

A Low-Cost Workstation with Enhanced Performance and I/O Capabilities

Various entities involved in product development came together at different times to solve a design problem, evaluate costs, and make adjustments to their own projects to accommodate the cost and performance goals of the low-cost HP 9000 B-class workstation.

by **Scott P. Allan, Bruce P. Bergmann, Ronald P. Dean, Dianne Jiang, and Dennis L. Floyd**

The design and development of the HP 9000 B-class workstation is a good example of cooperative engineering. In cooperative engineering, the various entities involved in product development come together at different times to solve a problem or make adjustments to their own projects to accommodate a common need. Examples of this cooperation for the B-class workstation include coordination between system designers and firmware developers, the addition of new functionality without impacting the development schedule, close ties with manufacturing, evaluation of implementation based on detailed cost models, and simplification of the PA 7300LC design by moving clocking functions onto a small chip on the system board.

Design Objectives

The design objectives for the B-class workstation were low cost, quick time to market, performance, functionality, longevity, and modularity. In addition to these objectives, the development team's main goal was to produce a workstation based on the PA 7300LC processor that would be comparably priced to the HP 9000 Model 715 workstation, but with superior performance and I/O capabilities. This goal and the design objectives remained the same throughout the project.

With low cost as the primary objective, any feature that was perceived as too costly or of limited value to our customer base was not included. Leveraged subsystems were reviewed in search of creative ways to reduce cost. This led to reductions in the cost of the clock circuitry and firmware interface and elimination of some legacy I/O interfaces. From a cost/performance perspective we were able to justify the addition of a PCI (Peripheral Component Interconnect) bus, a higher-speed memory technology, a second-level cache, and a higher-performance processor and graphics subsystem. Fig. 1 shows the B-class workstation.



Fig. 1. *The B-class workstation.*

Features and Capabilities

Based on the objectives for the B-class workstation, the following features are included in the product:

- PA 7300LC high-performance, low-cost microprocessor with two on-chip associative caches with 64K bytes for data and 64K bytes for instructions
- 1M bytes of ECC (error-correcting code) directly mapped second-level cache for additional performance
- HP VISUALIZE graphics technology from HP VISUALIZE-EG (entry-level graphics)
- HP VISUALIZE-IVX graphics on the B132 workstation (optional)
- Six memory slots that support up to 768M bytes of ECC memory, including fast-page mode (FPM) and extended-data-out (EDO) DRAM dual inline memory modules (DIMMs)
- General system connect (GSC) bus for high-speed I/O bandwidth
- Flexible I/O that includes two I/O slots, which can be configured as:
 - Two PCI slots
 - Two GSC slots
 - One EISA slot
- Optional fast-wide SCSI (20-Mbyte/s) card that supports internal and external disks without using an I/O slot.

In addition to these features, the B-class workstation's modular design provides simple installation, flexibility of use, and easy servicing. This is accomplished through design features such as:

- Simple tray design
- Built-in expandability
- Plug-in memory modules.

Fig. 2 shows a block diagram of the components that make up the B-class workstation.

Processor and System Design

Since the processor chip used in the B-class products is the PA 7300LC, one of the main areas of cooperation was between the PA 7300LC processor design team and the B-class system design team.

The previous-generation processor used in HP workstations of a comparable price was the PA 7100LC. The PA 7100LC was an extremely versatile processor, and many of its best points were leveraged into the PA 7300LC design (see [Article 6](#), [Article 7](#), [Article 8](#), and [Article 9](#)). However, the PA 7100LC was not without its challenges, such as the difficulty in synchronizing the processor clock with the GSC (general system connect) bus.

Clock Frequency

The GSC bus is a general-purpose synchronous bus used to communicate between the processor and I/O. Its phase is determined in relation to a nonexistent GSC clock. This imaginary clock runs at half the frequency of the clock sync signals driven to each GSC device. Its rising edge is defined by the rising edge of reset during initialization, and each GSC device is responsible for keeping track of the current phase of the GSC clock starting from initialization.

