantage of the Cray-1's new architecture. C-4's group leader, Paul Brown, argued that, despite the upfront costs, creating a new OS from scratch would allow C Division to develop a modular and

³⁷¹ Worlton, "Guidelines for the Cray-1 Operating System Decision"; R. N. Thorn, TD-DO, to D. P. MacDougall, ADW, "Joint LASL-LLL Development of Operating System for Cray-1 Computer," June 5, 1975, Records of C-DO, box 1, folder 1, (March-June, 1975) Cray-1, LTSS, LANL Archives.

portable OS design. For the sake of rapid development, and portability between computer architectures, C-4 raised the possibility of writing the majority of the new OS in a high-level programming language. This was an unusual choice, because operating systems were usually coded in a low-level assembly language. In agreement with Worlton's suggestion that the OS design, whichever operating system was ultimately chosen, should build upon and expand LASL's existing expertise in OS development, the C-4 proposal argued for a project that would ensure Los Alamos attracted and retained sufficient programming talent necessary to support the Lab's OS requirements in the long term. Brown believed the innovative approach of coding the OS in a high-level language would help LASL push the boundaries of OS development, and, in addition to the draw of creating an OS for the already attention-getting Cray-1, would aid in the staff recruitment and retention effort.³⁷²

Clare Nielson, alternate group leader of the new Controlled Thermonuclear Research Division's CTR-6 Plasma Theory group, represented another viewpoint in the search for an OS. CTR-6 strongly objected to LTSS as the basis of an ongoing LASL OS effort, and considered it an irresponsible use of the Cray-1. Nielson, a time-sharing user, contended that LTSS lacked many basic features of other time-sharing operating systems, and was "absolutely inadequate" for performing state-of-the-art computations. Nielson ascribed most of LTSS' efficiency and other problems as a time-sharing OS to its "dreadful" LRLTRAN compiler, which Nielson argued was substantially less efficient for highly interactive code-development work than LASL's other times-sharing OS,

³⁷² Paul J. Brown, C-4, to Jack Worlton, C-DO, "Cray-1 System," May 29, 1975, Records of C-DO, box 1, folder 1, (March-June, 1975) Cray-1, LTSS, LANL Archives.

KRONOS, sourced from Control Data. Nielson contended that much of the debate over LTSS versus other operating systems in use at LASL was due to the unwillingness of "partisans" to allow performance comparisons to be conducted between LTSS, CROS, and KRONOS. With a proper comparison, Nielson continued, the argument for LTSS would have no legs. CTR-6, as a result, favored KRONOS as a basis for the Cray's new OS, because its implementation of time-sharing provided better features and performance than the Livermore-sourced LTSS. The various parties involved in the OS debate became increasingly polarized over the spring of 1975, largely in response to TD Division's increasingly hardline support of LTSS, with non-LTSS supporters feeling bullied by TD's political and financial power as the dominant user of LASL computing resources.³⁷³

The Cray Review Board, a collection of representatives from C Division and the major user divisions that had been organized to determine how LASL should use a Cray-1, met repeatedly after C Division instigated the purchase proposal for the new computer. A survey of the group provided data on the OS debate. TD Division was largely in favor of LTSS, while other major users, such as T Division, were firmly in favor of CROS, as most of their code work was designed around the batch OS. A few other divisions, like CTR, were split between KRONOS and CROS. Most time-sharing users (who were still the minority of computer users at LASL) outside TD Division were inclined toward KRONOS. C Division was not unified on a specific course of action, with C-DO favoring CROS, but other groups and individuals in the division favoring either a new

³⁷³ C. W. Nelson, CTR-6, to C. M. Gillespie, "LTSS System," May 25, 1975, Records of C-DO, box 1, folder 1, (March-June, 1975) Cray-1, LTSS, LANL Archives.

OS, CROS, or even KRONOS. A C-Division proposal during the second Cray Review Board meeting suggested that the development of a KRONOS-CROS link for the new computer could provide the best of both worlds for time-sharing and batch users. The Cray-1 was not designed to operate as a stand-alone system, requiring at least a minicomputer to serve as a front-end machine between the supercomputer and its users. In C Division's proposal, the Cray-1 would, itself, utilize a version of CROS, while its front-end machine, possibly a CDC Cyber-series mainframe, ran KRONOS. The KRONOS front-end machine would provide users with interactive time-sharing capability as it compiled, submitted, and monitored the batch jobs being processed on the Cray-1. This compromise arrangement would preserve the Cray-1's runtime for batch applications, which represented the majority of LASL's 80,000 jobs per month, while still providing interactive capabilities for the users who wanted it.³⁷⁴

The KRONOS-CROS link proposal began to gain traction among users, such as Alternate T-Division Leader Richard Morse, who wanted interactive capability, but considered LTSS "not acceptable for production execution." Robert Thorn attempted to seek managerial intervention in favor of LTSS by sending a contentious memo to the Assistant Director for Weapons, Duncan MacDougall. MacDougall had been the leader of GMX Division before his promotion during the Agnew reorganization of the LASL Weapons Program. In the memo, Thorn made a series of contentious points about LTSS, claiming it to be superior to the alternative time-sharing systems, which users of LASL's

³⁷⁴ "Cray Review Board Meeting, 5-16-75," Records of C-DO, box 1, folder 1, (March-June, 1975) Cray-1, LTSS, LANL Archives; R. L. Morse, T-DO, W. J. Worlton, C-DO, "Comments on Draft Memo "CCF Software Support Policies," 16 June 1975, Records of C-DO, box 1, folder 1, (March-June, 1975) Cray-1, LTSS, LANL Archives; Russell, "The Cray-1 Computer System," 71.

other time-sharing OS, KRONOS, flatly refuted. When Thorn stated that LTSS was all Livermore used on its large systems (indicating, for Thorn, that LASL should do the same), one critic responded, "if LTSS is all they [Livermore] have, I feel sorry for them." Thorn also claimed that LTSS was far more heavily used in the LASL CCF than the usage data indicated, and called it "the only game in town," with respect to the viable options on which to base a Cray-1 OS. Not for the first or last time, he would also insinuate that C Division and those who supported CROS or KRONOS were jeopardizing the Lab's mission. He recommended that the Director's Office should directly control CCF policies (and that TD Division should play a major role in determining those policies), because the computer operation was too mission-critical, Thorn contended, for C Division to make consequential policy decisions on its own.³⁷⁵

The memo was sufficiently inflammatory, and generated such strong responses from other divisions and computer users, that it effectively worked against Thorn's case with upper management. Approximately two years' worth of experience with LTSS on LASL's CDC 7600s had given many LASL computer users outside TD Division a negative impression of the Livermore-sourced OS, and few could see why it should be on the Cray-1. Many of its problems, critics cited, were due to the age of the OS, which had entered development in the mid-1960s, and had been ported to several Control Data computer models since that time. As a consequence, it was not as modern or as tightly

³⁷⁵ Morse, Worlton, "Comments on Draft Memo 'CCF Software Support Policies'"; "Cray Review Board Meeting, 5-16-75"; Nelson, to Gillespie, "LTSS System"; J. E. George, C-6, to J. Worlton, F. Dorr, C-DO, "June 3, 1975 Memo from R. N. Thorn to D. P. MacDougall, 'LTSS and Interactive Graphics,'" June 6, 1975, Records of C-DO, box 1, folder 1, (March-June, 1975) Cray-1, LTSS, LANL Archives; R. C. Lindberg, C-9, to W. J. Worlton, C-DO, "Memo, Thorn to MacDougall, 19 May 75," May 20, 1975, Records of C-DO, box 1, folder 1, (March-June, 1975) Cray-1, LTSS, LANL Archives.

coded as most of the other, newer options. Another issue was that, being created at Livermore, it was more difficult to maintain and support at Los Alamos, which lacked the in-depth knowledge the OS' developers possessed, a significant factor when tailoring codes to work efficiently in an interactive time-sharing environment. One of Thorn's critics referred to "LTSS Syndrome," exemplified by Thorn, in his memo, calling it "our" [TD Division's] operating system. Richard Lindberg, group leader of C-9 Computer Systems Engineering, expressed bewilderment at Thorn's many disputable claims, including one of the memo's justifications for why TD's preferred OS should ultimately go onto the Cray-1. Thorn contended that TD had insufficient LTSS capacity in the CCF, but Lindberg pointed out that the CDC 7600 set aside only for LTSS remained unsaturated every week, unlike the three CROS and three KRONOS systems. Lindberg wondered, in response to Thorn's statement that the Director's Office should set C-Division policies, why, if TD was so vital to LASL's mission, did the Director's Office not set TD policies, too?³⁷⁶

The internal C-Division discussion of Thorn's memo to MacDougall raised possibly the most pertinent, unanswered question about Thorn's position, asking *why* TD so fervently wanted LTSS. Thorn alluded to unspecified "serious consequences" if LASL failed to adopt LTSS on the Cray, prompting Bill Buzbee, group leader of C-3 Mathematical Applications Programming to respond with, "We must insist that TD start providing logic and facts to support LTSS. They have repeatedly said that they want

³⁷⁶ Lindberg, to Worlton, "Memo, Thorn to MacDougall, 19 May 75."

LTSS, but to my knowledge they have never said why (nor does this memo)."³⁷⁷ TD's position left many non-TD personnel, and even some dissenting personnel within TD, puzzled, in the face of considerable technical data and user experiences demonstrating sizeable issues with LTSS as a production platform. Weapons designer and TD-3 Weapons Outputs group leader, Paul Whalen, was among the dissenting voices in TD. An expert in the application of computational capabilities in weapons design, Whalen stated that the porting of LTSS and its codes to the Cray-1 that Thorn demanded would hobble the capabilities of the new computer. He suggested that TD should simply tell C Division "what our general goals are, and quit trying to tell them how to accomplish those goals." Whalen argued that all TD Division really wanted was interactive capability for large codes, while also being able to communicate with the 7600s and files on the Hydra storage network. Thorn's insistence that those goals could only be obtained through LTSS and the unsuitable LRLTRAN compiler, Whalen concluded, would "keep the Cray-1 as a toy for years."³⁷⁸

What justification Thorn did provide for LTSS came down to a desire to "gain parity" with Livermore, an ill-defined goal that would have ruffled feathers at the insinuation of Los Alamos not already being Livermore's equal. Richard Lindberg quipped, "Doesn't TD in some way 'compete' with some segment of LLL [Livermore]?" in response to the notion of intentionally becoming technically dependent upon, and

³⁷⁷ B. L. Buzbee, C-3, to W. J. Worlton, C-DO, "My Comments on Thorn's Memo to MacDougall," May 20, 1975, May 20, 1975, Records of C-DO, box 1, folder 1, (March-June, 1975) Cray-1, LTSS, LANL Archives.

³⁷⁸ Buzbee, to Worlton, "My Comments on Thorn's Memo to MacDougall"; Lindberg, to Worlton, "Memo, Thorn to MacDougall, 19 May 75"; P. P. Whalen, TD-3, to G. R. Spillman, TD-DO, "TD Meeting, Cray-1 Discussions," April 28, 1975, Records of C-DO, box 1, Cray-1, folder 2, (March-May, 1975) Cray-1, LANL Archives.

maintaining supposed "parity" with, LASL's rival. LASL observers developed the impression that important segments of the TD community and its leadership sought to reorient Los Alamos' weapons program toward that of Livermore's approach to research, with its riskier, but sometimes more ostentatious projects, compared with the stronger theoretical orientation and grounding at Los Alamos.³⁷⁹

LTSS was a product of, and was better suited to, Livermore's more rapidly iterative and experimental orientation, with less extensive computer-based modeling of a design concept's theoretical basis than at Los Alamos. LTSS' limited popularity at Los Alamos outside TD Division was an indicator of how poorly the OS meshed with LASL's predominantly batch-oriented computing demands. If TD's leadership hoped to replicate the Livermore approach to research at Los Alamos, adopting the other lab's computing technologies and policies might have been a key component. The problem in 1975 was that a large number of Los Alamos' computer users, and its Computing Division, were not interested in replicating Livermore's computing operation. As a consequence, unlike the successful managerial intervention in C Division policy in 1973, which brought classified time-sharing to the CCF, TD's attempt to coerce LASL into adopting LTSS as the Cray-1 OS resulted in widespread resistance.³⁸⁰

TD's interest in forcing LTSS onto the Cray-1, and its interest in establishing direct control of LASL's computing policies, coincided with a similar process occurring at Livermore. Sidney Fernbach reported that non-expert users and administrators from

³⁷⁹ Buzbee, to Worlton, "My Comments on Thorn's Memo to MacDougall"; Lindberg, to Worlton, "Memo, Thorn to MacDougall, 19 May 75"; Ghausi interview with Fernbach.

³⁸⁰ Buzbee, to Worlton, "My Comments on Thorn's Memo to MacDougall"; Lindberg, to Worlton, "Memo, Thorn to MacDougall, 19 May 75"; Ghausi interview with Fernbach.

that lab's weapons program became increasingly involved in directing Livermore's computing operation over the 1970s. The process occurring at Livermore was even more severe than that at Los Alamos in the early 1970s, when the AEC and the LASL weapons program began to impose stronger bureaucratic and financial constraints on Los Alamos computing. Livermore's largely uncritical support of the STAR-100 supercomputer, an error unjustifiably placed almost entirely on Fernbach's shoulders, was an early manifestation of the growing interventionism over Livermore computing policies. The escalating drive within LASL's weapons program to exert direct control over its local computing operation appeared to be an import from Livermore, much as LTSS had been.³⁸¹

The important differences between the two labs' approach to weapons research played an important role in LASL's resistance to software and policies that were better suited to Livermore. Indeed, it was Los Alamos' greater diversification of non-weapons research compared to Livermore that provided the financial roadblock to TD's effort to assert direct control over LASL's computing operation. Budget cuts in the late 1960s, and the AEC's crackdown on its labs straying beyond their central missions in the early 1970s, meant that the respective weapons programs at Los Alamos and Livermore grew significantly in relation to each lab's overall research funding. The consequences for the already more weapons-oriented Livermore were more severe, with its weapons program growing to represent nearly 80% of that lab's funding, versus 68% at Los Alamos. Although TD *was* LASL's single largest computer-using division, its command of 54%,

³⁸¹ Ghausi interview with Fernbach.

or \$3.9 million of the \$7.2 million, of the total machine usage in 1974, meant that a very large minority of computer users did not share TD's research goals. As a consequence, far more users outside LASL's weapons program had a vested interest in that lab's computing operation and policies than was the case at Livermore, in addition to the members of the weapons program who did not share the interest of TD's leadership in drawing closer to Livermore's technologies or computing policies.³⁸²

In June of 1975, Robert Thorn sent a memo to C-Division leader, Jack Worlton, walking back some of his claims and recommendations in his memo to MacDougall, including Thorn's statement that C Division should not be allowed to determine its own policies. Thorn ascribed much of the "fiasco" of the OS debate to a lack of proper communications between C and TD Divisions, and suggested that formal, monthly meetings take place between management of both divisions. Thorn was not the first to express frustration with C Division's communication with users. Jack Worlton acknowledged in the preceding week that C Division was insufficiently communicative regarding its planning, a problem that had given some users the impression that the division was "doing secret things behind their back" regarding the Cray-1 OS decision.³⁸³

The "secret things" that caused users concern had, in fact, been C Division's software groups working to determine the feasibility of the KRONOS-CROS link before

³⁸² Buzbee, to W. J. Worlton, "My Comments on Thorn's Memo to MacDougall"; Lindberg, to Worlton, "Memo, Thorn to MacDougall, 19 May 75"; Nicholas Nagy, C-1, to W. Jack Worlton, C-DO, "TD Computer Usage," June 17, 1975, Records of C-DO, box 1, Cray-1, folder 1, (March-June, 1975) Cray-1, LTSS, LANL Archives; Westwick, *The National Labs*, 321.

