VISUALIZE Workstation Memory Subsystem

By Kirsten Olsen
HP Workstation System Lab

With the new HP VISUALIZE UNIX® Workstations family (B1000, C3000, J5000), the customer will find a completely redesigned memory architecture. With this new architecture comes dramatic performance improvements in memory intensive applications over the HP PA-8200-based (C200, C240, J282, J2240) family of workstations. This paper will explain some of the key elements that contribute to the memory performance improvements in the HP VISUALIZE UNIX® Workstations. It will also compare the architecture of the new memory subsystem with that of the HP PA-8200 subsystem. Finally, and most importantly, application performance will be examined in comparison to the HP PA-8200 family of workstations and those of our competitors.

Overview of Memory Subsystem

The new memory subsystem architecture was designed with the goal of dramatically improving the memory subsystem performance in the HP VISUALIZE UNIX® family of workstations. At the heart of this new architecture are custom HP VLSI and memory DIMMs, which were designed to provide a low latency, high bandwidth memory subsystem. The HP VISUALIZE UNIX® Workstation systems have an idle system memory latency roughly 40% better than that of the HP PA-8200 products, as well as a memory bandwidth two times that of the HP PA-8200 products. The peak memory bandhwidth of the PA-8200 products was 960MB/s; the peak memory bandwidth of the HP VISUALIZE UNIX® Workstations family is 1.92GB/s.

The new subsystem was designed using the latest in synchronous DRAM technology, taking advantage of the memory performance benefits inherent to synchronous DRAM. Sophisticated scheduling algorithms were used to help improve the busy system memory latency and to take full advantage of the memory subsystem bandwidth. The system was also designed to allow maximum memory bandwidth with a minimum of memory in the system. An innovative clocking scheme used in the design of the system board and memory DIMMs enabled the memory subsystem to operate at a frequency much higher than typical for the SDRAMs used. All of these features of the new memory system combine to give the customer dramatic improvements in application performance with the new HP VISUALIZE UNIX® Workstations.

Key Design Features

New Architecture and VLSI

Two new chips were designed for the VISUALIZE memory subsystem: The memory controller chip, which also acts as the I/O controller; and a chip which acts as a large data mux (there are 2 of these chips in the system). In comparison, the HP PA-8200 memory subsystem has a master memory controller and slave memory controllers (3 SMCs in a 12 DIMM system and 4 SMCs in a 16 DIMM system), and data mux chips (4 chips in both the 12 and 16 DIMM systems). The leaner architecture greatly improves the idle system

latency of the HP VISUALIZE UNIX® Workstation memory subsystem, and contributes to an improved price/performance ratio.

The new memory architecture was designed so that it operates with a minimum of one DIMM loaded in the system, in contrast to the HP PA-8200 memory subsystem that requires a minimum of 2 DIMMs loaded in the system. This allows the HP VISUALIZE UNIX® Workstation systems to reach maximum memory bandwidth with just 1 DIMM in the system, allowing customers to buy a minimum of memory and still achieve maximum memory bandwidth.

New Synchronous DRAMs

A key feature in this new subsystem is the use of synchronous DRAMs (SDRAMs), a new DRAM architecture. Changes from asynchronous DRAMs, which were used in the HP PA-8200 workstations and are currently used in Sun's Ultrasparc5, Ultrasparc10 and Ultrasparc60 systems, include a synchronous interface between the DRAM and the outside world and multiple banks internal to the DRAM. One advantage of the synchronous interface of SDRAMs is that it latches address and control signals into the chip on a given clock edge. This frees up the address and control bus for other activities, such as issuing precharge and activate commands to other banks in the subsystem, while the SDRAMs decode the latched signals. With asynchronous DRAMs, the memory controller has to sit idle while these internal DRAM operations are performed. This extra address and control bandwidth provided by SDRAMs increases with higher frequencies and contributes to lower busy system memory latency in the HP VISUALIZE UNIX® Workstation systems.

The synchronous data interface of SDRAMs is another big advantage of SDRAMs over asynchronous DRAMs. With SDRAMs, data is clocked in or out of the SDRAM with the SDRAM or system clock, making the data rate of SDRAMs equal to the clock rate. With asynchronous DRAMs, other timing constraints come into play that limit the data rate of these devices. Comparing the SDRAMs used in the HP VISUALIZE UNIX® Workstation systems with the asynchronous DRAMs used in the HP PA-8200 systems shows that the per device bandwidth of the SDRAMs is close to 5 times higher than that of the asynchronous DRAMs. This synchronous data interface significantly contributes to the higher memory bandwidth of the HP VISUALIZE UNIX® Workstation systems.