On the PA 7100LC, the GSC bus was only permitted to operate at fixed ratios of the processor clock frequency, including some odd clock ratios such as 1.5:1 (see Fig. 3). All of the clock syncs and the resets used to initialize the GSC clock were external to the chip. Designing circuitry to maintain these ratios and timing margins with minimal clock skew and noise immunity became increasingly problematic. In addition, every frequency point of operation required a special clock design to ensure maximum performance. This limited our ability to select the frequency of operation based upon yield at a later point in the design process. For the PA 7300LC, the situation became more critical because the final processor frequency was still uncertain, and the final ratio between the processor frequency and the GSC clock was also undecided.

The first approach investigated was to bring the entire clocking solution into the PA 7300LC. It would be much easier to adjust the delays and control the skew within an ASIC rather than in discrete circuits. The proposal was to incorporate a phase-locked loop circuit within the PA 7300LC to generate the processor clocks from a low-frequency external crystal.

The GSC syncs could then be created by dividing the phase-locked loop output internally in the PA 7300LC. The PA 7300LC would also drive out the reset used to initialize the GSC phase. Upon further investigation, the PA 7300LC design team became concerned about the risk associated with the phase-locked loop. The phase-locked loop was considered a major component of the PA 7300LC design. This was significant because all post-fabrication verification and debugging of the chip would be dependent upon a functional phase-locked loop.

At this point, the B-class system designers and the PA 7300LC design team began to look at a mixed solution. The phase-locked loop was scrapped to avoid risk, and its die area recovered for other uses. The PA 7300LC would continue to drive the primary synchronizing reset to eliminate the need to synchronize the asynchronous power-on reset to the GSC's syncs. The generation of the syncs and the maintenance of their skew requirements would be moved to an external ASIC.