³⁸³ R. N. Thorn, TD-DO, to W. J. Worlton, C-DO, "C-Division/TD-Division Communications and Understanding of Computing Requirements," June 12, 1975, Records of C-DO, box 1, folder 1, (March-June, 1975) Cray-1, LTSS, LANL Archives; J. Worlton, to Bill Huntman, C-4, memo, June 4, 1975, Records of C-DO, box 1, folder 1, (March-June, 1975) Cray-1, LTSS, LANL Archives; Lewis interview with William J. Worlton.

presenting it to the Cray Review Board as an option. While the link concept proved popular among users as a potential solution to the OS problem, C Division's inward focus when researching the possibility made "users very nervous," in Worlton's words, and unnecessarily escalated the situation with TD Division. C Division generally had a spotty track record of communicating well with computer users, and the rapid deterioration of the Cray-1 OS debate into a noisy and ill-tempered fracas in early 1975 strongly demonstrated how that deficiency could make the tense relationship with some segments of LASL's weapons program even worse.³⁸⁴

Worlton called for more and better communications channels with user divisions, including regular meetings with users to elaborate on the projects and planning being undertaken in C Division. These meetings were in addition to the usual user group and procurement meetings, and the collaboration meetings that occurred between C-Division's development groups and those of other divisions. While C Division's communication efforts improved in some respects, the issue was not completely one-sided. Veterans of C Division noted that the tumultuous debates primarily with weapons-program users over operating-systems strategies and other software planning sometimes made attempts at communicating difficult, with some C-Division personnel becoming increasingly reluctant to risk altercations with particular users from the weapons program. One former C-Division programmer recalled regularly waiting in his office to avoid being berated by a specific member of TD Division in the parking lot at the end of

³⁸⁴ Thorn, to Worlton, "C-Division/TD-Division Communications and Understanding of Computing Requirements"; Worlton, to Huntman, memo, June 4, 1975; Lewis interview with Morrison; Lewis interview with William J. Worlton.

workdays. Thorn's animosity toward C Division was still clear in his memo, claiming that C Division was too ignorant of TD's complex research needs to appreciate why CROS and KRONOS were unsuitable for TD research. The pushback from non-TD users against LTSS, in Thorn's view, indicated that TD was "premature" in using LTSS to create a "new era" of cooperation between TD and C Division.³⁸⁵

Discussions and feasibility assessments over the Cray-1 OS continued throughout 1975. For the same reason ERDA forbade LASL from using the Cray-1 for production runs during the evaluation, C Division could not commit to the development of an OS specific to the Cray before or during the evaluation. The division's general OS R&D work could continue, exploring the possibilities for an OS that would carry the Lab into the next decade. The Cray-1 had to go through the same bid process as any other vendor's computer once the evaluation was completed. As the computer's arrival neared, site preparations also continued, with the CCF's power, cooling, and other facilities being made ready for the new system. Formal time tables for the evaluation solidified in the first months of 1976, and test-code development work began in close coordination between C Division and the users whose codes were selected.³⁸⁶

Unlike the arrival of most supercomputers at LASL, where C Division encouraged all interested users to redevelop their codes to work on the new system, the C

³⁸⁵ Thorn, to Worlton, "C-Division/TD-Division Communications and Understanding of Computing Requirements"; Worlton, to Huntman, memo, June 4, 1975; Jim Clifford, interview by Nicholas Lewis, Los Alamos, NM (July 23, 2014).

³⁸⁶ "Possible Chronology of Events for the Evaluation of the Cray-1," December 10, 1975, Records of C-DO, box 3, folder 12, Class VI Computer Installation at LASL, LANL Archives; W. Spack, C-DO, to D. Bryson, SP-DO, "Class VI Computer: Site Preparation Equipment," January 23, 1976, Records of C-DO, box 3, folder 12, Class VI Computer Installation at LASL, LANL Archives; "C-Division Newsletter," 3/25/1976, Edward A. Voorhees Papers, box 57, folder 3, C-Division Meetings, 1976-1978, series 42, LANL Archives.

Division Office had to consider carefully which codes could be used during the Cray-1 evaluation. The reason was that Cray Research would not have a full operating system available for the evaluation, or much software of any kind, aside from the Cray Assembly Language. The company's included Benchmark Operating System (BOS) was less an operating system, and more a collection of basic services to make interfacing with the supercomputer possible via the Data General Eclipse minicomputer that Cray Research included as a front-end machine. Per ERDA stipulations, LASL was not allowed to modify the Cray operating system more than was necessary for the evaluation. As a consequence, there were considerable limitations on the software that would be able to run successfully on the computer. The lack of FORTRAN or any other high-level language for the Cray-1 meant that all codes used on the machine would have to be written in assembly language, which was complicated by the fact that the computer's vector capabilities made code conversions less straightforward than on a typical supercomputer. Cray Research would later release a vectorizing FORTRAN compiler, but not in time to be of use during the evaluation. Despite the software limitations, C Division planned to have at least six production codes operating on the Cray by the end of its evaluation. The hope was, without violating the contract stipulations, to have codes ready to jump quickly into production use once the final selection and purchase had been completed.³⁸⁷

³⁸⁷ "C-Division Newsletter"; B. L. Buzbee, C-3, to W. J. Worlton and F. W. Dorr, C-DO, "Class VI Evaluation Parameters," January 13, 1976; Marge McCormick, ed., "Computing Requirements of the ERDA Advanced Computing Laboratories, July, 1977," 3.

Just before the evaluation began, Jack Worlton stepped down as C-Division leader, with Agnew appointing former alternate C-Division leader, Fred Dorr, in Worlton's place. Worlton cited the stress and politics of the position in his decision to step down, but he remained in C-Division management until 1981. The evaluation that Worlton helped to instigate began on schedule, in April 1976. Despite ERDA and the Department of Military Applications' overall interest in LASL acquiring a Cray-1, high-level ERDA and DMA management, wary of the evaluation being treated as a formality, warned C Division that "if LASL doesn't do a good job on the evaluation, we [Los Alamos] may not even get a machine." Fred Dorr commented to Harold Agnew that, by focusing on doing well with the evaluation, C Division had an opportunity to provide ERDA with more information about a pre-purchase supercomputer than ever before, while also potentially setting the standard by which supercomputer performance benchmarking would be done in the future.³⁸⁸

The Cray-1 evaluation was officially a cooperative effort among LASL, ERDA, and FEDSIM, an independent organization that conducted computer evaluation services for federal agencies. The involvement of ERDA and FEDSIM, both of which monitored the evaluation and had to approve LASL's proposed testing criteria and methods, was intended to ensure that the evaluation was conducted in good faith. ERDA was not going to risk another embarrassment, knowing that the US Congress was watching closely after the STAR-100 fiasco. The results of the evaluation remained confidential between CRI

³⁸⁸ Lewis interview with William J. Worlton; F. W. Dorr, C-DO, to Distribution, "Early Use of the Cray-1 Computer," March 30, 1976, Records of C-DO, box 3, folder 10, Cray-1 Evaluation, LANL Archives; F. Dorr, C-DO, to Agnew/Thorn, "Cray-1 Evaluation," May 24, 1976, Records of C-DO, box 4, Cray-1, folder 1, (May-July, 1976) Interim Report of the Cray-1 Evaluation, LANL Archives.

and the groups conducting the evaluation until the final report, with interim reports being produced throughout the testing process. The first of several preliminary reports appeared in May of 1976 to demonstrate that the computer was functional, and that testing was successfully underway. The evaluation criteria specified a minimum set of requirements for acceptable performance and reliability from a "Class VI" supercomputer.³⁸⁹

A "Class VI" supercomputer was defined as being capable of approximately 20 to 60 million floating point operations per second (the "class" designation was only roughly assigned by performance. Machines like the Cray-1 and STAR-100 were technically far beyond the 60 "Mflops" mark). Previous-generation supercomputers fell into lower class designations, with the CDC 7600 being a Class V, and the CDC 6600 being a Class IV. The preliminary results were in comparison with the same tests being conducted on one of LASL's CDC 7600s, which revealed that the Cray-1 was approximately three-times faster than the 7600 in scalar work, and closer to five-times faster in vector processing. The early Cray-1 results were considered extraordinary, particularly in light of the minimal, unoptimized software being used. Machine downtime and system errors were attributed to bugs that Cray Research could continue to resolve during the testing phase. Testing continued through September, with LASL running two types of sample workloads through the machine. The first included codes that represented a more general

³⁸⁹ J. Worlton, C-DO, to D. MacDougall, ADW, "FEDSIM Participation in the Cray-1 Evaluation," March 15, 1976, Records of C-DO, box 3, folder 10, Cray-1 Evaluation, LANL Archives; William J. Worlton, "Cray-1 Preliminary Evaluation Results," May 17, 1976, Records of C-DO, box 4, folder 2, (April-May, 1976), Preliminary Cray-1 Evaluation Plan, Performance Statistics, LANL Archives; T. W. Keller, "Cray-1 Evaluation: Final Report," Los Alamos Scientific Laboratory, LA-6456-MS (December 1976), 1-8.

workload that might appear in any scientific-computing setting, while the second consisted of LASL codes meant to represent computing throughout ERDA.³⁹⁰

The final results were mostly excellent, in favor of the new computer, but with one significant problem. As the May preliminary report indicated, the Cray-1 greatly exceeded the capabilities of the CDC 7600, particularly on heavily vectorized applications. On long-vector problems, with datasets in excess of 500 numbers, the new computer was over five times as fast as the 7600, while still boasting at least twice the older supercomputer's scalar performance. With hardware and software optimization, the differences would likely become even more significant. The Cray-1's performance, and most of its other capabilities, exceeded all of the minimum requirements LASL and ERDA stipulated for a Class VI production computer. The problem was the system's unacceptably high error rate. The majority of errors occurred in main memory.³⁹¹

The Cray-1 was Seymour Cray's first computer design to incorporate semiconductor-based memory, rather than the older core-memory technology used in his previous supercomputers. Across the 910 total hours in which the computer operated during the evaluation, the computer experienced 152 memory parity errors, meaning that a stored value was different when read than when it was originally written in memory. Those errors introduced inaccuracies that could invalidate results, and could halt the job being processed, potentially requiring operator intervention, and running parts or all of a

³⁹⁰ Don Bryson, "Letter from Don Bryson to Ray Anastos, Sperry UNIVAC," Records of C-DO, box 3, Cray-1, folder 11, Evaluation Schedule, LANL Archives; William J. Worlton, "Cray-1 Preliminary Evaluation Results"; Keller, "Cray-1 Evaluation: Final Report," 1-8.

³⁹¹ Keller, "Cray-1 Evaluation: Final Report," 3-4.

job again. A high error rate rendered the Cray-1 effectively unusable for LASL.

Unknown at the time, the Los Alamos environment, itself, caused the errors.³⁹²

Although not confirmed until 2010, at over 7,000 feet above sea level, the high altitude of Los Alamos Scientific Laboratory allowed atmospheric neutrons—at a level six times higher than at Cray's Chippewa Falls Lab in Wisconsin—to alter the supercomputer's memory. Atmospheric neutrons easily traveled through the minimal shielding provided by the CCF's thin concrete roof. The neutrons caused "single event upsets," or SEUs, when they encountered the computer's memory chips, resulting in a "bit flip," when a one or zero in memory would "flip" to its opposite. Bit flips resulted in data errors when the affected memory segments were read again by the computer. The larger the memory, the greater the chance of errors, an important issue, considering the Cray-1 would ship with a memory size of up to one million, 64-bit words. The Cray-1's predecessors at LASL would also have experienced effects from atmospheric neutrons, but research using the neutron-beam facility at Los Alamos in 2010 demonstrated that the type of semiconductor-based memory used on the Cray computer, sourced from Fairchild Semiconductor, was particularly susceptible to SEUs. Without built-in error-correcting capability, the Cray-1 could not recover from the SEU-induced errors.³⁹³

When C Division shared its results with Cray Research, the company halted its construction of Serial No. 2, a clone of Serial No. 1 at LASL. Although unaware of the cause of the memory errors, CRI's engineers modified the computer's design to

³⁹² Normand, et al., "First Record of Single-Event Upset on Ground, Cray-1 Computer at Los Alamos in 1976," 3114-3115.

³⁹³ Normand, et al., "First Record of Single-Event Upset on Ground, Cray-1 Computer at Los Alamos in 1976," 3114-3116.

incorporate an error correction scheme called SEC-DED, or single-error correction, double-error detection, which would detect when a single bit in a word had changed, and correct the error automatically. Customers could choose the error-correcting memory feature as an option when leasing or purchasing a Cray-1. Starting with Serial No. 3, Cray-1 models equipped with SEC-DED were four-inches taller than those lacking the feature. LASL had hoped to purchase Serial No. 1 after the evaluation, but its lack of error-correcting memory meant that it had to be returned to Chippewa Falls, as it could not meet ERDA's minimum reliability requirements. Despite the unexpected need to return the evaluation model, the modified Cray-1 design incorporating SEC-DED not only surpassed the performance of what little competition existed on the supercomputer market, but also proved to be more reliable than the CDC 7600.³⁹⁴

When LASL presented its case before ERDA to return the evaluation Cray-1 and purchase a version with error-correcting memory, the Sandia Lab in Albuquerque stepped forward to aid in LASL's case. LASL and Sandia proposed to share access to a LASL-sited Cray-1 via a remote link, which would allow the Albuquerque lab to meet most of its expected increase in computing demand until 1982 without the need to install new local systems, offsetting a portion of the Cray-1 purchase price. The joint-lab cooperative proposal reinforced LASL's justification for the Cray-1 to ERDA, and before Serial No. 1 had departed New Mexico, LASL and ERDA had already signed a contract

³⁹⁴ Normand, et al., "First Record of Single-Event Upset on Ground, Cray-1 Computer at Los Alamos in 1976," 3114-3116; Margaret Loftus, OH 316, Oral history interview by Judy E. O'Neill (18 and 25 April, 1995), Minneapolis, Minnesota. Charles Babbage Institute, University of Minnesota, Minneapolis. <http://hdl.handle.net/11299/107444>; Marge McCormick, ed., "Computing Requirements of the ERDA Advanced Computing Laboratories, July, 1977," 3-5; Dorr, "Computer Science and Services Division Activities and Plans, April, 1978," 7.

for Serial No. 4 to be delivered in 1977. Because of LASL's assistance in identifying a potentially fatal flaw in the Cray-1's design, CRI discounted the final purchase price of Serial No. 4 by \$400,000.³⁹⁵

Los Alamos conducted the first evaluation of the Cray-1 and signed a contract for delivery of Serial No. 4, but the National Center for Atmospheric Research (NCAR) in Boulder, Colorado became Cray Research's first paying customer when it ordered a Cray-1 during the summer of 1976. NCAR, an NSF-funded center established in 1960 to support research relating to the Earth's atmosphere, was a growing supercomputer user in the 1970s. NCAR experienced a rapidly increasing demand for computing resources in the early 1970s, and was quickly growing beyond its CDC 7600. Like Los Alamos, NCAR was dissatisfied with the supercomputing options on the market prior to the announcement of the Cray-1. Unlike Los Alamos in the summer of 1976, NCAR was financially and organizationally positioned to commit to ordering a Cray supercomputer without a full evaluation being completed first. Fortunately for NCAR, the incorporation of SEC-DED starting with Serial No. 3, the machine NCAR received, prevented the memory issues that had occurred at Los Alamos from happening at the, likewise, relatively high altitude of Boulder, Colorado.³⁹⁶

The NCAR purchase provided the financially precarious Cray Research with much-needed funds to remain in business even before the Los Alamos evaluation and

³⁹⁵ Ronald J. Detry, "Letter from Ronald J. Detry, Sandia Laboratories, to Dieter Fuss, Lawrence Livermore Laboratory, June 16, 1977," Records of C-DO, box 1, folder 6, (1976-1977) Cray-1, LANL Archives; McCormick, ed., "Computing Requirements of the ERDA Advanced Computing Laboratories, July, 1977," 3; Dorr, "Computer Science and Services Division Activities and Plans, 1978," 2, 7.

