

# HP 9000/735 Second Generation PA-RISC Snakes Workstation

Ed Keane, Pat McGuire

Hewlett-Packard Company  
Cupertino, California

## Abstract:

*This paper describes the second generation Snakes workstation, the HP 9000/735. The workstation includes a new PA-RISC processor operating at 50% greater frequency, providing large performance gains. Furthermore, integration techniques allow higher performance and greater function at lower cost.*

## 1.0 Introduction

In 1991 Hewlett-Packard introduced the first members of Series 700 PA-RISC workstations (Snakes). One member of the family, the model 730, led the marketplace for an unprecedented 18 months in desktop performance.

This workstation is logically partitioned as shown in Figure 1. There are four main subsystems in the workstation: the processor, the built-in I/O, the expansion graphics, and the expansion I/O (via industry standard EISA bus). Since the workstation is also physically partitioned according to the block diagram in Figure 1, board upgrades can be easily accommodated.




Figure 1: System Block Diagram

By late 1992, HP had announced upgrades to the processor, the I/O and the graphics boards. The

processor performance has been greatly increased, as a result of increased operating frequency, larger cache size, and larger supported memory. The I/O has been improved in terms of performance and function, including the addition of CD-quality audio, FDDI, and fast-wide SCSI. These I/O enhancements are the subject of an accompanying paper [1]. The graphics have also been improved in terms of performance and function [2].

This paper describes the new processor subsystem, and its improvements over the first generation processor [3]. The changes are described in terms of the desktop workstation model, the 735, but are germane to the desktside model, the 755, as well.

## 2.0 Processor Overview

The processor subsystem of the 735 workstation contains four separate functional units: the processor core, the system memory, the system bus interface, and the system clocks. The processor subsystem architecture is shown in Figure 2. Each of these subsystems will be described.




Figure 2: Processor Board Block Diagram

### 3.0 Processor Core

#### 3.1 Overview

The 735 is the first HP workstation to implement HP's PA-RISC 7100 CPU operating at 99 MHz [4,5]. This processor, introduced earlier in 1992, represents major improvements to the previous PA-RISC processors [6] in four separate areas: operating frequency, superscalar execution, cache and TLB optimizations, and floating point integration.

Improvements in HP's CMOS process technology, as well as improvements in static RAM speeds, have allowed a 50% increase in the processor operating frequency from 66 MHz to 99 MHz. This provides an almost immediate 50% performance gain.

Independent integer and floating point units, along with a double-word path for fetches, allow superscalar execution of instructions.

The cache and TLB have been improved in several areas. Cache optimizations include a reduction in cache write cycle timing, as well as features such as stall-on-use (the ability to continue instruction execution until missed data is needed) and instruction streaming. The TLB has been improved by use of a hardware TLB walker, greatly reducing the TLB miss latency.

Floating point add and multiply latency has been decreased from 3 to 2 cycles, and floating point load-store bandwidth has increased by 50%.

#### 3.2 Function

There are five major blocks in the processor core: the integer unit, the floating point unit, the unified TLB, the cache unit, and the P-bus interface. A simplified block diagram is shown in Figure 3.

The integer unit contains all integer data path and control including the ALU, the shift-merge unit and the general and special purpose register files. The six-stage integer pipeline is optimized for cache access.

The floating point unit contains the floating point data path and control, and is IEEE 754 compliant. The data path contains four blocks: a double precision ALU, a double precision multiplier, a divide/square root unit, and an eight port 28x64 bit register file.

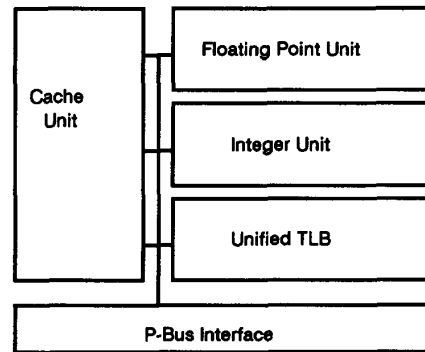


Figure 3: Processor Core Block Diagram

The unified instruction and data TLB has 120 fixed size and 16 variable size entries and is fully associative. The variable entries can be programmed to represent 1/2 MByte to 64 MByte memory blocks. A second level TLB in system memory can be accessed in only ten cycles, greatly reducing the TLB miss penalty.