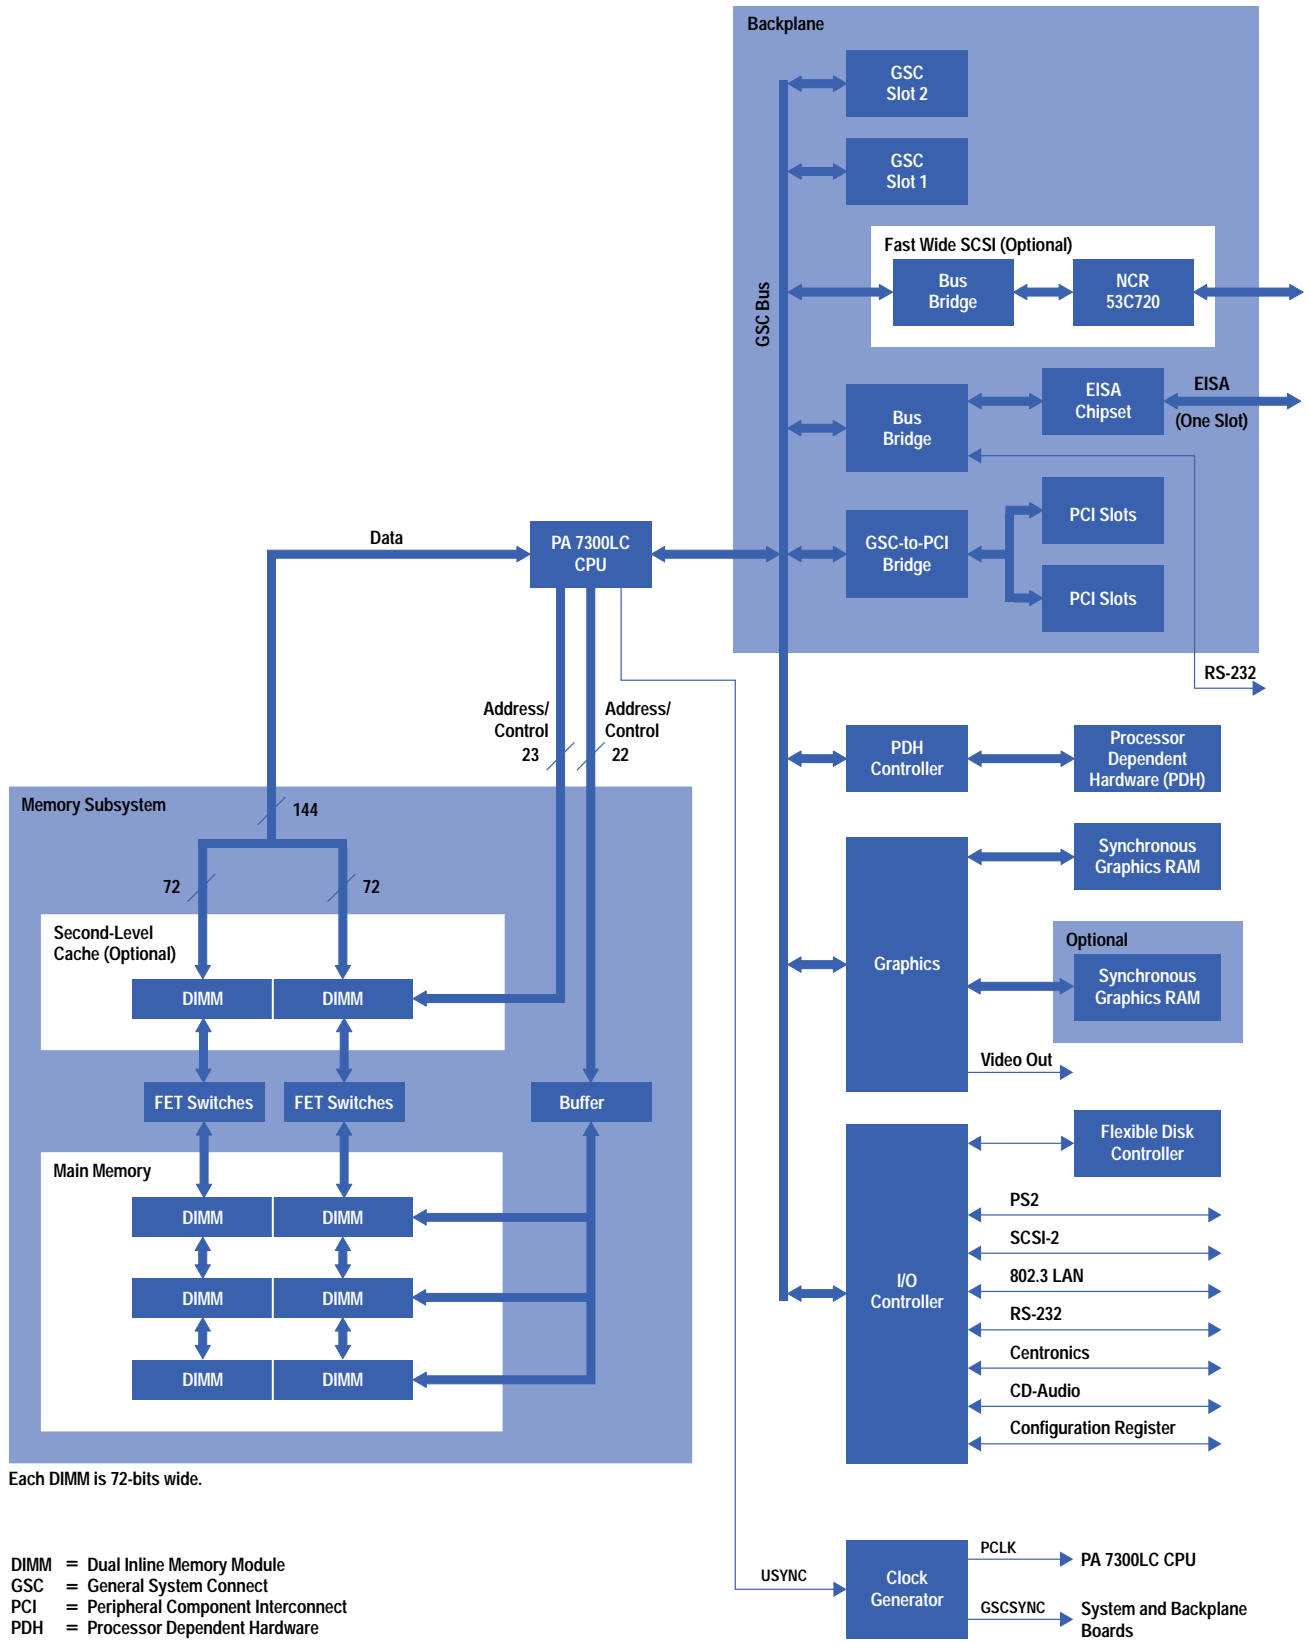


Fig. 2. The system block diagram for the B-class workstation.