³⁹⁶ "CRI Cray-1A S/N 3," Computational & Information Systems Lab. Accessed 4/12/19. <https://www2.cisl.ucar.edu/supercomputer/c1>; Dorr, "Computer Science and Services Division Activities and Plans, April, 1978," 7-8.

potential ERDA purchase were complete. An issue for CRI was NCAR's unwillingness to purchase the computer without a minimal set of software utilities being included, such as a full operating system and an automatic vectorizing FORTRAN compiler. This spurred Cray to reconsider what he had told Los Alamos before the evaluation, when he expressed little interest in developing a full OS or much of a software package for the new machine. Unlike the weapons labs, NCAR and many other relative newcomers to supercomputing lacked the capability to develop most or all of their software in-house, making software the key issue for the new computer's success. While the LASL evaluation was still underway, Cray formed a software-development group at CRI in order to meet his contractual agreement with NCAR. With the computer's most significant hardware issue identified and resolved due to the LASL evaluation, and CRI altering its software strategy to suit an expanding variety of supercomputing centers, the Cray-1 was poised to become both a resounding financial success and an icon that would bring Seymour Cray's company to the forefront of the supercomputing market for years to come.³⁹⁷

Although Los Alamos could not transition from the Cray-1 evaluation to production use of the same computer as hoped, signing the contract to take delivery of Serial No. 4 in the fall of 1977 gave LASL a focus for its short- and long-term software preparations. No longer merely hoping to incorporate a Cray-1 into the CCF, LASL made concrete plans for supporting and operating the new supercomputer, which meant starting the difficult work of modifying the Lab's code base for the new vector-based system and the

³⁹⁷ O'Neill interview with Loftus; McCormick, ed., "Computing Requirements of the ERDA Advanced Computing Laboratories, July, 1977," 2-5.

development of its operating system. While CRI had committed to developing a fully realized Cray Operating System (COS) to fulfill its NCAR contract, C Division and its users had little expectation of the new OS being adequate for LASL's needs and moved forward with plans for a local OS.³⁹⁸

The final section of this chapter focuses on the negotiation between C Division's supporters and TD Division over LASL's operating system strategy following the order of Serial No. 4. Much of that battle will hinge on the difficulty of developing an OS strategy capable of striking a balance between preserving the Cray-1's system performance for production use, while also providing the diverse services demanded of a modern OS. The services most central to the negotiation process were directly tied to the availability and technical means of accessing LASL's custom supercomputer data-storage and remote-access network, Hydra, which would evolve into the Integrated Computer Network (ICN). Before moving on to the final section of this chapter, because of the centrality of computer networking in shaping LASL computer users' perceptions of what services were most essential from a supercomputer OS in the late 1970s, it is first necessary to clarify how *networking* assumed an influential role in guiding LASL's OS and computer-science R&D strategies.

While the origins of Los Alamos supercomputer networking appeared in Chapter 4, the following section will provide a brief overview of the rapid evolution of networking technologies at Los Alamos after the completion of Hydra's initial configuration in 1974, and the remarkable speed with which users took up computer networking.

³⁹⁸ Dorr, "Computer Science and Services Division Activities and Plans, 1978," 7-8.

Supercomputer usage norms at the Lab quickly reoriented around shared data-storage, remote-access, and interactive capabilities made possible largely through R&D work custom to Los Alamos' C Division. As a consequence, understanding the difficulties and drives behind integrating the Cray-1 into LASL's computing operation requires a firm grounding in the rising importance of supercomputer networking at Los Alamos, and how networking concurrently shaped C-Division strategies and the demands of supercomputer users.³⁹⁹

V. The Route Less Taken

As discussed in Chapter 4, LASL's first supercomputer network, Hydra, entered development in 1970, following the success of the Multiple User eXperiment (MUX) remote-access system. Hydra entered regular use at LASL in 1974, one year later than planned. Ron Christman, Hydra's technical manager, argued that the delay was largely due to an overly ambitious early design, and insufficient coordination within the development team for such a complex system, all common issues in early networking projects. The relative simplicity and small scale of MUX had not prepared the C Division networking team for the challenges of interconnecting multiple supercomputers and their shared storage systems using largely custom hardware and software. The resulting network was, by necessity, greatly scaled-back from the original plan, with proposed features, such as graphical interface capabilities, being stripped from the final design. The highly centralized Hydra network initially connected each of LASL's four

³⁹⁹ Dorr, "Computer Science and Services Division Activities and Plans, April, 1978," 7-8.

CDC 7600 supercomputers to a shared pool of short-term, long-term, and archival storage systems through a central front-end machine, itself a Control Data 6600 supercomputer. The front-end machine, doing no production work of its own, also managed remote terminal connections to the four 7600s.⁴⁰⁰

With Hydra suffering from early reliability issues, C Division began to modify the centralized network soon after it first entered service. A CDC Cyber 73 replaced the 6600 as the front-end machine in 1975, when the funding became available. Because so few high-capacity commercial networking solutions existed in the early 1970s, minicomputers became the Swiss Army knives of Los Alamos networking, and a major focus of C Division's R&D efforts. For Los Alamos, supercomputer networking had, in many respects, become what electronic, general-purpose computing had been in the early 1950s: a still-maturing field that required both a strong justification and a considerable investment in local expertise to make the new technologies operational and useable for early adopters. C Division customized UNIX-powered DEC PDP-11 minicomputers, and later the DEC VAX, and minicomputer models from Systems Engineering Laboratories (SEL) and Gould, to fulfill a variety of networking functions where vendor solutions were either inadequate or missing altogether. For example, in order to connect remote users to Hydra, C Division developed the KCC (Keyboard Communications Concentrator), a custom device based on a PDP-11 and DEC interface unit to provide a single connection to the Hydra front-end machine for up to 96 keyboard-based terminals in the classified environment. The CCC (Computer Communications Concentrator)

⁴⁰⁰ Christman, "A Review of the Hydra Project"; "The Computer Sciences and Services Division Annual Report, January—December 1975," 5.

provided a similar connection for up to 18 computer-based terminals, which were usually other DEC minicomputers. Despite Hydra's early teething issues, user adoption rates skyrocketed during the network's first year of operation, primarily as a consequence of the pent-up demand for remote access to the CCF for classified users.⁴⁰¹

In 1975, when the number of simultaneous users demanding keyboard-based terminal access already exceeded the capacity of a single KCC, C Division developed the PDP-11-based Synchronous Concentrator (SYNC). The SYNC served as a concentrator and message switcher to interconnect the KCCs with the Hydra front-end machine. More importantly, the SYNC provided direct terminal connections with each of the networked supercomputers in the CCF. This innovation reduced the centralization of the Hydra system, which C Division recognized as an unreliable network model. The CDC 6600's channel couplers proved to be the single most unreliable component of Hydra, as they were never intended to connect more than two CDC computers together.

Troubleshooting a connection error required maintenance staff to take both the originating machine and the Hydra front-end computer offline. Based on its first-hand networking experiences in the mid-1970s, C Division's networking groups, quickly growing in size and importance, developed a plan to ameliorate the existing network's most significant shortcomings. The first goal was to reduce further LASL's dependence on Hydra, seeking to decentralize the network in stages. Another critical goal was to

⁴⁰¹ "The Computer Science and Services Division Annual Report, January—December 1975," 5; Fred W. Dorr, "Computer Science and Services Division of the Los Alamos Scientific Laboratory," Los Alamos, NM: Los Alamos Scientific Laboratory, LA-UR-74-1566, (1974), 2; Leo Romero and Bill Buzbee, "Overview of the Los Alamos Integrated Computing Network," Los Alamos, NM: Los Alamos Scientific Laboratory, LALP-85-34 (1985), 23-24.

increase overall network throughput speeds to support supercomputers with faster I/O capabilities, such as the incoming Cray-1. Finally, the rapidly growing popularity of the Lab's existing networked services pressed C Division to expand the network well beyond its original scope. The network had to accommodate a variety of new services and users, including the addition of unclassified and administrative users. The plan meant adding new transport, switching, and security systems, all of which would have to be custom to Los Alamos.⁴⁰²

One of the first products of C Division's new networking plan was the Los Alamos File Transport machine, introduced in 1977. While the SYNCs provided complete terminal access to the supercomputers of the CCF independent of Hydra, Los Alamos lacked similar direct connectivity for file access. The File Transport (FT) machine used the same architecture and low-level protocols as the SYNC to provide high-speed data links of up to 40Mbit/s between the CCF supercomputers. The first FT linked the Cray-1, which arrived in the spring of 1977, and two CDC 7600s, which served as the Cray's front-end machines. (The purpose and consequences of this technical arrangement for accessing the Cray-1 will be discussed later.) Multiple FT machines, when linked together, formed the File-Transport Network (FTN), a high-speed backbone for direct file transfers across the network, which reduced the need for routing data through the Hydra front-end machine. The FTN also provided the much-needed infrastructure to support the

⁴⁰² Romero and Buzbee, "Overview of the Los Alamos Integrated Computing Network," 23; "The Computer Science and Services Division Annual Report, January—December 1975," 5-6; Lewis interview with Christman; Fred W. Dorr, "Computer Science and Services Division CCF Five-Year Plan," Los Alamos Scientific Laboratory, LA-6921-MS, (1977), 3-4, 13, 43-45, 79; Fred W. Dorr, "Computer Science and Services Division Activities and Plans, April, 1978" Los Alamos Scientific Laboratory, LA-7093, (1978), 4, 12, 19.

growing demand at LASL for high-speed graphics terminals, a technology that C Division, its users, and collaborating partners in industry, such as the Evans and Sutherland company in Salt Lake City, began to explore in the early 1970s as a means of providing supercomputer-driven visual representations of simulation data on a remote terminal. A minicomputer-based terminal would capture simulation data streamed from a supercomputer, then convert that data into images and even video for the user. The graphical terminals entering use in the late 1970s terminals, once connected to the FTN, allowed users to stream color graphical data at 150 kbits per second (compared to the 9.6 kbits per second of conventional CRT-based terminals) from the supercomputers of the CCF directly to their offices.⁴⁰³

Increasing the range of services and users on the Integrated Computer Network (ICN), as the network was known following its decentralization from the Hydra storage-network platform, in the late 1970s required a significant reconfiguration of network *security*. The original Hydra network's approach to security was simple; only classified users were allowed access. After the rapid adoption of network use by classified users, C Division faced a growing demand to accommodate unclassified supercomputer users, and to facilitate administrative Lab functions via networked systems, but creating separate networks for the *different* security levels would have been cost prohibitive. Instead, C Division developed a custom, minicomputer-based Network Security Controller (NSC) to monitor and restrict file and resource access based upon user passwords.⁴⁰⁴

⁴⁰³ Romero and Buzbee, "Overview of the Los Alamos Integrated Computing Network," 24; Lewis interview with Morrison.

⁴⁰⁴ Romero and Buzbee, "Overview of the Los Alamos Integrated Computing Network," 24.

The NSCs separated the ICN into three partitions: Secure, Administrative, and Open. Several of the older supercomputers in the CCF were moved to the Open partition for unclassified use, while the Administrative partition users had two dedicated CDC Cyber 73s, replacing the 1960s-vintage IBM 1401 mainframes previously dedicated to LASL administrative use. Users of the Administrative or Open partitions could share files with users on the Secure partition, but not the other way around, in order to prevent a "write down" of classified data to a lower classification level. Each data packet ran through two NSC checks, providing redundancy to guard against error. The NSCs did not run any user code, and prevented the need to rely upon production computers to enforce security rules. Using the NSC to monitor and control network security on the single physical network was far less costly than developing discrete networks, and allowed classified and unclassified users to share the same resources.⁴⁰⁵

By 1977, the ICN had surpassed 2,200 users. Most were local, but the number of remote users from Sandia and other labs sharing resources began to grow as well. As networked resources became more widely used and more mission-critical to LASL, the demands of connecting larger numbers of machines and services to the ICN grew more difficult and time consuming. The various supercomputers, minicomputers, storage systems, printing services, and other devices employed at Los Alamos were never intended to function together on a single network. A variety of incompatible interconnect technologies and proprietary communication protocols meant that C Division had to

⁴⁰⁵ Romero and Buzbee, "Overview of the Los Alamos Integrated Computing Network," 24; Robert H. Ewald, et al., "Computing Division Two-Year Operational Plan, FY 1981-1982," Los Alamos, NM: Los Alamos National Laboratory (1981), 41-42.

develop expensive hardware and software conversion systems that could translate the native I/O of one machine into that of every other device with which it needed to connect on the network. This engineering task grew increasingly impractical as the network became larger and more heterogeneous in the late 1970s. C Division's response was to standardize communications across the ICN, marking one of the first efforts toward network-interface standardization in the supercomputing arena.⁴⁰⁶

Beginning development in 1978, the Los Alamos High-Speed Parallel Interface (HSPI) converted the native I/O of a computer or another networked device into the standard ICN higher-level protocol suite, which was also developed at Los Alamos. The suite primarily consisted of SIMP (Simple Inter-Machine Protocol), PTP (process-to-process), and AFT (File Transport), which would remain in use for well over a decade, much like the HSPI. The custom protocols were designed to accommodate large data blocks with low latency, high burst rates, extensive error detection and correction, and minimal message passing. C Division had to design and construct a HSPI, including the hardware and drivers, for each new device placed on the network, but each HSPI only had to convert the native I/O interface and protocols of the new device into the standard ICN protocols. The HSPI of each receiving device then converted the ICN protocols back to that device's native I/O channels. Once connected to its HSPI, a device had full

⁴⁰⁶ Dorr, "Computer Science and Services Division Activities and Plans," 5; Ewald "Computing Division Two-Year Operational Plan, FY 1981-1982," 36; Romero and Buzbee, "Overview of the Los Alamos Integrated Computing Network," 43; Don Tolmie, oral history interview by Nicholas Lewis, Los Alamos, NM (July 14, 2014); Andrew and David Dubois, oral history interview by Nicholas Lewis, Los Alamos, NM (July 17, 2014).

access to the entire Los Alamos network, which significantly reduced the time required to integrate new equipment onto the ICN.⁴⁰⁷

HSPIs provided reliable, full-duplex point-to-point transfers to the various machines of the CCF, implementing extensive error correction integrated into the HSPI hardware, and facilitating data transfers of up to 60 Mbits per second. The fast HSPI data-rate was intended to exceed the I/O capabilities of the fastest devices on the network, so that the network interfaces would not act as bottlenecks. In order to accommodate different data rates (for example, 3 Mbits per second for the PDP-11, and 50 Mbits per second for the Cray-1), a handshake between HSPIs determined the maximum data rates of the exchanging machines. The HSPI employed hardware error detection and correction while data was on the fly across the network, so as not to slow the data transfer. When double-bit errors were detected, the HSPI hardware aborted the transfer. The first completed HSPI was for the ubiquitous PDP-11, with others following in short succession.⁴⁰⁸

The late 1970s introduction of higher networking speeds, new security partitions, and standardized interfaces and protocols on the ICN coincided with the retirement of Hydra in 1979, completing the Lab's transition to a modular network configuration. Replacing Hydra as the networked storage system on the ICN was the far more sophisticated, but still custom, Common File System (CFS). C Division used off-the-shelf IBM storage

⁴⁰⁷ Donald E. Tolmie, et al., "Interconnecting Computers with the High-Speed Parallel Interface," Los Alamos, NM: Los Alamos National Laboratory, LA-9503-MS, (1982), 2-12; Ewald "Computing Division Two-Year Operational Plan, FY 1981-1982," 43; Granville Chorn, et al., "The Standard File Transport Protocol," Los Alamos, NM: Los Alamos Scientific Laboratory, LA-7388-MS, (1978).