The cache unit contains instruction and data caches with independent 64-bit data paths. Both caches are direct mapped and have hashed addresses to increase hit rates. Each cache can be read on every cycle for a read bandwidth of 792 MBytes/second. The data cache can be written on every other cycle for a write bandwidth of 396 MBytes/second. Each cache is 32K double-words deep, for a total of 256 Kbytes of instruction cache and 256 Kbytes of data cache. Each cache is parity protected.

The interface from the processor to the system memory and the System Graphics Connect (SGC) bus is via the Processor bus (P-bus). P-bus is a 66 MHz synchronous 32-bit multiplexed address and data bus. The P-bus interface allows the processor to run at 99 MHz on a 66 MHz P-bus. P-bus peak bandwidth is 264 MBytes/second.

#### 3.3 Implementation

All functions in the processor core except the instruction and data cache memory are implemented in an 850,000 transistor custom processor (HP7100). The processor is fabricated in HP's .8 micron CMOS process and is packaged in a custom 504 pin interstitial ceramic PGA. The processor consumes about 25 watts.

The instruction and data cache are implemented in standard off-the-shelf 32Kx8 9 nanosecond static RAMS. Precise cache timing is achieved through careful circuit design and printed circuit board trace delay lines.

## 4.0 System Memory

### 4.1 Overview

The memory subsystem supports both on-board and expansion memory using common control and data paths. Sixteen Mbytes are included on the processor, with connectors for expansion up to 400 Mbytes.

### 4.2 Function

The system memory interface is focused around the primary memory performance contribution: the cache fill/flush characteristics of the processor. The architecture of the memory subsystem, a two bank interleaved design, is found in Figure 4.

The memory subsystem is connected to the processor via the 32-bit multiplexed address/data bus (P-bus) running at 66 MHz, two-thirds the processor frequency. The processor transactions on this bus are of the cache line size (8 words) only. These eight-word transactions are completed as two quad-word DRAM reads or writes.

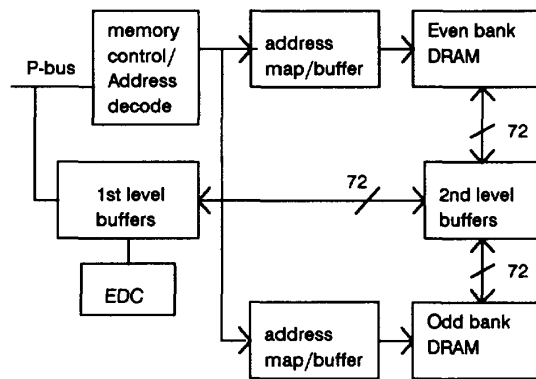


Figure 4: Memory Block Diagram

The expansion memory card for the 735 workstation is a 72-bit wide semi-custom SIMM. Memory cards are available in 8 MByte, 16 MByte

and 32 Mbyte sizes. The 8 and 16 Mbyte cards use 4 Mbit DRAMs and the 32 Mbyte card uses 16 Mbit DRAMs. Cards are added in pairs, one card for each of the two interleaved banks. Each card also includes an address decode/buffer ASIC.

Sixteen Mbytes of system memory is incorporated onto the processor card itself. This allows a system to function without requiring additional memory, and also increases the maximum amount of memory possible in the system. This memory consists of two interleaved 8 Mbyte banks of 4 Mbit DRAMs and an address decode/buffer ASIC.

On the data path between the P-Bus and the DRAM are two levels of buffers. The first level consists of three separate buffers: a DMA buffer, a write buffer, and an instruction pre-fetch buffer. These buffers are used to collapse sequential word transactions into double-word bursts.

The second level of buffers converts these double-word bursts into interleaved quad-word transactions for the DRAM banks. Error detection and correction (EDC) is performed on the DRAM data. The eight bit error correction code is capable of detecting and correcting single bit errors, detecting double bit errors, and detecting some multiple bit errors.

The control portion of the memory subsystem provides several functions including DRAM and buffer control, address mapping and hardware support for graphics acceleration.

