VISION FOR SEQUOIA 2000 PHASE II:

THE SEQUOIA COMPUTATIONAL INFRASTRUCTURE

September 26, 1994

Principal Investigators:

Joseph Pasquale
Department of Computer Science and Engineering, C-014
University of California
La Jolla, CA 92093-0114
(619) 534-2673
pasquale@cs.ucsd.edu

Randy Katz
Department of Electrical Engineering & Computer Science
University of California
Berkeley, CA 94720-1776
(510) 642-8778
katz@cs.berkeley.edu

Jeff Dozier
Center for Remote Sensing and Environmental Optics
University of California
Santa Barbara, CA 93106-3060
(805) 893-8116
dozier@crseo.ucsb.edu

1. Introduction

Sequoia 2000 is a collaboration of computer scientists and Earth scientists at the University of California, whose goal is to design, build, and put into practical use a distributed information system that effectively supports intensive computational and storage requirements, such as those evidenced by Earth science applications. In addition, the system will provide tools which support collaboration among distributed teams of scientists. The Sequoia project is beginning its second three year phase.

During Phase I, the project focused on four technology areas: databases, networking, storage, and visualization. The project has made significant contributions, most notably in how very large data objects (tens of megabytes to hundreds of gigabytes) are queried and searched, communicated between processes within a computer and between computers over a network, stored and retrieved between different levels of storage, and visualized and animated with a high degree of user interaction. In addition (and perhaps most important), these advances are being incorporated in working systems which are helping Earth scientists do their science more effectively.

As we begin Phase II, there is a need for integration and a more systematic approach toward building a coherent wide-area distributed "system" (rather than a collection of loosely-connected

subsystems). While system integration efforts are already underway (e.g. a NASA-sponsored grant to solve various "end-to-end" problems), Phase II will begin from basic principles and seek to provide an "infrastructure" that not only supports these subsystems, but addresses a fundamental goal, "power harnessing" (discussed below), that makes the Sequoia project uniquely important.

Sequoia scientists have Earth science applications with massive computational and storage requirements, measured in teraops and terabytes. These requirements stress all known information systems beyond their limits. They generally far exceed the power provided by a scientist's workstation; even the combined resources provided by the scientist's local area networking environment are commonly inadequate. To begin satisfying these demands, one must look to the cumulative power of very large-scale distributed systems, e.g. environments rich in computational and storage resources connected by wide-area high-speed networks.

It is noteworthy that the Sequoia project has a significant amount of computational and storage resources and has in place an operational high-speed wide-area network. Moreover, the project has fostered a strong partnership between applications experts and computer technologists. These factors present a key opportunity for us to explore ways to take advantage of economy of scale by allowing any single Sequoia application to harness as much power as it needs, potentially all the power provided by the effective cooperative use of all the resources in the system. With the right system design approach this exceeds the sum of the power provided by each individual resource because of synergistic effects.

"Power harnessing", the ability to concentrate as much of the distributed power available in a large-scale system to meet the demands of any single application, is the major goal for Phase II. Power harnessing is difficult for numerous reasons, including the physical distribution, the decentralized control and differing ownership, the varying performance and reliability, etc, of the resources. Solutions to these 21st century resource sharing issues are important to Sequoia as well as all projects where scalability to very large numbers of resources and users are important. Indeed, we believe these solutions will be crucial for the ultimate success of the National Information Infrastructure, for which we view Sequoia as an early prototype.

1.1. The Sequoia Computational Infrastructure

To meet this goal, we will take the approach of building an underlying foundation or basic framework that supports building and operating such a system, i.e. an infrastructure. The Sequoia Computational Infrastructure supports the development of an environment that provides power harnessing, and has the following characteristics:

1) An architecture based on service-providers that interact with clients (and other service-providers) through a well defined interface that supports the following basic system wide operations.

- service advertising (for resource discovery)
- pricing of services (with provisions of varying levels of quality)
- service requests and responses
- contracting (price negotiation + establishment of connection)

2) A wide-area network that supports high-bandwidth (Gb+/s) communication across the entire system, i.e. between all clients and service-providers (we currently have a CNRI/NSF "Planning Grant" to plan the upgrade of the Sequoia network to Gb/s speeds).

- at Gb/s speeds (and solutions to latency reduction), the network can effectively act as a "backplane" for all service-providers (the "devices" in the backplane analogy);
- service-providers "plug in" to the backplane through a well-defined interface, promoting resource portability and ease of integration
- latency is reduced through a variety of techniques, e.g. caching, fastpaths, and prefetching

- allows the dynamic formation of tightly-coupled interactions to deliver as many resources to a problem on demand

3) A contracting system that controls complexity of services by integrating composition of service-providers and pricing of services.