⁴⁰⁸ Norman R. Morse and Joseph L. Thompson, "Bringing A Large Computer Network Into Focus," paper presented at the 6th International Conference on Software Engineering, Tokyo, Japan, (September 13-16, 1982), 8-11; Tolmie, "Interconnecting Computers with the High-Speed Parallel Interface," 1-5.

hardware and operating systems, with the tape-cartridge-based IBM 3850 replacing the gigantic IBM Photostore, which the company no longer supported. LASL had to develop the CFS application and interface software entirely in-house to accommodate the complex storage demands of the Los Alamos network, which included users on all three security partitions. While CFS still required a central storage controller, unlike the repurposed CDC front-end computers of Hydra, the CFS controller was designed and used specifically for storage, offering a maximum data transfer rate of about 5 Mbits per second, and storing about 384 gigabytes of combined online and offline data on the network as of 1980, with a growth rate of about 175 gigabytes per year.⁴⁰⁹

After the introduction of CFS, the C Division leader at the time, Robert Ewald, noted a gratifying drop in the number of tapes at the Central Computing Facility, as the storage system reduced the number of redundant, partially filled tapes in circulation. Much of the CFS development team originally worked on the Hydra project, and had feared a repeat of the design difficulties that had delayed the rollout of Hydra. As a consequence, the team leaders embraced a "structured" process of allocating and reviewing programming work. Individual programmers on the project were assigned specific parts of the CFS coding effort, with their code being checked at least three times by other programmers for errors and efficiency before the code was added to the CFS application suite. The

⁴⁰⁹ Bill Collins, Marjorie Devaney, and David Kitts, "Profiles in Mass Storage: A Tale of Two Systems," *Ninth IEEE Symposium on Mass Storage Systems, 1988. Storage Systems: Perspectives* (Oct. 31, 1988 – Nov. 3, 1988), 61-67; Bill Collins, interview by Nicholas Lewis, Minneapolis, MN (2/19/2015).

process itself was time-consuming, but team members recalling the effort cited it as the key reason for the on-time, smooth roll-out of CFS, and its long-term reliability.⁴¹⁰

After the initial success of CFS at Los Alamos, in which the storage system received letters of praise from many of its users, General Atomics Corporation sold CFS commercially as DataTree, bringing CFS into use at several dozen data-intensive computing sites worldwide over the next decade. CFS during the 1980s represented a transition period in networked data storage, when a growing number of computing sites began to approach the networked storage requirements that Los Alamos first experienced in the early-to-mid 1970s. The success of CFS beyond the Lab indicated the continued maturation of a particular arena of networking technology, as the demand for high-end storage solutions grew outside Los Alamos. CFS also marked the important role that Los Alamos would play as a leader in networked storage development, collaborating with a growing range of partners in industry on high-speed and high-capacity storage systems, such as HPSS and Panasas, through to the present.⁴¹¹

⁴¹⁰ Lewis interview with Christman; Ewald, "Computing Division Two-Year Operational Plan, FY 1981-1982," 37.

⁴¹¹ M. Blood, R. Christman, and B. Collins, "Experience with the LASL Common File System," *Fourth IEEE Symposium on Mass Storage Systems*, New York: The Institute of Electrical and Electronics Engineers, Inc., (1980); Bob Ewald, "Overview," *User Interface* 67 (May 1981): 3-4; Lewis interview with Collins; Gary Grider, interview by Nicholas Lewis, Los Alamos, NM (8/21/2015); Fred McClain, "DataTree and UniTree: Software for File and Storage Management," *Digest of Papers. Tenth IEEE Symposium on Mass Storage Systems: Crisis in Mass Storage* (1990): 126-28. Lawrence Livermore National Laboratory also developed its own networked-storage system, which sold commercially as UniTree. As the name indicates, UniTree ran atop the Unix operating system, while Los Alamos' DataTree/CFS relied on IBM's MVS OS. While MVS was a more mature platform in the late 1970s, when both networked-storage solutions were under development, it left DataTree/CFS tied to IBM operating system development and support at a time when the national labs and supercomputing in general were moving toward Unix in the late 1980s and 1990s.

VI. DEMOS and the End of the Beginning.

The growth and increasing importance of supercomputer networking at Los Alamos directly influenced the OS development decisions and considerations surrounding the Cray-1, even before Serial No. 4 arrived, and was an important backdrop to the evolution of computer use at LASL during the pivotal events of the late 1970s (discussed in this section) that permanently reshaped the distinctive Los Alamos approach to computing. Part of the influence of networking capabilities on the LASL Cray-1 OS strategy was predictable, with users, such as Paul Whalen, understandably desiring interactive, remote-terminal access to the Cray-1, and the ability to move files between the new computer and Hydra's data-storage devices, services that were already in place for the CDC 7600s of the CCF. Another influence of LASL networking on the Cray-1 OS was less expected, and involved C Division and LASL management addressing the possibility that C Division had drifted from a key aspect of its original R&D mission. With its increasingly heavy focus on practical networking development projects, C Division no longer hosted the diverse range of basic computer science R&D programs expected of it in the division charter. The concern over LASL's local R&D capabilities were in connection with a larger ERDA initiative to support ERDA-wide computer management efforts in order to reduce the duplication of development work across the ERDA labs. As a consequence, C Division's OS strategy for the Cray-1 developed out of the division's responsibility to provide user-demanded services, and an ERDA initiative to rejuvenate a perceived loss of basic computer science research capacity at LASL. When both ERDA and LASL's administrative structures once again

underwent significant restructuring, managerial support for C Division's OS strategy evaporated, leaving C Division vulnerable to its most vocal, and increasingly influential critics.

During the Cray Review Board meetings of 1975, the proposal to develop a version of CROS for the Cray-1, possibly linked to a mainframe or larger front-end machine running KRONOS, had gained support among most LASL computer users, with TD Division being the noted exception. By late 1976, the Cray-1's architectural details and system capabilities, which the Computing Division and ERDA had experienced firsthand, spurred C Division to revisit C-4's suggestion of developing an entirely new operating system, rather than adapting an existing OS for the Cray. Jack Worlton and John Morrison, group leader of C-11 System Software, presented a technical argument stating that no existing operating system could adequately exploit the Cray-1 architecture. Porting LTSS, CROS, KRONOS, UNIX, or others, would fail to optimise the new supercomputer's performance, because they were not designed specifically for the Cray. Out of necessity, C Division planned to run the Cray Operating System (COS) temporarily on the new computer. While COS *was* designed specifically for the Cray-1, Cray Research was developing the OS with a wide range of customers in mind, meaning that, as C Division's software experts predicted, it would not be well tailored, or easily modified, to meet LASL's requirements. Worlton, Morrison, others in C Division, and members of the user community, particularly from T Division, argued that LASL had an opportunity to develop an OS from scratch. Creating a new, custom OS would foster the best use of the Cray-1, would garner ERDA-level support and possible funding, given the

administration's interest in promoting software development and cutting-edge computer science research among its labs, and would bring LASL to the forefront of supercomputer OS design.⁴¹²

While C Division was completing the evaluation, and had begun to prepare for the arrival of Serial No. 4, ERDA began an initiative meant to maximize the benefits of supercomputer usage among its labs, focusing heavily on software design and development. Convening the Committee on Advanced Computing, ERDA's headquarters and field offices met to discuss the unique computing problems facing its labs. The unique problems included the frequent need of the ERDA labs to develop their own software and other computing resources due to their "specialized computing environments," referring to the conjoined technical and security issues. The labs, particularly LASL and Livermore, often had to tailor their operations and technical approaches to unusual research requirements, but security restrictions severely curtailed how the labs shared their resources and software, and limited how outside resources and expertise could be brought into the secure environments. The committee concluded that, out of necessity, the labs would have to continue producing a large portion of their own software in-house if they were to make the best use of state-of-the-art computers that were entering the market. After the LASL evaluation, ERDA management had strong hopes of the Cray-1 becoming widely used in the lab system, but wanted to avoid software fragmentation issues similar to what had occurred under the AEC's leadership. The committee determined that ERDA needed to support these in-house software

⁴¹² Lewis interview with Morrison.

development efforts at the highest level. With better high-level support and the promotion of ERDA-wide cooperation, the labs could also avoid duplication in the design and development of software resources. The conclusions, and resulting initiatives, of the Committee on Advanced Computing had two almost immediate implications for Los Alamos computing and its policies.⁴¹³

The first implication for Los Alamos was a review of the Lab's computer science R&D capabilities, as part of a series of reviews being conducted. This occurred outside the usual C-Division annual review, in which independent experts from outside LASL assessed the overall performance and capabilities of C Division in meeting its responsibilities, and would make suggestions for improvement where necessary. The review instigated by the Committee on Advanced Computing focused specifically on C Division's computer science research activities and capabilities. The review concluded that C Division had invested less and less effort into basic computer science R&D over the preceding several years. Fred Dorr's decision in early 1976 to cancel the MANIAC II R&D project as of 1977 reflected that shift in investment. The emphasis on basic research, such as projects pushing the state-of-the-art in OS and programming-language development, had declined, while practical development efforts, such as the production of networking hardware and software devices for immediate LASL requirements, had assumed the lion's share of funding and attention within C Division.⁴¹⁴

⁴¹³ "Committee for Advanced Computing Meeting, October 12, 1976," C Division Papers, box 1, folder 6, (1976-1977) Cray-1, LANL Archives; M. C. Greer, "Committee on Advanced Computing, November 12, 1976," C Division Papers, box 1, folder 6, (1976-1977) Cray-1, LANL Archives;

⁴¹⁴ James C. Brown, et al., "C-Division Computer Science Research Review Board Report, January 17, 1977," Edward A. Voorhees Papers, box 58, Computers and Computing, C-Division Memos and Reports, folder 4, Reports, 1975-April 1977, series 42, LANL Archives; Fred Dorr, "Policy on MANIAC," March 8,

The shift in C Division's focus was partly in response to pressure from weapons program users who were opposed to LASL conducting "computer science" R&D of any sort. The strained relationship particularly with TD Division following the 1973 reorganization of C-Division management made some staff members in the computing operation disinclined to stoke further animosity by pursuing research projects that lacked some immediate application. Adding to the pressure against basic research, in May 1976, Robert Thorn replaced Duncan MacDougal as the Associate Director for Weapons, placing C Division's most outspoken critic in charge of the associate directorship responsible for the entire weapons program. Fred Dorr maintained close communications and a good working relationship with Thorn to avoid issues similar to the 1975 operating-system "fracas," but few in C Division could ignore the past issues with Thorn, or the current tensions and the potential for future unrest his position of power caused within the computing operation.⁴¹⁵

While the diversity of, and level of support for, C-Division basic research was on the decline, the growing lab-wide demand for more and better supercomputer networking capabilities heavily contributed to an emphasis on practical R&D efforts. Not all CCF users or members of C Division were pleased with the division's overall direction. In addition to the transition away from basic CS research, C Division began a scheduled phase-out of CROS from the CCF, to be complete by 1978. Organizational pressure from TD Division, now with support from the new Associate Director for Weapons, Robert

1976, Edward A. Voorhees Papers, box 57, Computers and Computing, C-Division Meetings, folder 3, C-Division Meetings, 1976-1978, series 42, LANL Archives.

⁴¹⁵ Brown, et al., "C-Division Computer Science Research Review Board Report, January 17, 1977"; Norman Morse, oral history interview by Nicholas Lewis, Los Alamos, NM (July 30, 2018).

Thorn, to transition the CDC 7600s of the CCF entirely to LTSS coincided with the logistical problem of C Division potentially having to provide support for three custom operating systems, CROS, LTSS, and the new Cray-1 OS, which would have strained C Division's support staff and funding. TD had also increased its use of LTSS significantly since the OS "fracas" in 1975. These changing dynamics left C Division with fewer technical, financial, and political reasons for retaining CROS, while moving more heavily toward LTSS offered some distinct benefits.⁴¹⁶

Appealing to TD users, the solidifying OS plan in late 1976 (discussed below) raised the possibility of using front-end machines running LTSS to provide interactive Cray-1 access, while still equipping the Cray-1 itself with a custom LASL operating system that would better utilize the Cray's architecture. Many of the CCF users who opposed LTSS were worried about the forced shift away from CROS, but changing administrative power dynamics left little recourse. Members of C Division who were displeased with division's new directions instigated an attempted "coup" in September 1976, when several members of C Division approached Director Agnew to have Fred Dorr replaced as division leader. The attempt failed, but it left Dorr demoralized in his demanding leadership position. Dorr was being pulled in multiple directions simultaneously, with TD Division and Robert Thorn largely opposed to CS research at Los Alamos, but tolerating practical R&D where it immediately benefited the weapons program, while ERDA sought to bolster C Division's basic CS research capabilities,

⁴¹⁶ Brown, et al., "C-Division Computer Science Research Review Board Report, January 17, 1977"; C. M. Gillespie, ADWA, to R. N. Thorn, "The Future of CROS at LASL, December 7, 1976," Edward A. Voorhees Papers, box 58, Computers and Computing, C-Division Memos, folder 2, C-Division Memos, 1976-1978, series 42, LANL Archives.

particularly in operating-system development, the most contentious issue in Los Alamos computing since the origins of C Division. The added complication of disaffected C-Division personnel bypassing their division leader to seek his ouster ultimately contributed to Dorr's resignation in 1978.⁴¹⁷

The ERDA review board determined that C Division needed to augment its basic computer science research activities, particularly in software development, and adhere better to the original C-Division charter from 1968, which stipulated the maintenance of a heavy program in basic and applied CS research. The board's concern was that, with the continuing trend in C Division toward applied research, LASL was at risk of being caught unprepared for future developments in computing technologies, particularly if Cold War political or technical demands changed quickly. Of particular concern, the board noted, was that several group leaders within C Division did not believe that research in any form, basic or applied, had value for the division or the Lab. The board believed that developing worthwhile basic and applied research projects within the division would "overcome" this lack of commitment to R&D. The overall goal of the review was to enhance the local R&D capabilities of all ERDA labs. LASL, with its strong history of computing R&D, was among the most vital links in forming an ERDA-wide software collaboration effort, as the Committee for Advanced Computing sought to create.⁴¹⁸

⁴¹⁷ Brown, et al., "C-Division Computer Science Research Review Board Report, January 17, 1977"; J. Worlton, C-DO, to H. Agnew/R. Thorn, "C-Division Management Issues, March 13, 1978," cabinet 10-2 (Jack Worlton), HPC Archive; Gillespie, to Thorn, "The Future of CROS at LASL, December 7, 1976."

⁴¹⁸ Brown, et al., "C-Division Computer Science Research Review Board Report, January 17, 1977."

The review board noted one bright spot of combined basic and applied CS research in C Division: the proposed development of a Cray-1 operating system, called DEMOS, short for Deus Ex Machina Operating System. The Committee for Advanced Computing, in anticipation that the Cray-1 would enter widespread ERDA use, requested that the ERDA labs submit proposals for the development of an ERDA-wide OS for the new supercomputer. The goal was to avoid OS fragmentation, as happened with LTSS and CROS on the CDC 7600, to enable better software sharing across the lab system, and to foster cutting-edge OS development, ensuring that ERDA maintained a capable R&D capacity for supercomputer software moving forward. The DEMOS concept at Los Alamos grew out of C Division's leadership and the C-11 System Software group revisiting the idea of developing an all-new OS for the Cray-1. The concept was bold, featuring a high-speed file system new to DEMOS, an unusual "link-based" messaging system (to be discussed later in detail), and the development of a high-level programming language, called MODEL, specifically for programming the OS. While the computer science research review of LASL was in progress, C Division, with approval Fred Dorr secured from Robert Thorn and Harold Agnew (thanks, in part, to the promise of ERDA funding), proposed the DEMOS concept to be the ERDA-wide Cray-1 OS.⁴¹⁹

The Committee for Advanced Computing accepted the DEMOS proposal, and C Division began to meet with the ERDA Advanced Computing Working Group to

⁴¹⁹ Brown, et al., "C-Division Computer Science Research Review Board Report, January 17, 1977"; Lewis interview with Morrison; Forest Baskett, John H. Howard, and John T. Montague, "Task Communication in DEMOS," in *Proceedings of Sixth ACM Symposium on Operating Systems Principles* (November, 1977), 23-31; Michael L. Powell, "The DEMOS File System," *Proceedings of Sixth ACM Symposium on Operating Systems Principles* (November 1977), 33-42; F. W. Dorr, C-DO, to R. N. Thorn, ADW, "Cray-1 Operating System Proposal, March 2, 1977," Records of C-DO, box 1, Cray-1, folder 6 (1976-1977) Cray-1, LANL Archives.

collaborate on the OS design. The working group was created to find practical means of implementing the Committee for Advanced Computing's software initiative, focusing on the development of an ERDA-wide operating systems for Class VI supercomputers. As the primary developer, LASL would receive ERDA funding to act as the official OS support center. The DEMOS design, as it matured in 1977, reflected both LASL and ERDA's interests in a Cray-1 OS, with the project residing organizationally within C Division's C-4 and C-11 groups. C-11 group leader, John Morrison, also became the DEMOS project leader. Most of DEMOS' advanced features aimed at pushing the limits of cutting-edge OS design, with a heavy focus on optimizing data and job throughput via innovative programming and design choices. The use of a high-level language, MODEL, for most of the programming work was intended to make development faster, while also making the OS easier to maintain, and more flexible across multiple ERDA sites.⁴²⁰

An important, early design consideration, which ERDA mandated in its selection of the OS for its lab system, was that DEMOS would *not* be a time-sharing system. The decision to make DEMOS a batch-only OS grew out of C Division and the Committee for Advanced Computing determining that the Cray-1 marked a fundamental shift in supercomputer architectures, and that it should be treated differently from preceding machines. While the Cray-1 evaluation was underway, and C Division continued to explore its options for an OS for the new machine, the C-4 group brought to Jack Worlton's attention a growing belief in the supercomputing community that time-sharing

⁴²⁰ John Morrison, C-11, to J. Worlton, N. Morse, and F. Baskett, "Meeting with ERDA Advanced Computing Working Group, June 22, 1977," Records of C-DO, box 1, Cray-1, folder 6 (1976-1977) Cray-1, LANL Archives; Dorr, to Thorn, "Cray-1 Operating System Proposal, March 2, 1977"; Lewis interview with Morrison.

on supercomputers was becoming an outmoded concept. The argument was that relatively inexpensive, networked minicomputers could easily provide local computing and time-sharing services to users, leaving supercomputers free for number crunching.⁴²¹

The Cray-1's exceptional performance and design offered a chance to break away from previous time-sharing usage models, and use supercomputers more effectively for their intended purpose. Offloading time-sharing use to minicomputers appealed to the Advanced Computing Committee's interest in optimizing supercomputer use across the ERDA labs, resulting in the ERDA committee stipulating that DEMOS be a batch OS. Because the Cray-1 required at least one front-end machine, the DEMOS plan, similar to the KRONOS-CROS link C Division proposed in 1975, would allow remote, interactive capabilities with the Cray-1 via a front-end computer running LTSS, which appealed to TD-Division users. The design plan called for the development of DEMOS interfaces for other front-end operating systems, such as UNIX (for minicomputers), over time.⁴²²

With the DEMOS development plan and ERDA collaboration (along with ERDA funding) in place, C-4 and C-11 prepared for an official project starting date of October 1977, allowing C Division time to recruit new personnel needed for the programming effort. While the project leader, John Morrison, was an existing LASL staff member, most of the twenty people devoted to the DEMOS project were new hires, consisting largely of recent computer science graduates. Many of the new hires came from Stanford, and were eager to work on a state-of-the-art operating system for the Cray-1. Much like the original MANIAC and other local computing-development projects at

⁴²¹ Lewis interview with Morrison; Dorr, "Cray-1 Operating System Proposal, March 2, 1977."