DRAM address and control is decoded from the P-bus transaction encoding. Memory address mapping occurs at two levels. First, the transactions are distinguished from the memory-mapped I/O addresses and checked against the limit of actual memory. Second, each memory card checks the transaction address against its location in the memory space.

Graphics features such as Z-buffering, Z-interpolation, and color interpolation are also implemented in the memory control PLA.

### 4.3 Implementation

The first level data buffers, the first level address decoding and the control PLA, including the graphics enhancements, are all implemented as part of a 185,000 transistor ASIC. This ASIC is implemented

in HP's CMOS26 (1 micron) process. It is packaged in a 272 pin PGA and consumes about six watts.

The second level data buffer was implemented using standard off-the-shelf Advanced BiCMOS Technology (ABT) parts.

The second level address mapping and buffering is implemented in a 2400 transistor ASIC. This ASIC was implemented in HP's CMOS26 process and packaged in a 68 pin PQFP.

The on-board memory banks are implemented using standard off-the-shelf 80 nanosecond DRAMs.

## 5.0 System Bus Interface

### 5.1 Overview

The processor card also includes the interface to the SGC, a 33 MHz 132 MByte/second bus. This interface supports the built-in I/O, the expansion I/O via EISA, and multiple graphics masters. SGC includes support for pipelined and burst transactions. An overview for the bus interface subsystem is shown in Figure 5.

### 5.2 Function

The control function of the bus interface provides the translation of the P-bus I/O transactions (outbound) to SGC. The control function also services inbound transactions from I/O and graphics masters to system memory. Outbound transactions include byte and word I/O reads and writes. Inbound transactions include pipelined and burst memory reads and writes.




Figure 5: System Bus Interface Block Diagram

Several system resources, such as arbitration and error support are implemented in the bus interface.

System arbitration is provided for five separate functions: the CPU, the I/O, the EISA interface and the two Graphics interfaces. The overall arbitration scheme is round-robin, with arbitration masking and priority setting available under software control.

Some system error functions such as SGC time-out, memory error, and interruption logging and reporting are included in the control function as well.

The data path for the bus interface converts the 32-bit multiplexed address/data P-bus to the 32-bit demultiplexed SGC. The mux/demux circuit must serve the two-thirds CPU frequency (66 MHz) P-bus on the processor side and the one-third CPU frequency (33 MHz) on the system side. Significant signal buffering is also required on the signals.

### 5.3 Implementation

The data path mux/demux and signal buffers are implemented in a 7000 transistor ASIC. The ASIC provides 16 bits of data path, so two are used in each system. The ASIC is manufactured in HP's CMOS26 process, and is implemented in a 100 pin plastic quad-flat-pack. These ASICs consume about 1 watt each.

The control functions are implemented in the 272 pin Memory control ASIC described previously.

## 6.0 System Clocks

### 6.1 Overview

A workstation operating in the performance range of the model 735 requires high-precision clocking circuitry [7]. The 735 system clock generation is resident on the processor board; its organization is shown in Figure 6.

### 6.2 Function

There are only three active components in the clock generation circuitry: a high performance 396 MHz ECL oscillator, an ECL clock divider/buffer ASIC, and an off-the-shelf ECL/TTL translator.

ECL technology was chosen for the clock system in order to achieve high accuracy and low skews in delivery of clocks to the system's components. Clock signals are driven using differential pairs to assure accuracy of period and duty cycle in the system environment. The clock subsystem delivers clocks at

skews of less than 250 picoseconds to all receiving circuits.

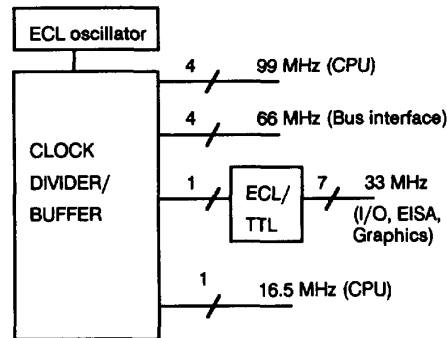


Figure 6: System Clocks Block Diagram

Since the processor, bus/memory interface and System Graphics Connect operate in a 3:2:1 frequency ratio, and all require a 50% duty cycle, the system clock oscillator operates at four times the processor frequency. The oscillator is the single input to the clock divider/buffer chip, which then drives nine pairs of low-skew differential clocks through 50 ohm stripline transmission lines to the VLSI and ECL/TTL translator on the processor board. The ECL/TTL translator drives multiple TTL clocks onto the system backplane to I/O, graphics, and EISA subsystems.