Fig. 3. The PA 7100LC clocking scheme.

Any necessary turns to a small ASIC would be quicker and less expensive. In addition, the clocking solution could be completely bypassed to allow continued verification and debugging of the PA 7300LC if necessary.

Working with Motorola, the PA 7300LC design team, and our materials organization, the system designers specified the device that became the Motorola MPC992 (the clock generator in Fig. 2). This device uses a phase-locked loop and an external low-frequency crystal to generate differential clocks that provide clocking to the processor and the other GSC devices. As an added benefit, its cost is relatively low in relation to the external clock oscillator and ECL devices used in previous products. The USYNC signal, which comes from the PA 7300LC processor, is the synchronizing signal that is responsible for aligning the GSCSYNC with the processor clock signal.

Memory and I/O Controller

The proximity and working relationship between the PA 7300LC and B-class system design teams allowed us to communicate design specifications with relative ease. This working environment allowed us to view the product as a whole rather than designing the system around an existing chip.

The design of the memory and I/O controller (MIOC) was the first area affected by this arrangement. The PA 7300LC is designed to support optional second-level caches of different technologies and sizes. When the PA 7300LC chip design team began investigating each of these second-level cache options, the B-class system designers were able to check the appropriateness of their solution with the design. One of the first decisions under this arrangement was to make the second-level cache optional and locate it on a DIMM (dual inline memory module) on a separate board. This provides the B-class workstation with several benefits:

- Lower-performance systems are not burdened with the cost of the second-level cache.
- Systems with and without a second-level cache can share system boards, reducing development and verification time.
- The exact configuration of the second-level cache can be altered at a later date if market conditions warranted.
- Less space is required on the board, permitting a lower-cost system board.

It was important for the PA 7300LC design team to know that a DIMM solution was being considered since it would have a big impact on the I/O pad design of the PA 7300LC.

Another area of concern within the MIOC involved the impact of the expanded data bus on the PA 7300LC (144 bits) compared to the PA 7100LC (72 bits). This would require additional pins and incur additional packaging costs. The PA 7300LC design team wanted to share the memory data bus with the second-level cache data bus to reduce the number of external I/O pins. However, the additional load associated with the memory would degrade the response of the second-level cache. The PA 7300LC design team suggested FET switches, which could be dynamically opened and closed to isolate the second-level cache from main memory.

The B-class system designers were able to verify using FET switches in a system environment. However, the only devices available that met the enable/disable speed requirements were 8-bit devices. This was viewed as an unwieldy and expensive solution in the B-class system. Working with our materials organization and Texas Instruments, the B-class system designers were able to make minor specification changes to an existing 24-bit Texas Instruments part to improve this speed parameter and cut the quantity and cost of the FET switches significantly. The B-class system designers verified the signal quality of the memory data and second-level cache data of these devices in a system environment.

As the configuration of the second-level cache solidified, the B-class system designers were able to provide the PA 7300LC design team with specific information concerning the electrical environment in which the PA 7300LC would be operating. With this information they were able to run simulations of their I/O pad drivers operating within the actual system. This led to some changes in their pad designs, eliminating potential problems later.

Memory

As with most projects, the PA 7300LC design team and the B-class system designers had their share of resource shortages. One such issue involved the memory family. The PA 7300LC is designed to support both fast-page mode and extended-data-out DRAMs. In fast-page mode, sequential data is driven from the DRAMs on successive column addresses following a single row address, rather than requiring both the row and column address to be driven on each data access. Extended-data-out DRAMs are an enhancement to fast-page mode DRAMs in which the data remains valid until the column address changes or a new column address strobe occurs, rather than becoming invalid when the column address strobe disappears. This allows a longer time period over which to latch incoming data and saves processor states in memory accesses.

Unfortunately, resource conflicts and schedule constraints made it impossible for the PA 7300LC design team to verify functionality of the chip for both memory technologies. The PA 7300LC design team wanted to qualify the extended-data-out DRAM technology because it would provide a higher-performance memory technology. The B-class system design team wanted the fast-page mode DRAM technology qualified to be compatible across the workstation family, rather than having a unique memory component for the B-class systems. The compromise solution was to have the PA 7300LC design team qualify the fast-page mode DRAM technology for first release. At a later point in the design phase, the B-class system designers would qualify the operation of extended-data-out mode DRAMs to be introduced as a performance enhancement.