⁴²² Lewis interview with Morrison; Dorr, "Cray-1 Operating System Proposal, March 2, 1977."

LASL, the DEMOS project was an effective recruiting tool for augmenting the Lab's expertise in a cutting-edge technical field. Among the new hires for DEMOS was Forest Baskett, a computer scientist and former Stanford professor, who contributed heavily to the innovative DEMOS message-passing system. Jobs and programs, called "tasks" on DEMOS, communicated with one another and the OS kernel through "links," which identified where in memory a task existed, what a task was requesting via the link, and what permissions a task had on the system. The purpose of link-based message-passing communication was to minimize the system calls being sent to and from tasks, as tasks could only communicate with other tasks and parts of the OS for which they had links. The use of links also protected the system memory (a problem on the Cray-1, which had little hardware-based memory protection), as the kernel was the only part of the operating system that required hardware access to the main memory.⁴²³

One of Baskett's former students, Michael Powell, was a key developer of the DEMOS file system, which had the difficult task of maximizing I/O performance when the use environments for the operating system across the ERDA labs would be difficult to predict, and would vary widely among the different sites, and even among the types of jobs at each lab. With the difficulty of building from scratch a state-of-the-art OS for a new supercomputer architecture, particularly one that had to meet the requirements of both LASL and ERDA, DEMOS was not originally scheduled to enter early production use until 1979, and would not be feature-complete until 1980. Because the new

⁴²³ Lewis interview with Morrison; Dorr, "Cray-1 Operating System Proposal, March 2, 1977"; Basket, et al., "Task Communication in DEMOS," 23-26; Harry McGavran, oral history interview by Nicholas Lewis, White Rock, NM (July 17, 2014); Paul Iwanchuk, oral history interview by Nicholas Lewis, Los Alamos, NM (July 7, 2014).

computer's time was too valuable to remain idle while DEMOS was under development, C Division created a Cray-1 link for LTSS, which entered production use in January 1977. Users submitted jobs via the Integrated Computer Network (ICN) to one of the two CDC 7600s working as the Cray-1 front-end machines. A FORTRAN cross-compiler then executed under LTSS on the 7600, generating code for the Cray Operating System (COS) on the Cray-1. DEMOS was not to begin user testing on the Cray-1 until 1979.⁴²⁴

The major DEMOS project landmarks were scheduled to take place on a quarterly basis. The resident file system was scheduled to be available in the first quarter of 1978, the multitasking system in the second quarter, and the network interface, which would allow any production computer on the ICN to connect to the Cray-1, not simply the front-end machines, was scheduled for the fourth quarter. Just as the DEMOS project began development work in late 1977, a political reorganization of the ERDA system significantly altered the pressures on LASL's weapons program, and, consequently, on the DEMOS project.⁴²⁵

In late 1977, the passage of the Department of Energy Act consolidated the functions of ERDA and several other administrations, agencies, and programs under the umbrella of the Department of Energy (DOE). The federal effort to contend with the energy crisis, which had prompted the elimination of the AEC and the creation of ERDA,

⁴²⁴ Powell, "The DEMOS File System," 33-35; Fred Dorr, "Computer Science and Services Division Activities and Plans, April, 1978," Los Alamos, NM: Los Alamos Scientific Laboratory, LA-7093, (1978), 8-9.

⁴²⁵ Dorr, "Computer Science and Services Division CCF Five-Year Plan," 29; Dorr, "Computer Science and Services Division Activities and Plans, April, 1978," 8-9.

reached its climax with the formation of the DOE. The DOE inherited ERDA's oversight of the US nuclear-weapons labs. New administrators and priorities within the department brought changing directives for the labs, which would quickly worsen the relationship between LASL's weapons program and the Computing Division, but for the time being, the DOE continued to sponsor DEMOS as the DOE-wide operating system for the Cray-1.⁴²⁶

In the first quarter of 1978, the modular components of DEMOS began to take shape on schedule, and the project's relationship with users and the office of the Associate Director for Weapons were generally good. TD Division had been by far the largest user of the Cray-1 since the installation of the LTSS link, using between 80% and 90% of the computer's time each month. TD-Division users looked forward to an OS more suited to large LASL codes than the barebones COS from Cray Research. The limitations of COS meant that C Division closely controlled which codes could be run on the Cray-1. As the effects of ERDA's reorganization began to be felt at Los Alamos, administrative pressures on the Lab's weapons program, and, consequently, on the DEMOS project, began to change. The formation of the DOE had brought heightened oversight of LASL's weapons program, and increasingly tight production timetables for its weapons designers, which directly translated into enhanced pressure on C Division and DEMOS' developers.⁴²⁷

⁴²⁶ Westwick, *The National Labs*, 269-270; Department of Energy Organization Act, Public Law 95-91, 95th Congress (Aug. 4, 1977): 91 STAT. 565-579; Lewis interview with Morrison.

⁴²⁷ J. Worlton, "Cray-1 Meeting with Thorn and Agnew, 2-7-78," Cabinet 10-2 (Jack Worlton), Reading File, Nov. 1977-Dec. 1978, HPC Archive; Dorr, "Computer Science and Services Division Activities and Plans, April, 1978," 8-9; Lewis interview with Morrison; Worlton, "C-Division Management Issues, March 13, 1978"; "Cray-1 Schedule and Production Guidelines," In *Central Computing Facility Newsletter*, vol. 23 (September 1977), 12.

Heightened demands on LASL's weapons program from the administrative structure, in large part, drove TD Division's heavy use of the Cray-1, which greatly surpassed what TD or C Division had predicted. The supercomputers in the CCF were generally operating at full capacity in the late 1970s, primarily from TD's escalating requirements. The added computing demand spurred C Division to perform a \$1 million upgrade of the Cray-1's memory, and to begin making its case to the DOE and Congress to purchase a second Cray-1 for the CCF. Himself under pressure from DOE management and TD-Division computer users, the Associate Director for Weapons, Robert Thorn, began to press John Morrison, leader of the DEMOS project, for the OS to enter production use as quickly as possible, because the limitations of COS were hampering TD's timetables. The Cray Operating System lacked many of the features and capabilities users had grown to expect from the familiar, purpose-built LTSS. The use of two CDC 7600s as the Cray-1's front-end machines while DEMOS was under development also cost the CCF significant production capacity. Once DEMOS was more feature-complete, the 7600s could be replaced with a DEC minicomputer as its sole front-end machine. Morrison recalled the meetings with Thorn about DEMOS becoming increasingly unpleasant by the spring of 1978, but there was only so much Morrison could do with an OS that was still in such an early stage of development. Displeased with DEMOS not being ready for production use, and with the worsening state of computer access in the CCF, Thorn intervened in C-Division management.⁴²⁸

⁴²⁸ Dorr, "Computer Science and Services Division Activities and Plans, April, 1978," 8-9; Lewis interview with Morrison; Lewis interview with Mike Simmons; Lewis interview with McGavran; Worlton, "C-Division Management Issues, March 13, 1978."

When Fred Dorr resigned as C-Division leader in April 1978, Thorn appointed Kaye Lathrop in Dorr's place. A graduate of West Point who had, notoriously, ousted Bengt Carlson as leader of T-1 in 1972, Lathrop was appointed to "clean up" C Division. As part of that directive, Thorn and Lathrop pushed for accelerated availability of DEMOS whether the OS was ready or not. Originally scheduled for "friendly" user testing not to begin until 1979, early DEMOS production work began in May of 1978, in conjunction with the publication of the DEMOS User's Guide. Because of its missing networking features, users still had to utilize the two CDC 7600 front-end machines to access DEMOS on the Cray-1, just as they had with COS. Users logged onto DEMOS via LTSS on one of the 7600s, which compiled user code for DEMOS, and handled the data exchange between the 7600s and the Cray.⁴²⁹

Early production tests did not go well, largely due to missing features, and problems with DEMOS I/O performance. Writing the OS entirely in the high-level language MODEL had resulted in inefficient hard disk use, and other data-handling issues. More extensive testing with "friendly" users began in October, after more features were available, and after significant debugging had taken place. Michael Powell's team began rewriting parts of the file-system code in assembly language to correct the hard disk and other I/O performance issues, but the problem would not be fully corrected until 1979, after the second phase of testing. Although the OS was still incomplete, and

⁴²⁹ Thomas R. Hill, oral history interview by Nicholas Lewis, Los Alamos, NM (25 April, 2019); Worlton, "C-Division Management Issues, March 13, 1978"; Group C-2, "DEMOS User's Guide," Los Alamos Scientific Laboratory, LA-5525-M (May 1978), ix.

development was moving far more quickly than originally intended, early users were already forming and becoming set in their views of DEMOS.⁴³⁰

According to former weapons designers and other members of TD Division, much of the issue with DEMOS was its very different user experience compared with the more familiar LTSS. On LTSS, for example, users could assign their own job priority with the computer's scheduler based on an honor system. The DEMOS scheduler was more strict, and required users to abide by a formal justification procedure to prevent gaming the system, where users would place their job at the highest priority on the scheduler, whether that was justified or not. The interface was also markedly different from LTSS, with users becoming frustrated at having to learn how to use a new system in the midst of an increasingly stressful production environment from the DOE. Finally, many TD computer users who preferred LTSS resented the return to a batch-only OS, even if the increasing prevalence of minicomputers in their offices and on the Los Alamos network offset many users' former time-sharing requirements. Just as DEMOS entered early use at LASL, and its features began to mature under testing, the DOE's operating system priorities shifted away from DEMOS.⁴³¹

In October 1978, at the same time DEMOS entered its second phase of user testing at LASL, the DOE Scientific Computing Subcommittee (SCS) began to reevaluate the department's support for DEMOS. When Livermore was finally able to order its own

⁴³⁰ Group C-2, "DEMOS User's Guide," Los Alamos Scientific Laboratory, LA-5525-M (May 1978), ix; Dorr, "Computer Science and Services Division Activities and Plans, April, 1978," 8-12.

⁴³¹ Lewis interview with Hill; Lewis interview with McGavran; Lewis interview with Mike Simmons; J. Worlton, C-DO, to K. Lathrop, C-DO, "A Decision Method for the Cray-1 Operating System, 10-28-78," Cabinet 10-2 (Jack Worlton), Reading File, Nov. 1977-Dec. 1978, HPC Archive.

Cray-1, scheduled for delivery in 1979, the California-based lab decided to convert LTSS for use on the Cray, calling it the Cray Time-Sharing system, or CTSS. While CTSS was not yet ready for production use, because it was a conversion of an existing operating system, it would likely be complete before DEMOS, which had started from scratch. The Scientific Computing Subcommittee argued for a shift in DOE support from DEMOS to CTSS, spurring a debate amongst the labs over which OS they preferred. Jack Worlton wrote to Kaye Lathrop in response to the SCS' argument for CTSS, which Worlton described as "not only inadequate, but flippant." Worlton suggested that Lathrop and LASL management take great care before committing to any one course of action based on the SCS proposal.⁴³²

Worlton warned that canceling DEMOS and moving to CTSS would offer a short-term production benefit, but a long-term penalty in performance and capability, as CTSS would never be as well-optimized for the Cray architecture as DEMOS. For Los Alamos, specifically, the cancelation of DEMOS would have enormous consequences for morale, particularly if it meant importing another Livermore-sourced OS. In addition, canceling DEMOS in favor of CTSS would indicate to the DOE that Los Alamos was not capable of meeting its own OS requirements, and could result in LASL's newly hired programming talent departing for Livermore. Worlton's more dire warning invoked the recurring discussions that had taken place in the US Congress of shutting down one of the weapons labs, which, based on its geographical isolation, would likely be Los Alamos. Those discussions had never resulted in firm action, but they recurred about once per

⁴³² Lewis interview with Clifford; Worlton, to Lathrop, "A Decision Method for the Cray-1 Operating System, 10-28-78."

decade. Worlton cautioned that the cancelation of DEMOS could, if the possibility of closing one of the labs did become serious in Congress, be a tick in favor of Livermore, with its recurring successes in OS development.⁴³³

The arguments Worlton and others at LASL made in favor of DEMOS were sufficiently persuasive to continue director support for DEMOS, and the support of many non-TD computer users who were looking forward to the operating system's completion. Political concerns aside, the financial and labor investments involved with developing DEMOS and the conversion of LASL codes to run on it were substantial. Abandoning those investments because an inferior operating system would be completed marginally sooner than LASL's state-of-the-art OS would have been difficult to justify, particularly considering the benefits DEMOS had brought in expanding and improving the Lab's R&D capacity in a cutting-edge field of computing. Outside Los Alamos, the SCS debate went in favor of CTSS—and DOE support for DEMOS ended. Livermore had always resisted the idea of importing a batch-only OS from Los Alamos, and had pushed for the DOE to adopt CTSS as its preferred operating system. Despite the change in the DOE's position, Los Alamos remained at least temporarily committed to DEMOS as its Cray-1 operating system. While DEMOS would not be the DOE-wide OS as once had been hoped, it seemed to have a future at Los Alamos.⁴³⁴

Between January and September of 1979, DEMOS evolved significantly. The publication of the DEMOS Primer in January guided users through the new OS, including a walkthrough for vectorizing their FORTRAN programs. The Cray-1 had also

⁴³³ Worlton to Lathrop, "A Decision Method for the Cray-1 Operating System, 10-28-78."