### 6.3 Implementation

The divider/buffer ASIC is implemented in Hewlett Packard's HP-10 bipolar IC process, and packaged in a 44-pin PLCC. The chip contains 530 transistors and dissipates 1.2 Watts.

### 7.0 Board Assembly

The processor card is implemented on an eight by eleven inch twelve-layer fine-line printed circuit board. The physical design is shown in Figure 7.

Well-controlled impedances are achieved by use of multiple power and ground planes in a dual strip-line configuration.

The board consumes approximately 65 watts of power and is cooled with an average of 1-1/2 meters per second of air.

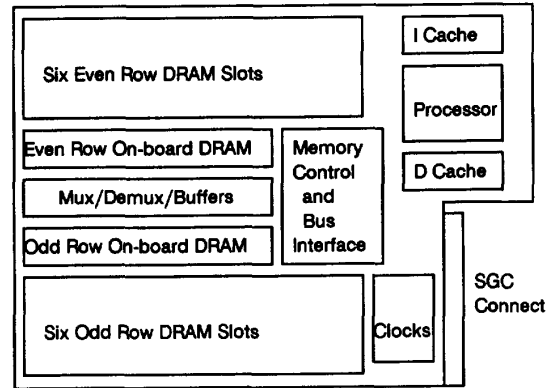


Figure 7: Processor Physical Design

The assembly is manufactured with a double-sided, fine-pitch, infrared reflow process.

## 8.0 Performance

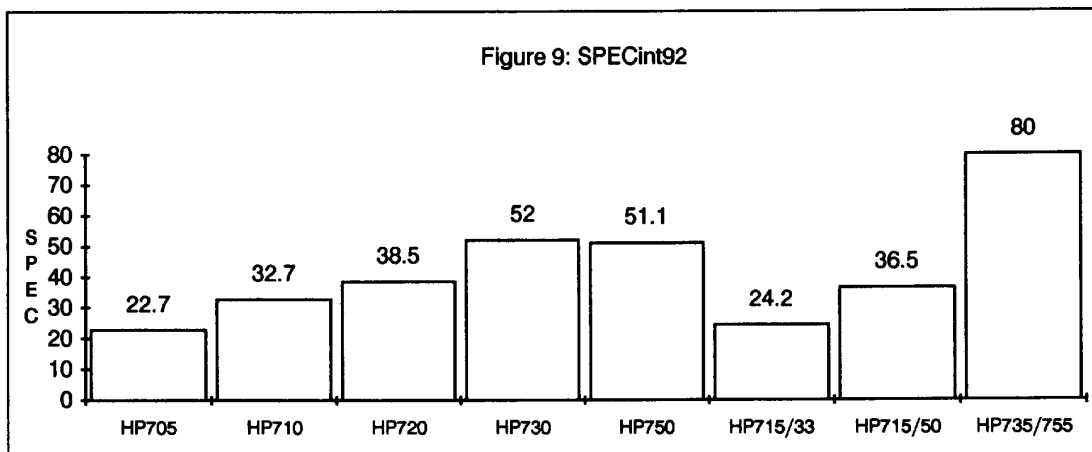

The 9000/735 is HP's highest performance workstation, as shown by the SPECfp92 and the SPECint92 results in Figures 8 and 9 [8]. The HP720, 730, and 750 were the first Snakes workstations introduced in March 1991 [9]. The HP710 and 705 were the low-end members of the family introduced in January 1992 [10]. The HP715, 735, and 755 are the second generation workstations introduced in November 1992. The HP750 and 755 are deskside, server models, while the others are desktop models.

SPECfp92 is SPEC's new floating point suite. It contains fourteen "real world" application benchmarks from a variety of typical application areas. Individual SPECfp92 results for the 735 are shown in Table 1.