Another advantage of SDRAMs is that they have multiple independent banks on a single chip. This allows memory interleaving to take place at the chip level. HP VISUALIZE UNIX® Workstation systems can now take advantage of interleaving performance gains with only one DIMM in the system, as opposed to needing multiple DIMMs in a system like the HP PA-8200 systems. This interleaving is advantageous in that with multiple internal banks, one page in a bank of an SDRAM can be accessed while a page in a second bank of the same SDRAM is activated, allowing the page in the second bank to be accessed as soon as possible after access to the first bank. This process hides the activate time of the page in the second bank in the access time of the page in the first bank. This feature of SDRAMs contributes to improved busy system memory latency.

Memory in the HP VISUALIZE UNIX® Workstations uses varying SDRAM densities—64Mb, 128Mb and, in the future, 256Mb—to provide customers with a range of memory DIMM sizes.

New Scheduling Algorithms

Another important design feature of the HP VISUALIZE UNIX® Workstation memory system is the use of performance enhancing scheduling algorithms. One key feature of the HP VISUALIZE UNIX® Workstation systems that allows more advanced scheduling algorithms to be used is the use of open pages. With the multiple internal banks of SDRAMs it became advantageous to leave memory pages open. In leaving memory pages open, the standard DRAM sequence *of activate -> column access -> precharge (close an open page)* is rearranged to be *precharge (close an open page) -> activate -> column access*. By leaving pages open and subsequently getting a read or write command to the same page, the precharge and activate parts of the SDRAM sequence can be avoided. The use of open pages can thus improve overall busy system memory latency. In the HP PA-8200 systems pages were not left open once they were accessed.

Another improvement to the HP VISUALIZE UNIX® Workstation memory subsystem is the use of out of order read returns. A good example of how out of order read returns are used is the reordering of read commands based on whether the read is to an open or closed page. If two reads are waiting to be issued, the first to a closed page and the second to an open page, HP VISUALIZE UNIX® Workstation systems will reorder the reads to allow the read to the open page to execute first, or to be issued "out of order." The HP PA-8200 systems did not use out of order read returns. As explained in the open page discussion above, this improves the busy system memory latency of the HP VISUALIZE UNIX® Workstation systems.

The HP VISUALIZE UNIX® Workstation systems also use a larger memory transaction queue. The larger queue is beneficial in that it allows a greater number of transactions to be prioritized. This feature couples nicely with the use of out of order read returns. With a larger transaction queue, the chance of finding a read or a write to an open page is much greater. The HP PA-8200 systems had a much smaller memory transaction queue than that of the HP VISUALIZE UNIX® Workstation systems. All of the scheduling features mentioned above contribute to the improved performance of the HP VISUALIZE UNIX® Workstation memory subsystem.

New Board and DIMM Designs

A key feature in the improved performance of the HP VISUALIZE UNIX® Workstation memory subsystem is the operation of standard PC100 SDRAMs at 120Mhz, as opposed to their standard operation at 100MHz. The memory subsystem is able to operate at this frequency due to a novel clocking scheme that essentially eliminates clock skew from the system's DRAM timing equations. Being able to operate the memory subsystem at this increased frequency greatly improved the performance of the HP VISUALIZE UNIX® Workstation systems.

Application Performance Data

Below you will find performance charts comparing HP VISUALIZE UNIX® Workstation systems to HP PA-8200 systems and to Sun's Ultrasparc60 systems. Performance gains in both EDA and MDA applications are shown. In the EDA applications, Verilog and QuickHDL are two memory intensive applications.

It should be noted that not all performance improvements are due to improvements in the memory subsystem. Faster processor speeds and improved I/O performance also influence the performance numbers below. The performance numbers below are based on proprietary customer application workloads. Larger numbers represent better performance. The EDA Composite workload is an average of the other EDA application results of Verilog, Design Compiler, QuickHDL, Calibre and Cyclone. Similarly, the MDA Composite is an average of the UG, SDRC and PTC workloads. The systems shown in the J5000 comparisons had 1GB of memory and the systems shown in the C3000 comparisons had 512MB of memory.