⁴³⁴ Worlton, to Lathrop, "A Decision Method for the Cray-1 Operating System, 10-28-78."

transitioned completely from COS to DEMOS as its primary operating system. As the I/O performance, reliability, and user-interface issues were resolved, and the LASL codes written or adapted for DEMOS were optimized, the system's total downtime and Mean Time Between Failures (MTBF) improved considerably over the summer, besting LTSS on the CDC 7600s by a wide margin. As the OS became more stable and reliable, and as more codes were ported to the operating system, more users transitioned to DEMOS on the Cray-1. The second LASL Cray-1, scheduled to arrive in November 1979, was slated to run DEMOS exclusively from the beginning, while the procurement of a third Cray-1 was already underway to meet the heightened user demand. The marked improvement of DEMOS' reliability and use statistics over 1979 indicated that the new OS was on track to become a cutting-edge workhorse at Los Alamos by 1980, when it would become feature-complete, including built-in networking capability that would allow the Cray-1 to shed its LTSS front-end machines. While the official reports on DEMOS offered optimistic appraisals of the operating system's future, heightened TD-Division user dissatisfaction with C Division, in conjunction with administrative intervention at the Lab, were about to change that outlook.⁴³⁵

In March of 1979, Harold Agnew resigned as Los Alamos director, and Robert Thorn assumed the role of Acting Director during the search for Agnew's replacement. Although Thorn was typically antagonistic toward C Division's independence, and was displeased with the prolonged development time needed for DEMOS, Thorn did not

⁴³⁵ "DEMOS Primer," Los Alamos Scientific Laboratory, LA-7499-M (January 1979), 1-4; "Plans and Progress," in *C Division News: Computer Science & Services*, vol. 47 (September 1979): 2, 5-7; "Plans and Progress," *C Division News: Computer Science & Services*, vol. 48 (October 1979): 2-3.

make significant changes to LASL's administrative structure or its computing policies during his tenure as Acting Director. The protocol for acting directors was to allow the incoming director to assess the Lab's organization, and make their own changes as needed. This holding pattern in LASL administration is likely what provided C Division time to continue its development of DEMOS without high-level intervention, despite a growing user discord from the weapons program, which will be discussed shortly. In July, Donald Kerr, an expert and former group leader in Los Alamos weapons testing, became the incoming director, and began organizational changes that would take effect when Kerr formally assumed his role as director in the fall. One significant organizational change was the formation of eight new associate directorships, effective in August. As a consequence, Kaye Lathrop stepped down as C-Division leader to become the Associate Director of the new Engineering Sciences Directorate. The new director appointed Robert "Bo" Ewald, former group leader of C-7 Systems and Language Support, and briefly the assistant division leader, as Lathrop's replacement. Ewald, who sought to inspire the best from his colleagues, brought a markedly different leadership style to C Division after Lathrop, who was often described as lacking in personal skills. In his first months as division leader, Ewald strongly promoted DEMOS as the future of LASL computing, changing its name in October to "ALAMOS," which stood for "A Los Alamos Modular Operating System."⁴³⁶

⁴³⁶ Lewis interview with Hill; "C Division - Computer Sciences and Services, Org Chart, 1979," Edward A. Voorhees Papers, box 1, folder 1, Organization Charts - C-Division, series 1, LANL Archives; The gym Ewald founded for Los Alamos personnel, now known as "Bo's Gym," is still open at the Lab today.

While C Division's optimistic official reports on the continuing refinement of ALAMOS accurately reflected the technical improvements and status of the OS as it reached its final stages of development, they did not reflect the turmoil stemming from TD-Division users as their discontent with the OS escalated over 1979. Because TD Division was the primary user of the Cray-1 since its arrival, most users who transitioned to DEMOS/ALAMOS from COS over 1979 were from TD, the LASL division that had most strongly argued for LTSS on the Cray from 1975 onward. Complaints about the new OS grew in relation to the number of new users who had to make the transition from the familiar interface and time-sharing capabilities of LTSS to a still-incomplete batch-only system, while still being required by the weapons program to meet their strict production schedules.⁴³⁷

C Division and DEMOS/ALAMOS project members worked to resolve the remaining technical and interface issues of the OS, while still working long hours to complete the operating system over the next year. However, the one intractable problem was DEMOS/ALAMOS, itself. TD-Division users largely did not want the new, unfamiliar operating system, and became increasingly vocal in their resentment of it. They wanted an easy transition to something familiar, even if it meant sacrificing performance with an OS that was less-well optimized than LASL's homegrown operating system. The argument from weapons program users was, because they were the primary financial supporters of Los Alamos computing, the Lab should simply buy more cycles to

⁴³⁷ Robert Ewald, "Overview," *C Division News: Computer Science & Services*, vol. 49 (November 1979): 3; Lewis interview with Morrison.

replace those lost to OS inefficiencies, if it meant the weapons program was able to meet its production deadlines.⁴³⁸

DEMOS/ALAMOS project members recalled feeling embattled, as TD users became more vocal in their opposition to the OS. There was little that members of the project could do to ameliorate the situation, as powerful users were opposed to the operating system itself. The directorship transition, and the lack of any other OS options for the Cray-1 at that moment in time, meant that TD-Division users and C Division were at a mutual impasse between the spring and fall of 1979. Neither group could effectively contend with the problems the other had with the new operating system. C Division could only continue developing and improving the OS it had been assigned to create two years earlier, and TD users could only complain to C Division and LASL management about the operating system they did not want.⁴³⁹

A further shift in the Lab's power balance toward the weapons program finally turned the tense situation into an open conflict. Donald Kerr's organizational changes dissolved TD Division, and consolidated the four most compute-intensive groups from TD, along with a design group each from T, the Laser Division, and the Field Testing (J) Division, to form the new Applied Theoretical Physics (X) Division. X Division became an even larger user of LASL computing than its predecessor, thanks to the intensive requirements of its constituent groups. While the formation of X Division further concentrated CCF funding and weapons-program users within a single, politically

⁴³⁸ Lewis interview with William J. Worlton; Lewis interview with Clifford; Lewis interview with Mike Simmons; Lewis interview with McGavran.

⁴³⁹ Lewis interview with Clifford; Lewis interview with McGavran.

powerful division, Kerr reestablished the position of Deputy Director, after that post had been abolished in 1974, with Robert Thorn becoming the new Deputy Director. As LASL's weapons program grew in power, Livermore began early testing of CTSS. Although still incomplete, the modified version of LTSS was on track to becoming production-ready in the next year, and provided not only an alternative to ALAMOS, but the same environment former TD users had wanted all along.⁴⁴⁰

In October and November 1979, while X Division and the Deputy Director's Office solidified in their new positions, C Division continued with its plans to finalize ALAMOS' development over the next year, and to install the new OS on the second Cray-1 when it arrived at the end of November. The Deputy Director, now in a place to respond to X-Division dissatisfaction with C Division and ALAMOS, called an emergency meeting over the 1979 Thanksgiving holiday break. Mostly X-Division members were in attendance at the meeting, along with some token representatives of C Division and the ALAMOS project. The new director, Donald Kerr, occasionally appeared during the proceedings to provide cookies and coffee as the group debated ALAMOS' future, or its lack of one. The assembled group overwhelmingly voted to end the ALAMOS project. The attending representatives from the ALAMOS project, including Jim Clifford, tired of frequent verbal altercations with members of the weapons program, likewise voted to end their own project. The meeting ultimately decided that ALAMOS would be phased out on the Cray-1 over the first several months of 1980,

⁴⁴⁰ "LANL Org Chart, 1943-1984," LASL, LANL, Org. Charts, HPC Archive; Lewis interview with Morrison; David Forslund, oral history interview by Nicholas Lewis, Los Alamos, NM (August 2, 2018); Lewis interview with Hill.

while C Division aided Livermore in completing the conversion of LTSS into CTSS. CTSS would be the Cray-1 OS at Los Alamos moving forward.⁴⁴¹

Although DEMOS/ALAMOS development ended at Los Alamos, the OS lived on outside the DOE. Michael Powell, one of the DEMOS file system developers, brought DEMOS with him to Berkeley in the mid-1980s, where he and his graduate students used the OS as a platform for distributed-computing research. Called DEMOS/MP (for Multi-Processor), the operating system's unusual link-based messaging system made it ideal for distributed computing, as the system kernel could easily keep track of where processes lived across multiple computers, because the links contained the data needed for identifying where processes were, which resources they needed, and where among the multiple, interconnected computers they needed to go to find them. The movement of links across the distributed system also provided an ideal means of monitoring the state of the DEMOS/MP system in real-time. Unlike supercomputer checkpointing today, which is only a snapshot of the status of a system at series of designated points, a DEMOS/MP system could be rolled back effectively to any previous state, just by restoring the links to that given moment.⁴⁴²

While the end of ALAMOS was a significant blow to C Division and its newly reclaimed computer science R&D capabilities, the Thanksgiving meeting was not the end of high-level Lab intervention in C-Division research. The administrative intervention in

⁴⁴¹ Ewald, "Overview (November 1979)," 3; Lewis interview with Clifford; Lewis interview with William J. Worlton; Lewis interview with Morrison; Robert Ewald, "Overview," in *C-Division News*, vol. 51 (January 1980): 3.

⁴⁴² Barton Miller, oral history interview by Nicholas Lewis, Los Alamos, NM (June 26, 2014). For a technical discussion of DEMOS/MP, see: Barton P. Miller, David L. Presotto, and Michael L. Powell, "DEMOS/MP: The Development of a Distributed Operating System," in *Software—Practice and Experience* 17, no. 4 (April 1987): 277-290.

LASL computing had only begun. In December 1979, before the C-Division newsletter had officially announced the end of ALAMOS, a former Livermore weapons designer, Robert Selden, became the first leader of X Division. As incoming X-Division leader, Selden had strongly supported the cancelation of ALAMOS and the import of CTSS. He also sought the imposition of a more Livermore-like relationship between LASL's weapons program and its computing operation, with the weapons program firmly in control of Lab computing. Beyond altering LASL's OS strategy, Selden and the Deputy Director, Robert Thorn, moved to reorganize C Division to eliminate its research and policy-making independence.⁴⁴³

Between December 1979 and February 1980, C Division underwent an almost total reorganization, with most of its research groups either being eliminated or fundamentally altered, and much of its staff being shuffled around the division. Two software R&D groups, including C-4 CCF Software Development and Support, which was in charge of the ALAMOS project, were dissolved entirely. John Morrison, leader of the ALAMOS project, was "sent off into the weeds," as one of his colleagues phrased it, and was denied another leadership position for several years afterward. Two new groups, the Systems Maintenance and User Interface groups, were formed in February, and, as their titles suggest, focused on computer maintenance and the user experience. Several groups had their R&D functions stripped away, while other groups had new functions added. Perhaps most telling of the forced restructuring of C Division and its operating philosophy was its renaming from the Computing Sciences and Services Division, to

⁴⁴³ Robert Ewald, "Overview," *C Division News: Computer Science & Services*, vol. 50 (December 1979): 3; Lewis interview with Clifford; Lewis interview with William J. Worlton; Lewis interview with Hill.

simply the "Computing Division," in January 1980. Other changes to C-Division groups and their leadership continued into early 1981, with many of the new group leaders being selected for their shared belief that Los Alamos should not be in the business of computer science research. While practical networking R&D continued largely as before, projects needed an X-Division sponsor before they could begin. As a culminating act of retribution, Robert Thorn issued an infamous memo, where he barred C-Division staff from attaining higher leadership positions at LASL.⁴⁴⁴

C Division became a fundamentally different organization at the start of the 1980s. Its leadership was primarily focused on maintaining and expanding user access. The division's remaining R&D efforts were largely directed at immediate problem solving, with less interest in LASL's standing or local capabilities in a particular field of computing, to say nothing of pushing computer science boundaries. LASL remained an important player in supercomputing, as C Division earned recognition among professional communities for its work in storage systems, network interconnects, and other home-grown technologies, but its work in fundamental computer-science research, like the MANIAC II and DEMOS, and its ability to plan its own development strategies, had effectively ended.⁴⁴⁵

The loss of LASL's inhouse computer science R&D capacity later proved nearly disastrous for the Lab at the end of the Cold War, when the US federal government and the Department of Energy began to promote the rapid advancement of computer-

⁴⁴⁴ Robert Ewald, "Overview," *C Division News: Computer Science & Services*, vol. 51 (January 1980), 3; Lewis interview with Clifford; Ann Hayes, oral history interview by Nicholas Lewis, Los Alamos, NM (August 3, 2018); Lewis interview with William J. Worlton.

⁴⁴⁵ Lewis interview with Forslund; Lewis interview with McGavran; Lewis interview with Clifford.

simulation capabilities in the US to contend with the decline and end of live nuclear testing. In order to bypass the weapons program's prohibition on C Division conducting its own basic computer science research, the division's computer benchmarking group (which developed an extremely small R&D capacity to aid in its benchmarking work) split off from C Division to form the Los Alamos Advanced Computing Laboratory (ACL), which redeveloped the Lab's computer science R&D tools, methods, and facilities between the late 1980s and late 1990s. The ACL's research program and methods were so successful in aiding with LANL's (Los Alamos National Laboratory's) transition into the Stockpile Stewardship era (in which the DOE weapons labs maintained the readiness of the nuclear stockpile, but without the aid of live testing) that the ACL was integrated back into the Lab's supercomputing operation, as its fundamental computer-science research operations had, once again, become part of LANL's standard computing practices by the early 2000s.⁴⁴⁶

VII. Chapter and Project Conclusion

The operating-system battles at Los Alamos during the 1970s were the top, technical layer of a deeper conflict over who would have the power to determine the future of Lab computing. The underlying points of contention stemmed from the long-established features of Los Alamos computing that C Division inherited from its

⁴⁴⁶ Tom Kitchens, "The US Department of Energy's "Grand Challenge" Program," in *The International Journal of High Performance Computing Applications* 4, no. 3 (September 1990): 3-5; Patrick Byrnes, et al, "Advanced Computing Laboratory," Los Alamos, NM: Los Alamos National Laboratory, LALP-91-66 (1991), 3; Tim Germann, oral history interview by Nicholas Lewis; Los Alamos, NM (August 22, 2018); David Daniel, oral history interview by Nicholas Lewis, Los Alamos, NM (July 5, 2018).

predecessor, and which C Division continued to practice to varying degrees until its dismantling through administrative intervention. The conflicts over CROS and DEMOS were the primary articulation points of resistance to C Division from the weapons program, but the diverging interpretations of who should control LASL's computing directions, and what those directions should be, were rooted in the unique development of the Los Alamos approach to computing decades earlier, and in a technical and political environment quite different from that of the 1970s.

The traditional Los Alamos approach to computing grew organically out of the early Cold War need for *users* of computing machines to play an active role in the development and refinement of the technologies they depended upon. As Los Alamos physicists, engineers, and other professionals sought new technologies and methods capable of resolving otherwise intractable mathematical and material-science problems underlying atomic- and thermonuclear-weapon design, they made Los Alamos a center for the promotion and advancement of the computational sciences and computer engineering. With few vendors or outside experts to draw upon in the incipient stages of the electronic, general-purpose computing field, the Los Alamos scientists working at the limit of existing mathematical modeling techniques, primarily in the Theoretical Physics Division, formed a highly concentrated collection of experts in scientific computing.

As the computing machines and mathematical models that T-Division experts depended upon evolved concurrently with one another, often far in advance of other scientific institutions, and even other divisions within Los Alamos, T Division scientist-managers placed computing on an equal footing with other key fields of scientific

research at the Lab. The thermonuclear-development effort of the early 1950s, stalled for lack of adequate computing power, conclusively demonstrated for T-Division scientists the centrality of computing technologies in LASL's mission. It also highlighted the need for Los Alamos to maintain a strong basic and applied research capacity in computing, because the Lab's ability to meet changing Cold War pressures could hinge upon the adaptability of the computing operation and its reserve of expertise.

The direct connection between T Division's active role in developing LASL's computing capabilities and the successful continuation of the Lab's central mission informed the Los Alamos approach to computing as both a science and as a service, and characterized the era of the expert user, in which personnel in need of machine-based computing resources had little choice but to become experts in those systems. As a consequence, the users and developers of LASL computing were often the same people, instilling among those expert users a strong belief in computing as a necessary field of Lab research. While Los Alamos was not isolated from outside centers of scientific computing, the classified nature of its work greatly restricted the Lab's freedom to engage openly with other sources of expertise. Security restrictions enhanced T Division's perceived need for self-reliance, and fostered a sense that "outsiders" were not in a position to understand the difficulties facing Lab computer users.

