Project Sequoia 2000

Joseph Pasquale

Computer Systems Laboratory
Department of Computer Science and Engineering
University of California, San Diego

Sequoia 2000 is DEC's $15 million flagship project of the '90s which addresses the massive data storage, network, and visualization requirements of Global Change researchers [1]. Faculty investigators include computer scientists and environmental scientists from various University of California campuses, led by database researcher Mike Stonebraker of UC Berkeley and environmental scientist Jeff Dozier of UC Santa Barbara. A group of us (Domenico Ferrari of UC Berkeley, George Polyzos of UC San Diego, and myself) are focusing on network protocols and operating system I/O software which will provide high-speed communication between a terabyte database being designed at UC Berkeley and any number of scientists' workstations located throughout the Sequoia 2000 network. Much of this data is in the form of a time-sequence of highly detailed images, collected by satellites and remote sensors located on the ocean or on land, which scientists wish to process, visualize, browse and manipulate interactively. Consequently, the network must support high-bandwidth transmission of these images and meet various delay constraints to simulate visual continuity. Some of the issues we are working on include design of protocols which support deterministic and statistical performance guarantees [2], integration of I/O system software and network protocol software to allow dynamic construction of in-kernel datapaths between the network device(s) and I/O source/sink devices (e.g. the frame buffer), determination of where to locate compression and decompression processes, and determination of how to take advantage of prioritized packets containing hierarchically coded data so that the network has the flexibility of selectively dropping packets if necessary (e.g. to prevent congestion).

Operating System and Network Support for Multimedia

We are also working on system software which supports multimedia applications, specifically those with multiple audio and video I/O components [3]. Given the intensive I/O demands made by digital video (and to a lesser extent, high-quality audio), we are working on the design of an I/O kernel which supports fast I/O switching of communication resources (e.g. the system bus) to support multiple I/O streams with time-varying bandwidth requirements and other real-time constraints. We have developed a multicast routing algorithm [4] which seeks to minimize used channel capacity and still meet delay constraints between source and all destinations. Finally, we are working on file system designs which are "multi-structured" [5], composed of different storage structures (e.g. striped, mirrored, log) each of which responds well to a particular type of file access pattern. These structures are integrated in a single design to simultaneously provide high throughputs for very large sequentially-accessed files (e.g. video), while providing fast response for more common file workloads containing smaller file requests as typified by Unix office/engineering environments.

---

[1] M. Stonebraker and J. Dozier, "Sequoia 2000: Large capacity object servers to support global change research," Sequoia 2000 Technical Report #1, U. C. Berkeley, July 1991.

[2] D. Ferrari, "Client requirements for real-time communication services," *IEEE Communications Magazine,* volume 28, number 11, pp. 65-72, November 1990.

[3] J. Pasquale and G. Polyzos, "System support for multimedia applications," *Proc. 1st International Workshop on Network and Operating System Support for Digital Audio and Video,* International Computer Science Institute Technical Report TR-90-062, Berkeley, CA, November 1990.

[4] V. Kompella, J. Pasquale, and G. Polyzos, "Multicasting for multimedia applications," *Proc. IEEE INFOCOM '92,* Florence, Italy, May 1992 (to appear).

[5] K. Muller and J. Pasquale, "A high performance multi-structured file system design," *Proc. 13th Symposium on Operating System Principles,* Asilomar, CA, pp. 56-67, October 1991.