

# Design Alternatives for Shared Memory Multiprocessors \*

John B. Carter, Chen-Chi Kuo, Ravindra Kuramkote, Mark Swanson

{retrac, chenchi, kuramkot, swanson}@cs.utah.edu

WWW: <http://www.cs.utah.edu/projects/avalanche>

UUCS-98-011

Department of Computer Science  
University of Utah, Salt Lake City, UT 84112

March 23, 1998

## Abstract

In this paper, we consider the design alternatives available for building the next generation DSM machine (e.g., the choice of memory architecture, network technology, and amount and location of per-node remote data cache). To investigate this design space, we have simulated six applications on a wide variety of possible DSM architectures that employ significantly different caching techniques. We also examine the impact of using a special-purpose system interconnect designed specifically to support low latency DSM operation versus using a powerful off the shelf system interconnect. We have found that two architectures have the best combination of good average performance and reasonable worst case performance: CC-NUMA employing a moderate-sized DRAM remote access cache (RAC) and a hybrid CC-NUMA/S-COMA architecture called AS-COMA or *adaptable S-COMA*. Both pure CC-NUMA and pure S-COMA have serious performance problems for some applications, while CC-NUMA employing an SRAM RAC does not perform as well as the two architectures that employ larger DRAM caches. The paper concludes with several recommendations to designers of next-generation DSM machines, complete with a discussion of the issues that led to each recommendation so that designers can decide which ones are relevant to them given changes in technology and corporate priorities.