Data Capture

Resource balancing was also evident in the development of a data capture board for the PA 7300LC. A data capture board is a device that is attached to a system board and is used to observe the high-frequency signals between the processor, second-level cache, and memory for debugging purposes. Since the B-class system designers were more familiar with board design tools and the board design environment, the B-class system design team developed the data capture board for debugging the PA 7300LC.

Hardware and Firmware Trade-offs

Design teams frequently look at trade-offs between optimizing resources and meeting the goals of the team. For the B-class workstation, the hardware and firmware teams fostered a close working relationship, allowing trade-offs to be made on a broader scale.

A most unusual but significant outcome of this close working relationship was the development of an unplanned ASIC for interfacing to the processor dependent hardware (PDH). The PDH consists of components such as the boot ROM, nonvolatile memory, and configuration registers. Although there was already a way to connect to the PDH functionality through part of the core I/O logic being leveraged from previous lower-end workstations, this interface did not provide the level of functionality that was implemented in the higher-end workstations. The firmware team could save significant resources by leveraging portions of code from the C-class workstation and the higher-end members of the D-class server family. Many of the basic I/O and graphics functions were similar between these platforms. However, the code leverage was predicated on having certain PDH functionality that could not be provided with the low-end solution. In addition, the high-end solution provided superior debug capabilities. These better debug capabilities were very attractive to help ensure a speedy startup of the new PA 7300LC processor, and hence help meet our time-to-market goals.

The key capability missing from the PDH interface used in previous lower-end workstations was the ability to perform word-wide write accesses to PDH devices. The PDH interface was optimized for reads, with only byte-write capabilities provided. The new PDH ASIC added the word-write capability to support a scratch RAM. This seemingly innocent scratch RAM was key, because in high-end workstation code it is used as a stack in the early stages of the boot process before main memory is initialized. The scratch RAM is also used for global information such as tables of I/O and graphics configuration information. It would have been very difficult to leverage code with the word-write capability to a platform without this capability.

The new PDH ASIC also provided additional address decode and the appropriate flexibility in timing to allow the direct connection of a serial port into the PDH hardware. This direct connection to a serial port, in conjunction with the capabilities offered with the scratch RAM, allowed a debugger to be operational even with hardware that was minimally functional. This serial port aided code and hardware debugging by allowing the hardware status to be monitored and the hardware configuration to be modified early in the boot process.

The risk for the new PDH ASIC was minimized by incorporating it into system simulation efforts and by keeping the design focused on the needed functionality and disallowing any unnecessary features.

Product Definition

The B-Class system was originally defined alongside the C-class workstations. The B-class system is essentially a smaller version of the C-class workstation. Our original intention for the B-class implementation was to use the same modular philosophy of separate I/O, CPU, disk interface, and human interface subsystems used in the C-class machines. However, when the time came to implement the B-class product, cost goals had become more important. When preliminary costs were evaluated, it became clear that we were not meeting the cost objectives with the existing product definition.

Many alternatives were generated and evaluated against product objectives. Finance and R&D reviewed their cost models to see where costs could be saved. Manufacturing reviewed the design alternatives for manufacturability and analyzed the supply chain for issues associated with parts procurement, assembly, material, and structure. Service was consulted to review serviceability and warranty implications of the various options, as well as issues with potential future upgrade products. The result of this analysis was a single-board integrated computer (see Fig. 4). The design, which was initially spread out over four separate boards in the C-class system for the sake of modularity, was now integrated onto one system board.




Fig. 4. The B-class system board and its components.

Single-Tray Concept

Like the C-class workstation, the B-class workstation uses a tray concept. However, instead of two trays (one for the disks and one for the boards), there is one tray that holds everything (see Fig. 5). For this reason, during the design phase it was important to consider keeping the weight down. Holes were added in the tray wherever possible to reduce the overall product weight.




Fig. 5. The main tray assembly.