- clients obtain service through contracts
- contracts are negotiated based on system-enforced bidding system
- service-providers can subcontract to other service providers

4) A market economy that supports the effective sharing of resources.

- users gain credits by making their resources available to others
- credits can be used to acquire resources when they are needed

1.2. Technologies

We expect that the following will be important technology areas in Phase II (in alphabetical order):

- collaboration tools
- distributed databases
- gigabit networking
- hierarchical MASS storage management
- multimedia
- ubiquitous access - wireless networks and mobile computing
- visualization
- wide-area file systems

Thus, the research areas that will be emphasized are those that either directly relate to the design and implementation of the infrastructure, or that involve other areas such as those listed above which are built on top of the infrastructure. Of course, the final research areas will ultimately depend on the interests of the Sequoia Phase II investigators.

1.3. Applications

Since its inception, Sequoia has been an Earth science application driven project, and continues to be so in Phase II. The applications include remote sensing, global climate modeling, and management and analysis of meteorological and hydrological data. It is worth noting that while the focus is on Earth science, we believe other application areas will also benefit from our model for a computational infrastructure.

Remote sensing applications drive the Sequoia architecture in several ways. Images are large (a Landsat Thematic Mapper frame is about 300 Mbytes), accessed from remote archives over networks, and produced in a plethora of different formats. Analysis requires quick perusal of multiple images (browsing through abstracts) followed by intensive calculations over large data sets. End-to-end analysis exercises the lineage problem, hence it is a problem in database management.

The climate modeling efforts involve a collaborative project to develop a first generation Earth System Model encompassing the coupled global atmosphere and ocean systems, including chemical tracers that are found in, and may be exchanged between, the atmosphere and the oce-

ans. This will be the first attempt we know of to extend a General Circulation Model with a chemical tracer model, and it will be used to study major current problems in climate, climate change, and climate/chemistry interactions. These problems include the general circulation of the coupled atmosphere/ocean system, and the global biogeochemical cycle of carbon. The model will require computational advances in parallelism, database management, and visualization.

Analysis of surface meteorological and hydrological data requires the ability to browse and select data sets, access and subset them, visualize, and conduct multivariate analyses.

To drive the requirements for the Sequoia infrastructure, we will create and maintain a quasi-operational information system that supports the Sequoia applications.

1.4. Funding

Funding will be very different in Phase II. A major problem during Phase I was the lack of funding for creating and maintaining an infrastructure. Thus, while investigators might obtain small amounts of funding from the project for their research, the actual benefits were low given that the infrastructure was not there in full force to support them.

The Phase II funding philosophy begins with the principle that the only direct provision by the project to investigators is the use of the Sequoia infrastructure. By focusing on creating a high quality infrastructure that provides capabilities that one would not otherwise typically have, the hope is to attract investigators who see benefits because it enhances their research and because it provides leverage for obtaining external research grants.

Consequently, there will be two Sequoia funding strategies: an infrastructure support strategy, and a research support strategy. Funds raised by the project (e.g. donations by industrial partners for the Sequoia project in general) will go directly toward supporting the infrastructure. This includes infrastructure related equipment (e.g. workstations for storage servers and network routers, disks and tapes for storage, network communication bandwidth), software, maintenance, and staff (e.g. administrative assistants, staff at each site for maintaining local infrastructure related equipment). It is important that the infrastructure be open and supportive of a wide range of computer equipment. Consequently, we will encourage the inclusion of and seek the support of a variety of computer/communication companies to be our industrial partners.

The research support strategy works as follows. Investigators who are part of the Sequoia project will form teams to obtain funds for their Sequoia related research. This strategy worked well in the latter part of Phase I, e.g. the NASA "End-to-End" 3 year grant of $2.4 million. The project directors will be actively involved in identifying funding sources and advertising them to the project investigators. The project will continue to actively seek support from state and federal government agencies, as well as from industry.

1.5. Summary

The Sequoia 2000 project, a collaboration of computer scientists and Earth scientists at the University of California, is beginning its second 3 year phase. The major goal for Phase II is to support power harnessing, the ability to dynamically concentrate as much of the cumulative resource power in a wide-area distributed system to meet the demands of any single application. The project will focus on creating an underlying hardware and software infrastructure that supports power harnessing, and above which software systems targeted to support Earth science applications can be built. As in Phase I, the project will seek to develop partnerships, both financial and intellectual, with the university, industry, and the state and federal government.