As general-purpose computing technologies and products from vendors matured, T Division refocused its local research and development efforts on new technical fields and methods that touched upon, or had the potential to touch upon, LASL's computing requirements. Relatedly, as Cold War political and strategic pressures on Los Alamos

enhanced or shifted the importance of the Lab's computing resources in new directions, T Division's chosen areas of computing research shifted to accommodate. These intertwined models of transitioning LASL's computer-research focus became a core component of T Division's computing strategy. While transitioning from local to vendor solutions served to reinvent the LASL computing operation and research focus over time, T Division's approach to computing reflected the broader Los Alamos conservatism toward its central mission. LASL's strong theoretical grounding in its weapons program, as opposed to the more rapidly iterative approach at Livermore, called for a heavy reliance on long, regimented rounds of computer-based modeling before a weapon's design would enter the physical-testing period. As a result, T Division was wary of embracing new technologies, methods, or organizational strategies that could jeopardize the stability of the computing operation and, by extension, LASL's mission.

The T Division approach to computing, including both its operating and research strategies, encountered little resistance at Los Alamos while T Division was both the provider and primary user of LASL's computing resources. As the key technologies behind scientific computing matured significantly over the 1960s, and as changing Cold War strategies brought new ranks of users to LASL's computing operation, T Division's inwardly focused development strategies came under scrutiny, spurring the formation of LASL's first dedicated computing division. Although the Computing Sciences and Services Division formed with the intent of easing past tensions over a powerful science division dominating LASL's computing resources and strategies, the traditional Los Alamos approach to computing found its way to the new division. The transition was

largely due to the almost complete transfer of T-Division computing personnel to C Division, with T Division's most accomplished expert users filling the ranks of C Division management.

In the early 1970s, although traditional batch computing comprised the majority of users, the rise of non-expert weapons program users, whose demands of and preferred strategies for the computing operation often differed from traditional users, became increasingly common at Los Alamos. This was due to a rapid expansion of the computing sciences at and beyond the Lab, diminishing returns on weapon-design efficiency improvements, and an intensification of administrative oversight for the AEC lab system. A growing cross-section of users who had little vested interest in the continuation of LASL's traditional approach to computing, which was no longer protected within a politically and financially powerful science division, brought C Division and its relative independence into repeated conflict with the assertive and powerful weapons program. Operating system development, an important point of basic and applied research at LASL due to the relative immaturity of vendor-sourced solutions, was a focal point of this conflict, because LASL's operating-system strategies directly shaped user experiences and capabilities with the Lab's supercomputers. As the OS priorities of conventional batch users and newer weapons-program users diverged, C Division was caught in the middle, attempting to satisfy the broadest range of Lab computing requirements with the resources available. C Division's retention of its predecessor's inward focus and often inadequate skills for communicating its plans with user divisions

intensified the discontent of users who disagreed with, or were skeptical of, C-Division policies.

The reorientation of LASL's power structure in the early 1970s favored newer, non-traditional computer users from the weapons program, and strengthened their demands for classified time-sharing in the computing facility. This led to the first top-level intervention in C-Division operating and R&D policies. While the managerial and OS policy changes resulting from the intervention did not completely undermine the continuation of C Division's traditional approach to computing, the C-Division emphasis on basic CS research began to decline in response to years of continuing pressure from weapons-program users, and the hefty resources required to develop the Lab's custom networking infrastructure. From the mid-1970s onward, many weapons-program users considered themselves underrepresented in C-Division policy decisions, in part due to the division's continued difficulties with communication, and in part due to the weapons program not fully recognizing the diverse user interests that C Division had to accommodate.

Despite the new-found administrative power of the weapons program to institute change in C Division's OS strategy in the early 1970s, and to pressure the division for a reduced emphasis on basic research through the mid-1970s, the Los Alamos weapons program still did not command the position of dominance over LASL computing that the Livermore weapons program possessed over that Lab's computing resources. LASL's greater diversification of research, and the long-standing history of the Los Alamos computing operation being largely autonomous, greatly curtailed the ability of key

weapons-program administrators to induce fundamental changes in the Lab's computing policies. The Cray-1, with its radical architecture, became a catalyst in the struggle to determine the future of LASL computing. Various factions, including C Division, LASL's traditional and non-traditional users, and ERDA, realized that the operating-system design that could first and successfully capture Cray-1 users in the ERDA lab system would shape the computing capabilities and directions of those labs for the foreseeable future.

With substantial high-level support, C Division in the late 1970s began to recapture its basic R&D capacity in OS development, first demonstrated with the Stretch operating system, and then with CROS and the Hydra OS, but the reprieve was only temporary. Before the end of the decade, an even more significant reorganization of LASL's power structure, following fast upon the end of ERDA support, brought C Division's most ambitious basic R&D software project, DEMOS, to an abrupt end. The same reorientation of power that ended DEMOS also ended the approach to computing that had made the Cray-1 OS possible. Reorganizing C Division's leadership, and stripping the division of its independence and its in-house basic research capacity, undercut the possibility of C Division continuing with the approach inherited from T Division over a decade earlier. The traditional Los Alamos approach to computing survived the transition to a new division, Lab-level reorganization, directed financial and political pressure, and marked changes in both the Cold War strategic and technological environments that drove Los Alamos computing. Its survival into the late 1970s reflected the strong commitment of T-Division staff and computing experts to ensure that Los

Alamos computing, and by association, the Lab's mission, would not fall prey to inadequacies of computing technology, methodology, or capabilities stemming from the rapidly changing requirements of the Cold War.

Bibliography

- "7600 Computer Now in Operation at Livermore Lab." *Magnet* (April 1969).
- Akera, Atsushi. *Calculating a Natural World: Scientists, Engineers, and Computers During the Rise of U. S. Cold War Research*. Cambridge, MA: MIT Press, 2007.
- _____. "Voluntarism and the Fruits of Collaboration: The IBM User Group, Share." *Technology and Culture* 42, no. 4 (2001): 710-736.
- Allred, John C., and Mary Mann, eds. *Behind Tall Fences: Stories and Experiences About Los Alamos At Its Beginning*. Los Alamos, New Mexico: Los Alamos Historical Society, 1996.
- Anderson, H. L. "Scientific Uses of the MANIAC." *Journal of Statistical Physics* 43, nos. 5/6 (1986): 731-748.
- _____. "Metropolis, Monte Carlo, and the MANIAC." *Los Alamos Science*, no. 14 (1986): 96-107.
- Armer, Paul. "SHARE—A Eulogy to a Cooperative Effort." *Annals of the History of Computing* 2, no. 2 (April 1980): 122-129.
- Aspray, William. *John von Neumann and the Origins of Modern Computing*. Cambridge, MA: MIT Press, 1990.
- Badash, Lawrence, Joseph O. Hirschfelder, and Herbert P. Broida, eds. *Reminiscences of Los Alamos: 1943-45*. Boston: USA, 1980.
- Bashe, Charles J., Lyle R. Johnson, John H. Palmer, and Emerson W. Pugh. *IBM's Early Computers*. Cambridge: The MIT Press, 1986.
- Baskett, Forest, John H. Howard, and John T. Montague. "Task Communication in DEMOS." *Proceedings of Sixth ACM Symposium on Operating System Principles* (November 1977): 23-31.
- Blood, M., R. Christman, and B. Collins. "Experience with the LASL Common File System." *Fourth IEEE Symposium on Mass Storage Systems*. New York: The Institute of Electrical and Electronics Engineers, Inc., 1980.

- Boghosian, Bruce. "Computational Physics on the Connection Machine." *Computers in Physics* 4, no. 14 (1990): 14-33.
- Bouricius, Willard G. "Operating Experience with the Los Alamos 701." *Proceedings of the 1953 Eastern Joint Computer Conference*. New York: ACM, 1953.
- Brockmeyer, John A., and A. Gene Dornhoff. "Computer Interfacing at LASL." Los Alamos Scientific Laboratory, LA-UR-76-2502, 1976.
- Burks, A. W., H. H. Goldstine, and J. von Neumann. *Preliminary Discussion of the Logical Design of an Electronic Computing Instrument*. Princeton, NJ: Institute for Advanced Study, Princeton, 1946.
- Campbell-Kelly, Martin, William Aspray, Nathan Ensmenger, and Jeffrey R. Yost. *Computer: A History of the Information Machine*, Third Edition. Boulder, CO: Westview Press, 2014.
- Carlson, Bengt G. "Solution of the Transport Equation by Sulfur Neutron Approximations." Los Alamos Scientific Laboratory, LA-01599, 1953.
- Carlson, B., C. Lee, and J. Worlton. "The DSN and TDC Neutron Transport Codes." Los Alamos Scientific Laboratory, LAMS-2346, 1959.
- Chambers, Marjorie Bell. "Technically Sweet Los Alamos: The Development of a Federally Sponsored Scientific Community." PhD diss., University of New Mexico, 1974.
- Chorn, Granville, Ronald Christman, and Thomas Klingner. "The Standard File Transport Protocol." Los Alamos, NM: Los Alamos Scientific Laboratory, LA-7388-MS, 1978.
- Christman, Ronald D. "MUX, a Simple Approach to On-Line Computing." *Communications of the ACM* 15, No. 5 (May 1972): 319-329.
- _____. "The LASL Computer Network." Los Alamos, NM: Los Alamos Scientific Laboratory, LA-5103-MS, December, 1972.

- Cochran, Thomas B. *Nuclear Weapons Databook: Volume II, U.S. Nuclear Warhead Production*. Ballinger Publishing Company: 1987.
- Collins, Bill, Marjorie Devaney, and David Kitts. "Profiles in Mass Storage: A Tale of Two Systems." *Ninth IEEE Symposium on Mass Storage Systems, 1988. Storage Systems: Perspectives* (Oct. 31, 1988 – Nov. 3, 1988).
- Dorr, Fred W. "Computer Science and Services Division Activities and Plans, April, 1978." Los Alamos Scientific Laboratory, LA-7093, 1978.
- _____. "Computer Science and Services Division CCF Five-Year Plan." Los Alamos Scientific Laboratory, LA-6921-MS, 1977.
- _____. "Computer Science and Services Division of the Los Alamos Scientific Laboratory." Los Alamos Scientific Laboratory, LA-UR-74-1566, 1974.
- Eckhardt, Roger. "Stan Ulam, John von Neumann, and the Monte Carlo Method." *Los Alamos Science* 15 (1987): 131-143.
- Edwards, Paul N. *The Closed World: Computers and the Politics of Discourse in Cold War America*. Cambridge, MA: MIT Press, 1986.
- Elzen, Boelie, and Donald MacKenzie. "The Social Limits of Speed: The Development and Use of Supercomputers." *IEEE Annals of the History of Computing* 16, no. 1 (1994): 46-61.
- Enslow, Philip, Jr. "Multiprocessor Organization—A Survey." *ACM Computing Surveys* 9, no. 1 (March 1977): 103-129.
- Feynman, Richard P. "Los Alamos From Below: Reminiscences of 1943-1945." *Engineering and Science* 39, no. 2 (1976): 11-30.
- Flamm, Kenneth. *Creating the Computer: Government, Industry, and High Technology*. Washington, DC: The Brookings Institution, 1988.
- Fitzpatrick, Anne. "Igniting the Light Elements: The Los Alamos Thermonuclear Weapon Project, 1942-1952." PhD diss., George Washington University, 1999.

- Furman, Nicah Stewart. *Sandia National Laboratories: The Postwar Decade*. Albuquerque: The University of New Mexico Press, 1990.
- Galison, Peter. "Computer Simulations and the Trading Zone." *The Disunity of Science: Context, Boundaries, Power*. eds. P. Galison and D. Stump. Stanford: Stanford University Press, 1996.
- Galison, Peter, and Barton J. Bernstein. "In Any Light: Scientists and the Decision to Build the Superbomb, 1952-1954." *Historical Studies in the Physical and Biological Sciences* 19, no. 2 (1989): 267-347.
- Grier, David Alan. *When Computers Were Human*. Princeton, NJ: Princeton University Press, 2005.
- Gusterson, Hugh. *Nuclear Rites: A Weapons Laboratory at the End of the Cold War*. Berkeley and Los Angeles: University of California Press, 1996.
- Haigh, Thomas, Mark Priestley, Crispin Rope. *ENIAC in Action: Making and Remaking the Modern Computer*. Cambridge, MA: The MIT Press, 2016.
- Haigh, Thomas, Mark Priestley, Crispin Rope. "Los Alamos Bets on ENIAC: Nuclear Monte Carlo Simulations, 1947-1948." *IEEE Annals of the History of Computing* 35, no. 3 (July-Sept. 2014): 42-63.
- Hall, Donald M. *Generation of Wealth: The Rise of Control Data and How It Inspired an Era of Innovation and Investment in the Upper Midwest*. Minneapolis, MN: Nodin Press, 2014.
- Hammel, Edward F. "Los Alamos Scientific Laboratory Energy-Related History, Research, Managerial Reorganization Proposals, Actions Taken, and Results, 1945-1979." Los Alamos National Laboratory, LA-13072-H, 1997.
- Harlow, Francis H., and N. Metropolis. "Computing & Computers: Weapons Simulation Leads to the Computer Era." *Los Alamos Science* 7, (1983): 132-141.
- Harlow, Frank. "J. Carson Mark." *Los Alamos National Laboratory Daily News Bulletin* (March 4, 1997).

- Hawkins, David. *Project Y: The Los Alamos Story, Part I, Toward Trinity*. San Francisco: Tomash Publishers, 1988.
- Hecht, David K. "The Atomic Hero: Robert Oppenheimer and the Making of Scientific Icons in the Early Cold War." *Technology and Culture* 49, no. 4 (2008): 943-966.
- Hewlett, Richard G., and Oscar E. Anderson, Jr. *The New World: A History of the Atomic Energy Commission, Volume I, 1939-1946*. University Park: Pennsylvania State University Press, 1962.
- Hewlett, Richard G., and Francis Duncan. *Atomic Shield: A History of the United States Atomic Energy Commission, Volume II, 1947-1952*. U.S. Atomic Energy Commission, 1972.
- Hewlett, Richard G., and Jack M. Holl. *Atoms for Peace and War, 1953-1961: Eisenhower and the Atomic Energy Commission*. Los Angeles: University of California Press, 1989.
- Hoddeson, Lillian, Paul W. Henriksen, Roger A. Meade, and Chatherine Westfall. *Critical Assembly: A Technical History of Los Alamos during The Oppenheimer Years, 1943-1945*. Cambridge [England]: Cambridge University Press, 1993.
- Holloway, David. *Stalin and the Bomb: the Soviet Union and Atomic Energy, 1939-1956*. New Haven: Yale University Press, 1994.
- Hughes, Thomas P. *Networks of Power: Electrification in Western Society, 1880-1930*. Baltimore, MD: John Hopkins University Press, 1983.
- Hunner, Jon. *Inventing Los Alamos: The Growth of an Atomic Community*. Norman: University of Oklahoma Press, 2008.
- Hurd, Cuthbert C. "A Note on Early Monte Carlo Computations and Scientific Meetings." *Annals of the History of Computing* 7, no. 2 (1985): 141-155.
- _____. "Early IBM Computers: Edited Testimony." Jeffrey R. Yost, ed. *The IBM Century: Creating the IT Revolution*, IEEE Computer Society: CS Press, 2011: 71-89.

- Johnson, Kenneth J., ed. "Stretch, the Super Computer is Gone." *The Atom* 8, no. 6 (July-August, 1971): 23-26.
- Kaufmann, William J., III, and Larry L. Smarr. *Supercomputing and the Transformation of Science*. New York: Scientific American Library, 1993.
- Keller, T. W. "Cray-1 Evaluation: Final Report." Los Alamos Scientific Laboratory, LA-6456-MS, 1976.
- Kitchens, Tom. "The US Department of Energy's "Grand Challenge" Program." *The International Journal of High-Performance Computing Applications* 4, no. 3 (September 1990).
- Knutsen, Eric. *Wasn't That a Time: Stretch/Harvest Retrospectives*. Harvest/Stretch Reunion, September 28 and 29, 2002, Poughkeepsie, NY. Riverside, CT: LookBack Publishing, 2002.
- Lazarus, Roger B., Edward A. Voorhees, Mark B. Wells, and W. Jack Worlton. "Computing at LASL in the 1940s and 1950s." Los Alamos, NM: Los Alamos Scientific Laboratory, LA-6943-H, 1978.
- Lazarus, Roger, Mark Wells, and John Wooten, Jr. "The Maniac II System." *Symposium on Interactive Systems for Experimental Applied Mathematics: Proceedings of the Association for Computing Machinery, Inc.* New York: ACM, 1967.
- Lecuyer, Christophe. *Making Silicon Valley: Innovation and the Growth of High Tech, 1930-1970*. Cambridge, MA: MIT Press, 2006.
- Lincoln, Neil R. "Supercomputer Development—The Pleasure and the Pain." *Datamation* 23 (May 1977).
- MacKenzie, Donald. *Inventing Accuracy: A Historical Sociology of Nuclear Missile Guidance*. Cambridge, MA: MIT Press, 1990.
- _____. "Nuclear Weapons Laboratories and the Development of Supercomputing." In Mackenzie, Donald, ed. *Knowing Machines: Essays on Technical Change*. Cambridge: The MIT Press, 1998: 99-129.