EMI

The tray assembly slides into a metal can. With this approach, the EMI (electromagnetic interference) interface is limited to the perimeter of the rear panel. Once the tray is removed, there is easy access to the option boards, memory modules, second-level cache modules, optional fast-wide SCSI interface board, power supply, disk drives, speaker, fans, and the CPU chip. The system board is accessible by removing the disk bay, which is secured by only one screw and a few cables.

Disk drives can be accessed without removing the disk bay from the main tray simply by removing the snap-on cover. The disks are mounted using plastic brackets so that they can be changed without tools (Fig. 6). A fan was added to the bottom of the disk bay to provide enhanced disk cooling since successive generations of disks consume more power. Removing the backplane is slightly more difficult, requiring all modules to be removed first.




Fig. 6. *The disk bay.*

Manufacturing

Working with manufacturing included performing a supply chain analysis¹ as part of the total system cost analysis. Design efforts produced detailed material lists that were used to determine an overall system cost. Several design scenarios were developed with mechanically exploded models and material lists. The cost model for the B-class workstation was not limited just to the material content of the product, but also included system interconnect costs, parts procurement costs, part placement costs, printed circuit board electrical and system functional testing costs, and system support costs.

Manufacturing and field-support representatives were involved in defining the system for manufacturability, inventory control, and configurability to reduce the system cost. Design scenarios were then evaluated against each design objective.

Initial prototypes were assembled and disassembled by manufacturing personnel to provide a hands-on critique of the designs. These inputs were fed back into the design in the early stages of development.

Serviceability

One of the challenges of this single-board solution was to make the system board accessible for service. We wanted to have the board slide in and out, but there were connectors and switches on both edges of the board. In addition, the connectors had to be accessible through the rear panel. To allow the board to slide into the package we added a small tray to the bottom of the system board that could slide along card guides. One of the design requests from service representatives was to be able to service the system board without removing the rear panel. To accomplish this, the rear-panel connectors were recessed and a separate small bulkhead attached with a sliding EMI interface to the rear panel. This bulkhead remains attached to the board in a board replacement.

Another serviceability concern was the alignment of the power button, mute button, volume knob, audio jacks, and LEDs located on the front of the product. In the chassis, the main tray engages alignment pins, which serve to lock the tray to the chassis during vibration and shock. Because of the tolerance stackup from the front panel through the can, tray, backplane, system board, and all the connectors and buttons, we were concerned that the cosmetics at the front would be unacceptable. To improve alignment, we mounted these connectors on a long, thin section of the printed circuit board that would flex and be supported by a metal brace so that the front section could move relative to the rest of the board. We added aligning forks to the front panel to position just that section of board. With this method, we were able to locate these connectors accurately.

Manufacturing also assisted in improving the design through their participation in design reviews. One suggestion led us to abandon the captive rear-panel fasteners that we had been planning to use. If the captive screws are not properly aligned, they can be cross-threaded and stripped, or the captive nut on the chassis may be damaged. Consequently, the whole tray would need to be replaced just for a simple nut or screw. Instead, we designed custom thumbscrews with an unthreaded nose section to align the screw before the threads engage. This minimizes cross-threading. To save labor costs we also used a coarse thread to reduce the number of rotations necessary to remove and install them.

Another goal was to reduce the number of screw types. We tried to standardize on a single screw used in our earlier option boards because this was the one screw that we could not change. We used it to attach the power supply and the disk bay. To reduce screw count, the main fans and speaker snap in place. The backplane slides in place with keyhole standoffs and forks in the main tray. When the power supply is installed, two pins from the power supply trap the backplane in place. The power supply is supported with two screws and the two pins that are routed through the backplane into the backplane support. The

power supply has floating connectors so that stresses from vibration and shock are not transmitted between the backplane and the power supply via the connectors.

One of the primary objectives was upgradability. Upgrades can be easily accomplished by a simple swap of the system board. Since everything is on one board, there are no issues with incompatibilities between different versions of the I/O and CPU. The small I/O bulkhead stays with the board so the main tray assembly need not change. Sufficient extra height remains where the CPU and memory are located so that future high-power CPUs have room for larger heat sinks or even small daughter boards if more board real estate is needed.

Processor and System Verification

The verification effort for the PA 7300LC and the B-class and C-class products was also a joint effort. Shmoo tests were conducted simultaneously on both the B-class and C-class workstations.