SPECint92 is SPEC's new integer suite. It contains six "real world" application benchmarks from a variety of typical application areas. Individual SPECint92 results for the 735 are shown in Table 2.

Other popular benchmark results are displayed in Table 3.

Dhrystone is an integer benchmark designed to represent a programming environment. It typically shows processor and compiler efficiency. The results shown in Table 3 are reported in K dhrystones per second.



Whetstone is designed to represent small engineering or scientific applications. The results shown in Table 3 are reported in Whetstone instructions per second.

Linpack is a benchmark used to represent engineering and scientific applications. The results shown in Table 3 are reported in MFLOPS.

## 9.0 Summary

The first generation PA-RISC workstations, Snakes, has been significantly improved in a new design.

Large performance gains were accomplished by improving processor efficiency and by increasing operating frequency by 50% to 99 MHz. Large memory configurations up to 400 MBytes are supported, with 16 MBytes implemented on board.

This new processor design has extended the industry leading desktop performance of HP's PA-RISC workstation family

| test     | SPECmark |
|----------|----------|
| spice2g6 | 91.9     |
| doduc    | 142.0    |
| mdljdp2  | 192.1    |
| wave5    | 112.1    |
| tomcatv  | 138.0    |
| ora      | 276.9    |
| alvinn   | 176.8    |
| ear      | 258.4    |
| mdljsp2  | 92.3     |
| swm256   | 79.3     |
| su2cor   | 177.2    |
| hydro2d  | 166.1    |
| nasa7    | 123.3    |
| fpppp    | 237.1    |
| SPECfp92 | 150.6    |

| test      | SPECmark |
|-----------|----------|
| espresso  | 92.3     |
| li        | 86.4     |
| eqntott   | 90.9     |
| compress  | 66.0     |
| sc        | 71.7     |
| gcc       | 76.7     |
| SPECint92 | 80.0     |

| test                       | results |
|----------------------------|---------|
| Linpack Single Precision   | 41.2    |
| Linpack Double Precision   | 40.8    |
| Whetstone Single Precision | 158.7   |
| Whetstone Double Precision | 149.3   |
| Drystone 2.0               | 194.2   |
| Drystone 1.0               | 215.5   |

## 10.0 Acknowledgments

As part of the systems team, we would first like to acknowledge the contribution of the VLSI designers in HP's Engineering Lab in Ft. Collins Colorado. The contribution of their back-to-back VLSI designs

to workstation systems success cannot be understated.

Also, we would like to thank all the members of the systems teams of the Entry Systems Lab in Cupertino, California, who have produced two generations of industry leading desktop workstations.

Finally, we would like to thank the management team: Denny Georg, Cliff Loeb, Steve Foster, Kathy Wheeler, Andy Debaets, and Alan Wiemann.

## 11.0 References

- [1] Debaets, A., et. al, "High Performance PA-RISC Snakes Motherboard I/O" Digest of Papers, COMPCON Spring 1993.
- [2] Dowdell, C., Thayer, L., "Scalable Graphics Enhancements for PA-RISC Workstations" Digest of Papers, COMPCON Spring 1992.
- [3] Gleason, C., et. al, " VLSI Circuits for Low-End and Midrange PA-RISC Computers" Hewlett-Packard Journal, August 1992.
- [4] Lee, R. B., "Precision Architecture" IEEE Computer, Vol. 22 No. 1, January 1989, pp 78-91.
- [5] Delano, E., et. al, "A High Speed Superscaler PA-RISC Processor" Digest of Papers, COMPCON Spring 1992.
- [6] Forsyth, M., et. al, "CMOS PA-RISC Processor for a new family of Workstations" Digest of Papers, COMPCON Spring 1991.
- [7] Lettang, F., "ECL Clocks for High Performance RISC Workstations", Hewlett-Packard Journal, August 1992.
- [8] "HP Apollo 9000 Series 700 Workstation Systems Performance Brief" Third Edition, November 1992.
- [9] Horning, R., et. al, "System Design for a Low Cost PA-RISC Desktop Workstation", Digest of Papers, COMPCON Spring 1991.
- [10] Frink, C. R., et. al, "Low Cost Design for a PA-RISC Color Workstation", Digest of Papers, COMPCON Spring 1992.