Final Report on
Research in Parallel Computing: 1990-95

Very Large Scale Operating Systems
The Gigabit Nectar Project
Network-Based Multicomputers

Dale James and C. Roy Taylor, Eds.
CMU-CS-96-100
12 January 1996

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Richard F. Rashid, Brian Bershad, Carl Love
Peter Steenkiste
Thomas Gross

This research was sponsored by the Advanced Research Projects Agency, DoD, through ARPA Order 7330, and monitored by the Space and Naval Warfare Systems Command under contract MDA972-90-C-0035. Views and conclusions contained in this document are those of the authors and should not be interpreted as representing official policies, either expressed or implied, of the Advanced Research Projects Agency or the United States Government.

Keywords: Operating systems, parallel computing, Nectar, Warp, networks, multicomputers.

Abstract

The first section of this report documents progress Carnegie Mellon's Mach Group made in research in very large scale operating systems, for the period 1990-1993. Parts of the Mach Operating system have been incorporated in a number of commercial operating systems, including Encore's Multimax, NeXT OS, MachTen for the Macintoshes, Omron's Luna, OSF/1 for the DEC Alpha, and IBM's OS/2 for the RS6000 based machines.

Multicomputers consisting of off-the-shelf computers connected by commercial high-speed networks form an economically attractive computing platform for a large class of applications. However, while high-speed networks are fairly widely available (e.g., HIPPI and ATM), many computer systems have problems delivering this high bandwidth to the applications, thus limiting the class of applications that multicomputers can support effectively. The Gigabit Nectar project has developed a network interface architecture that supports efficient high-bandwidth end-end communication. This architecture has been implemented for workstations (DEC Alpha) and distributed-memory systems (iWarp) and has been deployed in the Gigabit Nectar testbed. The second section of this report describes the Nectar network interface, its implementation and performance, and summarizes our application experience in the testbed.

No existing parallel architecture (including SIMD machines, MIMD machines, vector processors, machines with a common, globally-addressable memory, machines with private memories) has demonstrated its superiority for all possible applications. That is, given different applications, different parallel machines may deliver the highest absolute performance or the best price-performance. A good solution to this problem is to link a number of different systems (ranging from workstations to parallel supercomputers) with a high-speed network to build a single, heterogeneous system. The final section of this report details our research to address this issue during the period 1990-95.