A shmoo test is designed to verify the product under voltage, temperature, and frequency extremes. Its intention is to electrically stress the system under test to within and beyond its operating limits. This process is part of our electrical characterization of the processor and system. A shmoo test is an important part of our product development cycle. By pushing the system to its electrical extremes, we hope to reveal any design weaknesses that could affect the operation and performance of the system under extreme operating conditions. It often uncovers weaknesses in both chip and board designs. These might include signal cross talk, chip-drive capability, slow-speed paths at high temperatures, or board-level clocking problems.

To achieve superior product quality, both processor and system shmoo tests were performed on B-class systems. The processor shmoo test focused on the core processor, caches, memory, and GSC bus. The system shmoo test emphasized peripherals and I/O, including the expansion I/O on the GSC, EISA, and PCI buses.

Since the PA 7300LC was designed to work in both B-class and C-class systems, it was tested in both systems. Processor characterization was performed in the C-class systems by the PA 7300LC design team. Simultaneously, the B-class system design team completed the processor shmoos in a B-class system. Both the B-class and C-class system design teams completed system shmoos with the PA 7300LC in their respective environments.

The parallel verifications of the PA 7300LC in the B-class and C-class systems complemented each other, providing opportunities for leveraging and making the debug process go smoother. One of the issues discovered during the processor shmoo test was the limited operating frequency of the GSC bus. This was caused by the length and load on the bus and a threshold problem on the PA 7300LC. The combined efforts of the PA 7300LC processor and B-class system design teams extended the operating frequency of the GSC bus in our systems and provided the desired performance. The PA 7300LC design team corrected the threshold problem and the B-class team shortened the GSC bus, which slightly changed its characteristic impedance and helped to alleviate the problem.

Processor-level electrical verification has three main goals: uncover electrical (nonfunctional) bugs in the system, find critical speed paths that limit the maximum frequency of the processor, and provide correlation between the IC tester frequency and the eventual system frequency. The third goal had the biggest impact on costs. As development progressed, it became obvious to the PA 7300LC and B-class teams that the frequency mix (132 MHz to 160 MHz) between the IC tester and the system was not meeting marketing requirements. The correlation effort between the teams uncovered ways to enhance the system electrical and thermal environments to bring the yield mix and market demand together. The close cooperation between the two teams enabled the quick identification of a solution to the problem. We made alterations to the system's thermal cooling environment, allowing us to run the PA 7300LC at a higher frequency, something we could not do in the original cooling environment.

Over the years, many efforts have been made to address and improve the shmoo test process at both the processor and system level. While processor shmoo testing reveals many system level problems, its primary focus is still the processor, cache, and memory subsystems, rather than the I/O subsystems. As I/O bus speed and peripheral interface IC complexity has increased, it has become more important to address the I/O subsystems in shmoo testing. The PA 7300LC was designed to make complete system shmoos more practical for this reason. The clock circuitry for the PA 7300LC was designed to permit overriding the nominal clock frequency while maintaining the correct synchronous relationship between the processor and I/O clocks. This allowed us to vary the frequency of operation more easily over a larger range of operation than in past products.

One of the challenges for system shmoo testing in B-class systems was the range of new system components that had to function correctly together during testing. As with the processor shmoo, system testing attempted to stress the electrical design of the new components by operating them under extremes of temperature, voltage, and frequency. In addition to the core I/O components, various expansion I/O cards were selected to verify complete system functionality.

The extensive system shmoo testing of the B-class system led to the optimization of several circuits and resulted in a higher-performing, more robust system. We have come to believe that shmoo tests are an indispensable part of our product development. Besides helping to catch potential problems before introduction, shmoo tests also make post-product support and maintenance easier.

Conclusion

Cooperative efforts between many functional areas such as manufacturing, service and support, marketing, firmware development, and the PA 7300LC chip development team together with the electrical and mechanical system designers have produced the B-class workstations. The closely coupled system design approach has yielded a workstation that provides significant value to our customers.

References

1. G. A. Kruger, "The Supply Chain Approach to Planning and Procurement Management," *Hewlett-Packard Journal*, Vol. 48, no. 1, February 1997, pp. 28-38.
-
-

- ▶ [Go to Next Article](#)
- ▶ [Go to Journal Home Page](#)