- _____. "The Influence of the Los Alamos and Livermore National Laboratories on the Development of Supercomputing." *IEEE Annals of the History of Computing* 13 (1991): 179-201.
- Mark, J. Carson. "A Short Account of Los Alamos Theoretical Work on Thermonuclear Weapons, 1946-1950." Los Alamos, NM: Los Alamos Scientific Laboratory, LA-5647-MS, 1974.
- Masco, Joseph. *The Nuclear Borderlands: The Manhattan Project in the Post-Cold War New Mexico*. Princeton University Press, 2006.
- McClain, Fred. "DataTree and UniTree: Software for File and Storage Management." *Crisis in Mass Storage. Digest of Mass Storage Systems, 1990* (1990): 126-128.
- McMurrin, Marshall William. "Achieving Accuracy: A Legacy of Computers and Missiles." Xlibris Corporation, 2008.
- Metropolis, N. *A History of Scientific Computing*. New York: ACM, 1990.
- _____. "The Age of Computing: A Personal Memoir." *Daedalus* 121, no. 1 (Winter, 1992): 119-130.
- _____. "The Beginning of the Monte Carlo Method." *Los Alamos Science*, no. 15 (1987): 125-130.
- _____. "The Los Alamos Experience, 1943-1954." N. Metropolis. *A History of Scientific Computing*. New York: ACM, 1990: 237-250.
- _____. "The MANIAC." In N. Metropolis, J. Howlett, and G. C. Rota, eds. *A History of Computing in the Twentieth Century*. New York: Academic Press, 1980: 457-464.
- Metropolis, N., and E. C. Nelson. "Early Computing at Los Alamos." *Annals of the History of Computing* 4, no. 4 (October 1982): 348-357.
- Metropolis, N., J. Howlett, and Gian-Carlo Rota, eds. *A History of Computing in the Twentieth Century: A Collection of Essays*. New York: Academic Press, 1980.

- Metropolis, N., and J. Worlton. "A Trilogy on Errors in the History of Computing." *First USA-Japan Computer Conference Proceedings: October 3-5, 1972, Tokyo, Japan*, American Federation of Information Processing Societies (1972): 683-691.
- Metropolis, N., and S. Ulam. "The Monte Carlo Method." *Journal of the American Statistical Association* 44, no. 247 (September 1949): 335-341.
- Metropolis, N., et al. "Equations of State Calculations by Fast Computing Machines." *Journal of Chemical Physics* 21, no. 6 (1953): 1087-1092.
- Metropolis, Nicholas, D.H. Sharp, W.J. Worlton, and K.R. Ames, eds. *Frontiers of Supercomputing*. Berkeley: University of California Press, 1986.
- Miller, Barton P., David L. Presotto, Michael L. Powell. "DEMOS/MP: The Development of a Distributed Operating System." *Software – Practice and Experience* 17, no. 4 (April 1987): 277-290.
- Misa, Thomas J. *Digital State: The Story of Minnesota's Computing Industry*. Minneapolis: University of Minnesota Press, 2013.
- _____. "Military Needs, Commercial Realities, and the Development of the Transistor, 1948-58." Merritt Roe Smith, ed. *Military Enterprise and Technological Change*. Cambridge, MA: MIT Press, 1985: 253-87.
- Misa, Thomas J., ed. *Building the Control Data Legacy: The Career of Robert M. Price*. Minneapolis: Charles Babbage Institute, 2012.
- Mitchell, Charlie. "C Division Plugs in CFS System Early." *LASL Newsletter* (June 27, 1979): 7.
- Morse, Norman R. "The Role of Supercomputers at LASL." Los Alamos, NM: Los Alamos Scientific Laboratory, LA-6750-MS, 1977.
- _____. "C-Division Annual Review and Operating Plan, August 1990." Los Alamos, NM: Los Alamos National Laboratory, LA-119545-MS, 1990.

- Morse, Norman R., H. Ferman Kelso, and John Morrison. "Computer Science and Services Division, Semiannual Report of the CCF Coordinating Committee." Los Alamos, NM: Los Alamos Scientific Laboratory, LA-6967-MS, 1977.
- Murray, Charles. *The Supermen: The Story of Seymour Cray and the Technical Wizards Behind the Supercomputer*. New York: John Wiley, 1997.
- "New Division Formed at LASL." *The Atom* 5, no. 5 (May 1968): 17-18.
- Norberg, Arthur L. *Computers and Commerce: A Study of Technology and Management at Eckert-Mauchly Computer Company, Engineering Research Associates, and Remington Rand, 1946-1957*. Cambridge, MA: MIT Press, 2005.
- Norberg, Arthur L., and Judy E. O'Neill. *Transforming Computer Technology: Information Processing for the Pentagon, 1962-1986*. Johns Hopkins Studies in the History of Technology, ed. Merritt Roe Smith. Baltimore: The Johns Hopkins University Press, 1996.
- Normand, Eugene, Jerry L. Wert, Heather Quinn, Thomas D. Fairbanks, Sarah Michalak, Gary Grider, Paul Iwanchuk, John Morrison, Stephen Wender, and Steve Johnson. "First Record of Single Event Upset on the Ground, Cray-1 Computer Memory at Los Alamos in 1976." *IEEE Transactions on Nuclear Science* 57, no. 6 (December 2010): 3114-3120.
- Owens, Larry. "Where Are We Going, Phil Morse? Changing Agendas and the Rhetoric of Obviousness in the Transformation of Computing at MIT, 1939-1957," *IEEE Annals of the History of Computing* 18, no. 4 (1996): 34-41.
- Pederson, Jeffery L. "MUX, adieu, faithful friend." *The Atom* 16, no. 8 (October, 1979): 14-16.
- Pinkston, J. T. "Supercomputer Trends and Needs: A Government Perspective." Nicholas Metropolis, et al., eds, *Frontiers of Supercomputing*. University of California Press, 1986.
- Powell, Michael L. "The DEMOS File System." *Proceedings of Sixth ACM Symposium on Operating Systems Principles* (November 1977): 33-42.

Powell, Michael L., and Barton P. Miller. "Process Migration in DEMOS/MP." *SOSP '83 Proceedings of the Ninth ACM Symposium on Operating Systems Principles*. New York: ACM, 1983.

Pugh, E. W. *Building IBM: Shaping an Industry and Its Technology*. Cambridge, MA: MIT Press, 1995.

_____. *Memories that Shaped an Industry: Decisions Leading to an IBM System/360*. Cambridge, MA: MIT Press, 1984.

Pugh, E. W., and William Aspray. "Creating the Computer Industry." *IEEE Annals of the History of Computing* 18, no. 2 (1996): 7-17.

Purcell, Charles J. "The Control Data STAR-100—Performance Measurements." *AFIPS '74 Proceedings of the May 6-10, 1974, National Computer Conference and Exposition*. New York: ACM, 1974.

Rankin, Joy. *A People's History of Computing in the United States*. Cambridge, MA: Harvard University Press, 2018.

Redmond, Kent C., and Thomas M. Smith. *Project Whirlwind: The History of a Pioneer Computer*. Bedford: Digital Equipment Corporation, 1980.

Reilly, Edwin D. *Milestones in Computer Science and Information Technology*. Westport, CT: Greenwood Press, 2003.

Rhodes, Richard. *Dark Sun: The Making of the Hydrogen Bomb*. New York: Simon and Schuster, 1995.

_____. *The Making of the Atomic Bomb*. New York: Simon and Schuster, 1987.

Russell, Richard M. "The Cray-1 Computer System." *Communications of the ACM* 21, no. 1 (January 1978): 63-72.

Seidel, Robert W. "A Home for Big Science: The Atomic Energy Commission's Laboratory System." *Historical Studies in the Physical and Biological Sciences* 16, no. 1 (1986): 135-175.

_____. "From Mars to Minerva: The Origins of Scientific Computing in the AEC Labs." *Physics Today* 49 (October 1996): 33-39.

_____. *Los Alamos and the Development of the Atomic Bomb*. Los Alamos: Otowi Press, 1995.

_____. "The National Laboratories of the Atomic Energy Commission in the Early Cold War." *Historical Studies in the Physical and Biological Sciences* 32 (2001): 145-162.

Shroyer, Jo Ann. *Secret Mesa: Inside Los Alamos National Laboratory*. New York: John Wiley & Sons, Inc., 1998.

Slotnick, D. L., C. W. Borck, and R. C. McReynolds. "The SOLOMON Computer." *Proceedings of the 1962 Fall Joint Computer Conference*. New York: ACM, 1962.

Small, James S. *The Analogue Alternative: The Electronic Analogue Computer in Britain and the USA, 1930-1975*. London: Routledge, 2001.

Stone, H. S. *High-Performance Computer Architecture*. Reading, MA: Addison-Wesley, 1987.

Tarter, C. Bruce. *The American Lab: An Insider's History of Lawrence Livermore National Laboratory*. Baltimore, MA: John Hopkins University Press, 2018.

"The Computer Science and Services Division Annual Report, January—December 1975," Los Alamos Scientific Laboratory, Los Alamos, NM, LA-6228-PR, March, 1976.

"The Shrinking of Stretch." *Datamation* (June 1961): 13-17.

Thornton, J. E. "The CDC 6600 Project." *Annals of the History of Computing* 2, no. 4 (Oct. 1980): 338-348.

Thornton, J. E. *Design of a Computer: The Control Data 6600*. Glenview, Illinois: Scott, Foresman and Company, 1970.

- Tolmie, D. E., A. G. Dornhoff, and S. C. Tenbrink. "Interconnecting Computers with the High-Speed Parallel Interface." Los Alamos, NM: Los Alamos National Laboratory, LA-9503-MS, 1982.
- _____. "The LASL Integrated Computer Network." Los Alamos, NM: Los Alamos Scientific Laboratory, LA-UR-76-2321, 1976.
- Tomash, E., and A. A. Cohen. "The Birth of an ERA: Engineering Research Associates, Inc., 1946-1955." *Annals of the History of Computing* 1, no. 2 (Oct. 1979): 83-97.
- Ulam, Stanislaw. *Adventures of a Mathematician*. New York: Charles Scribner's Sons, 1976.
- United States Nuclear Tests: July 1945 through September 1992*. DOE/NV-209, Rev. 14, December, 1994.
- Vogel, William F. "'The Spitting Image of a Woman Programmer': Changing Portrayals of Women in the American Computing Industry, 1958–1985." *IEEE Annals of the History of Computing* 39, no. 2 (April-June 2017): 49-64.
- von Neumann, John. "First Draft of a Report on the EDVAC." *Papers of John Von Neumann on Computing and Computer Theory*. eds. William Aspray and Arthur Burks. Cambridge: The MIT Press, 1987.
- Voorhees, Edward A. "Recollections of the 701 at Los Alamos." *Annals of the History of Computing* 5, no. 2 (April-June 1983): 177-178.
- Westwick, Peter J. *The National Labs: Science in an American System, 1947-1974*. Cambridge, MA: Harvard University Press, 2003.
- Worthy, James C. "Control Data Corporation: The Norris Era." *IEEE Annals of the History of Computing* 17, No. 1 (1995): 47-53.
- Yood, Charles N. *Hybrid Zone: Computers and Science at Argonne National Laboratory, 1946-1992*. Boston, MA: Docent Press, 2013.

Yost, Jeffrey R. *Making IT Work: A History of the Computer Services Industry*.
Cambridge, MA: MIT Press, 2017.

_____. "The Origin and Early History of the Computer Security Software Products
Industry." *IEEE Annals of the History of Computing* 37, no. 2 (April-June 2015):
46-58.

Archival Collections

Buildings Collection, HPC Archive, Los Alamos National Laboratory, Los Alamos, New Mexico.

Control Data Corporation Collection (CBI 80), Charles Babbage Institute, University of Minnesota, Minneapolis.

Margaret R. Fox Papers (CBI 45), Charles Babbage Institute, University of Minnesota, Minneapolis.

LASL, LANL Org. Charts Collection, HPC Archive, Los Alamos National Laboratory, Los Alamos, New Mexico.

Nicholas Metropolis Papers (A-98-011), LANL Archives, Los Alamos National Laboratory, Los Alamos, New Mexico.

Records of C-DO (A-93-015), LANL Archives, Los Alamos National Laboratory, Los Alamos, New Mexico.

Edward A. Voorhees Papers (A-83-011), LANL Archives, Los Alamos National Laboratory, Los Alamos, New Mexico.

Jack Worlton Papers, HPC Archive, Los Alamos National Laboratory, Los Alamos, New Mexico.

A Note on LANL Archival Collections

The original sources from the LANL Archives and the HPC History of LANL Supercomputing Project Archive are all subject to, and available through, the Freedom of Information Act (FOIA), making these sources subject to scholarly review. The HPC Archive of physical materials is located in a single room of Building 132 in TA-3 of LANL. HPC's digital archive is located on LANL's secure, unclassified "Yellow" network.

Cited Original Oral History Interviews

- Blanchard, Sean. Interview by Nicholas Lewis. Los Alamos, NM, 6/19/2017.
- Bryson, Don. Interview by Nicholas Lewis. Los Alamos, NM, 7/1/2014.
- Christman, Ronald. Interview by Nicholas Lewis. Los Alamos, NM, 7/2/2014.
- Clifford, Jim. Interview by Nicholas Lewis. Los Alamos, NM, 7/23/2014.
- Collins, Bill. Interview by Nicholas Lewis. Los Alamos, NM, 2/19/2015.
- Daniel, David. Interview by Nicholas Lewis. Los Alamos, NM, 7/5/2018.
- Dubois, Andrew. Interview by Nicholas Lewis. Los Alamos, NM, 7/17/2014.
- Dubois, David. Interview by Nicholas Lewis. Los Alamos, NM, 7/17/2014.
- Folkner, Charlie. Interview by Nicholas Lewis. Los Alamos, NM, 8/13/2015.
- Forslund, David. Interview by Nicholas Lewis. Los Alamos, NM, 8/2/2018.
- Grider, Gary. Interview by Nicholas Lewis. Los Alamos, NM, 8/21/2015.
- Hayes, Ann. Interview by Nicholas Lewis. Los Alamos, NM, 8/14/2015.
- Hill, Thomas R. Interview by Nicholas Lewis. Los Alamos, NM, 4/25/2019.
- Iwanchuk, Paul. Interview by Nicholas Lewis. Los Alamos, NM, 7/7/2014.
- McGavran, Harry. Interview by Nicholas Lewis. White Rock, NM, 7/17/2014.
- Miller, Raymond. Interview by Nicholas Lewis. Los Alamos, NM, 6/25/2014.
- Morrison, John. Interview by Nicholas Lewis. Los Alamos, NM, 7/1/2014.
- Morse, Norman. Interview by Nicholas Lewis. Los Alamos, NM, 7/30/2018.
- Simmons, Margaret. Interview by Nicholas Lewis. Los Alamos, NM, 8/20/2015.

Simmons, Michael. Interview by Nicholas Lewis. Los Alamos, NM, 8/26/2015.

Spack, William. Interview by Nicholas Lewis. Los Alamos, NM, 6/25/2014.

Thompson, Buck. Interview by Nicholas Lewis. Los Alamos, NM, 9/3/2015.

Tolmie, Don. Interview by Nicholas Lewis. Los Alamos, NM, 7/14/2014.

Worlton, Reid. Interview by Nicholas Lewis. St. George, UT, 11/29/2015.

Worlton, William J. Interview by Nicholas Lewis. St. George, UT, 11/29/2015.