J5000 Performance 2-CPU Systems (Larger numbers are better) EDA Applications

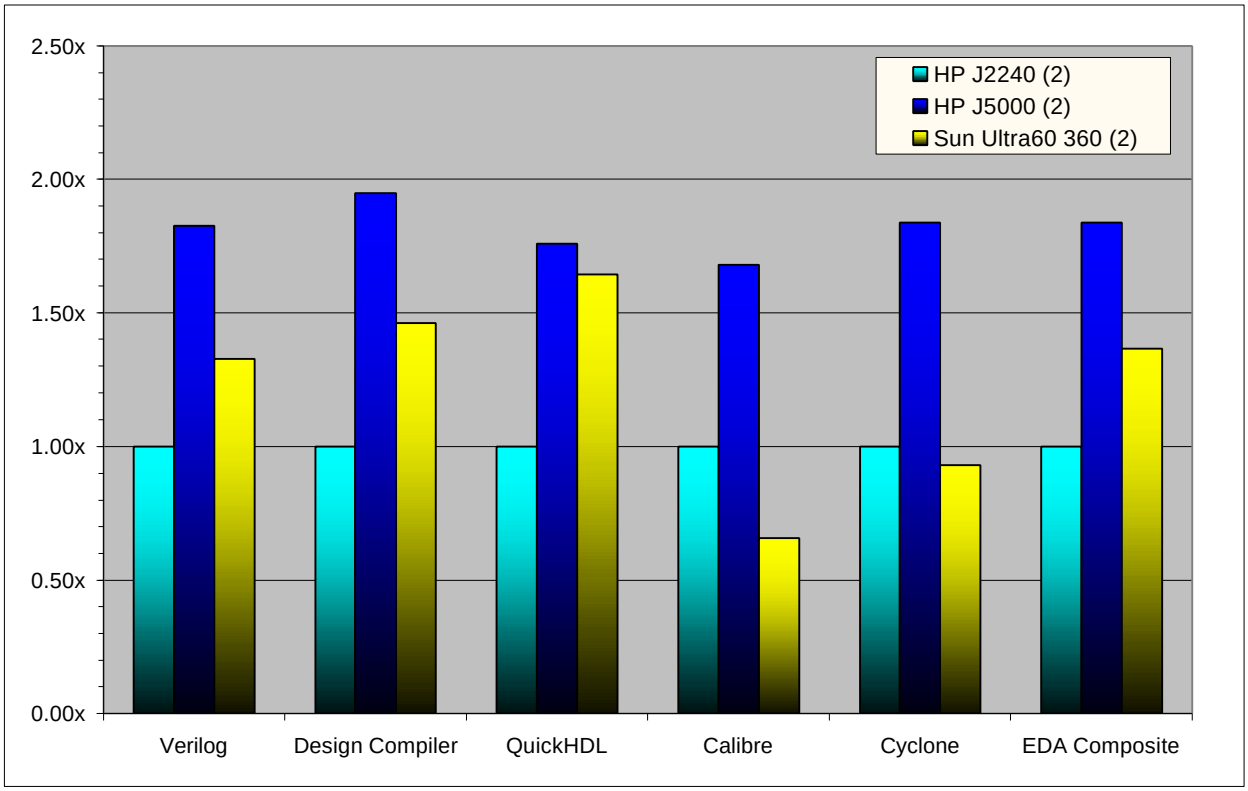

C3000 Performance (Larger numbers are better)
EDA Applications

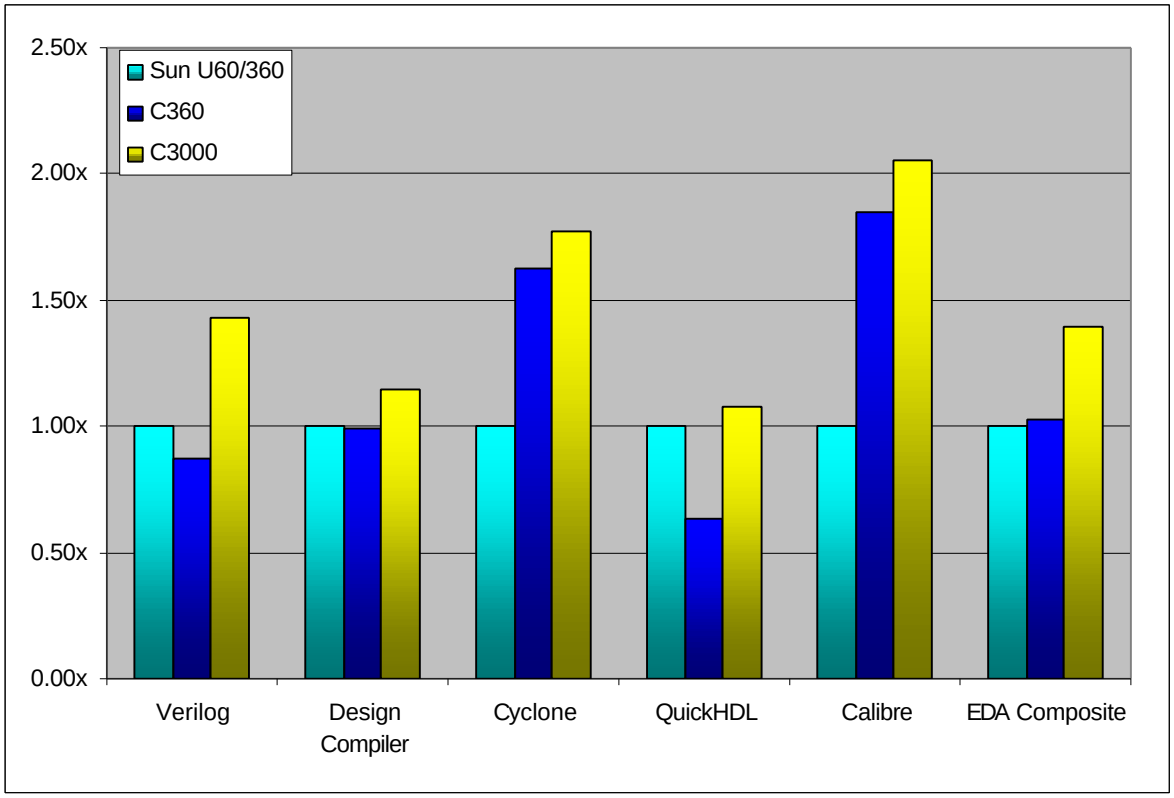

C3000 Performance (Larger numbers are better)
MDA Applications

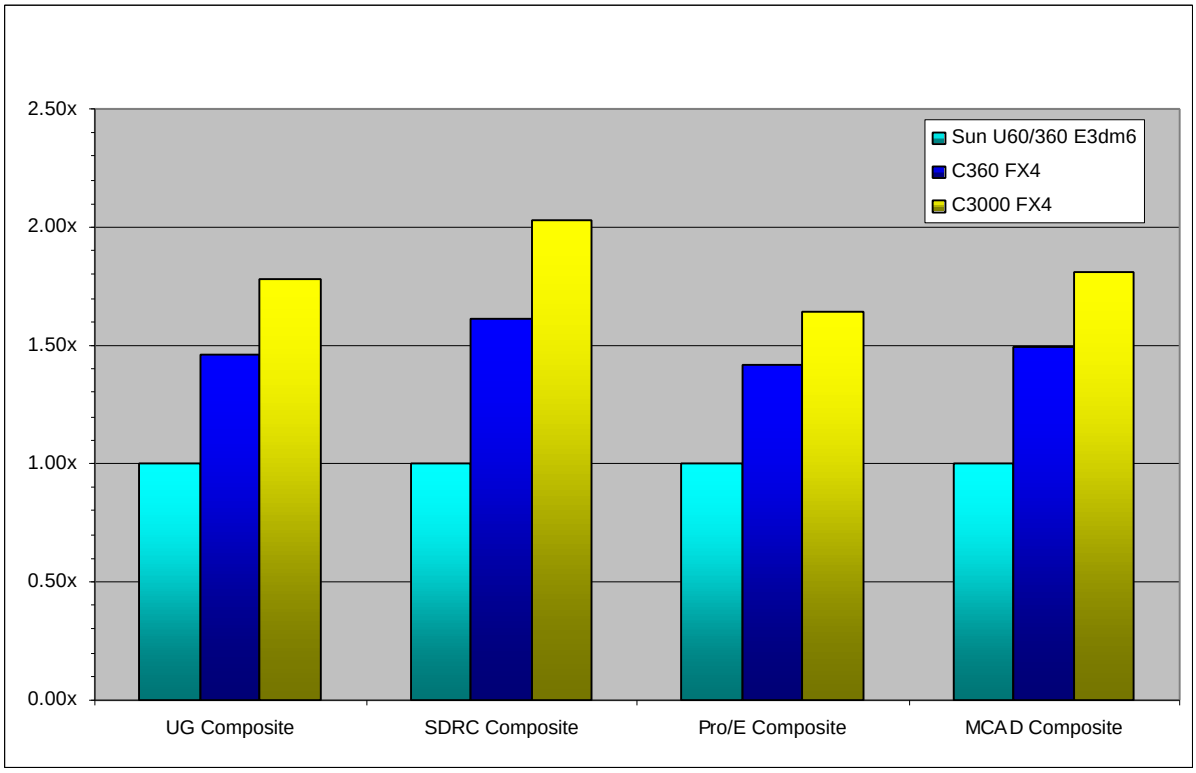

Summary

It is exciting to look at the application performance numbers and see the great results of the improved HP VISUALIZE UNIX® Workstation memory architecture. The improved application performance results are the direct result of a total 2X performance improvement plan based on:

- The 2X performance of this new HP VISUALIZE memory subsystem
- The greater than 2X performance of the new HP VISUALIZE I/O system
- The 2X performance of the 440MHz PA-8500 relative to the 236MHz PA-8200 CPU