HEWLETT-PACKARD JOURNAL

SEPTEMBER 1986

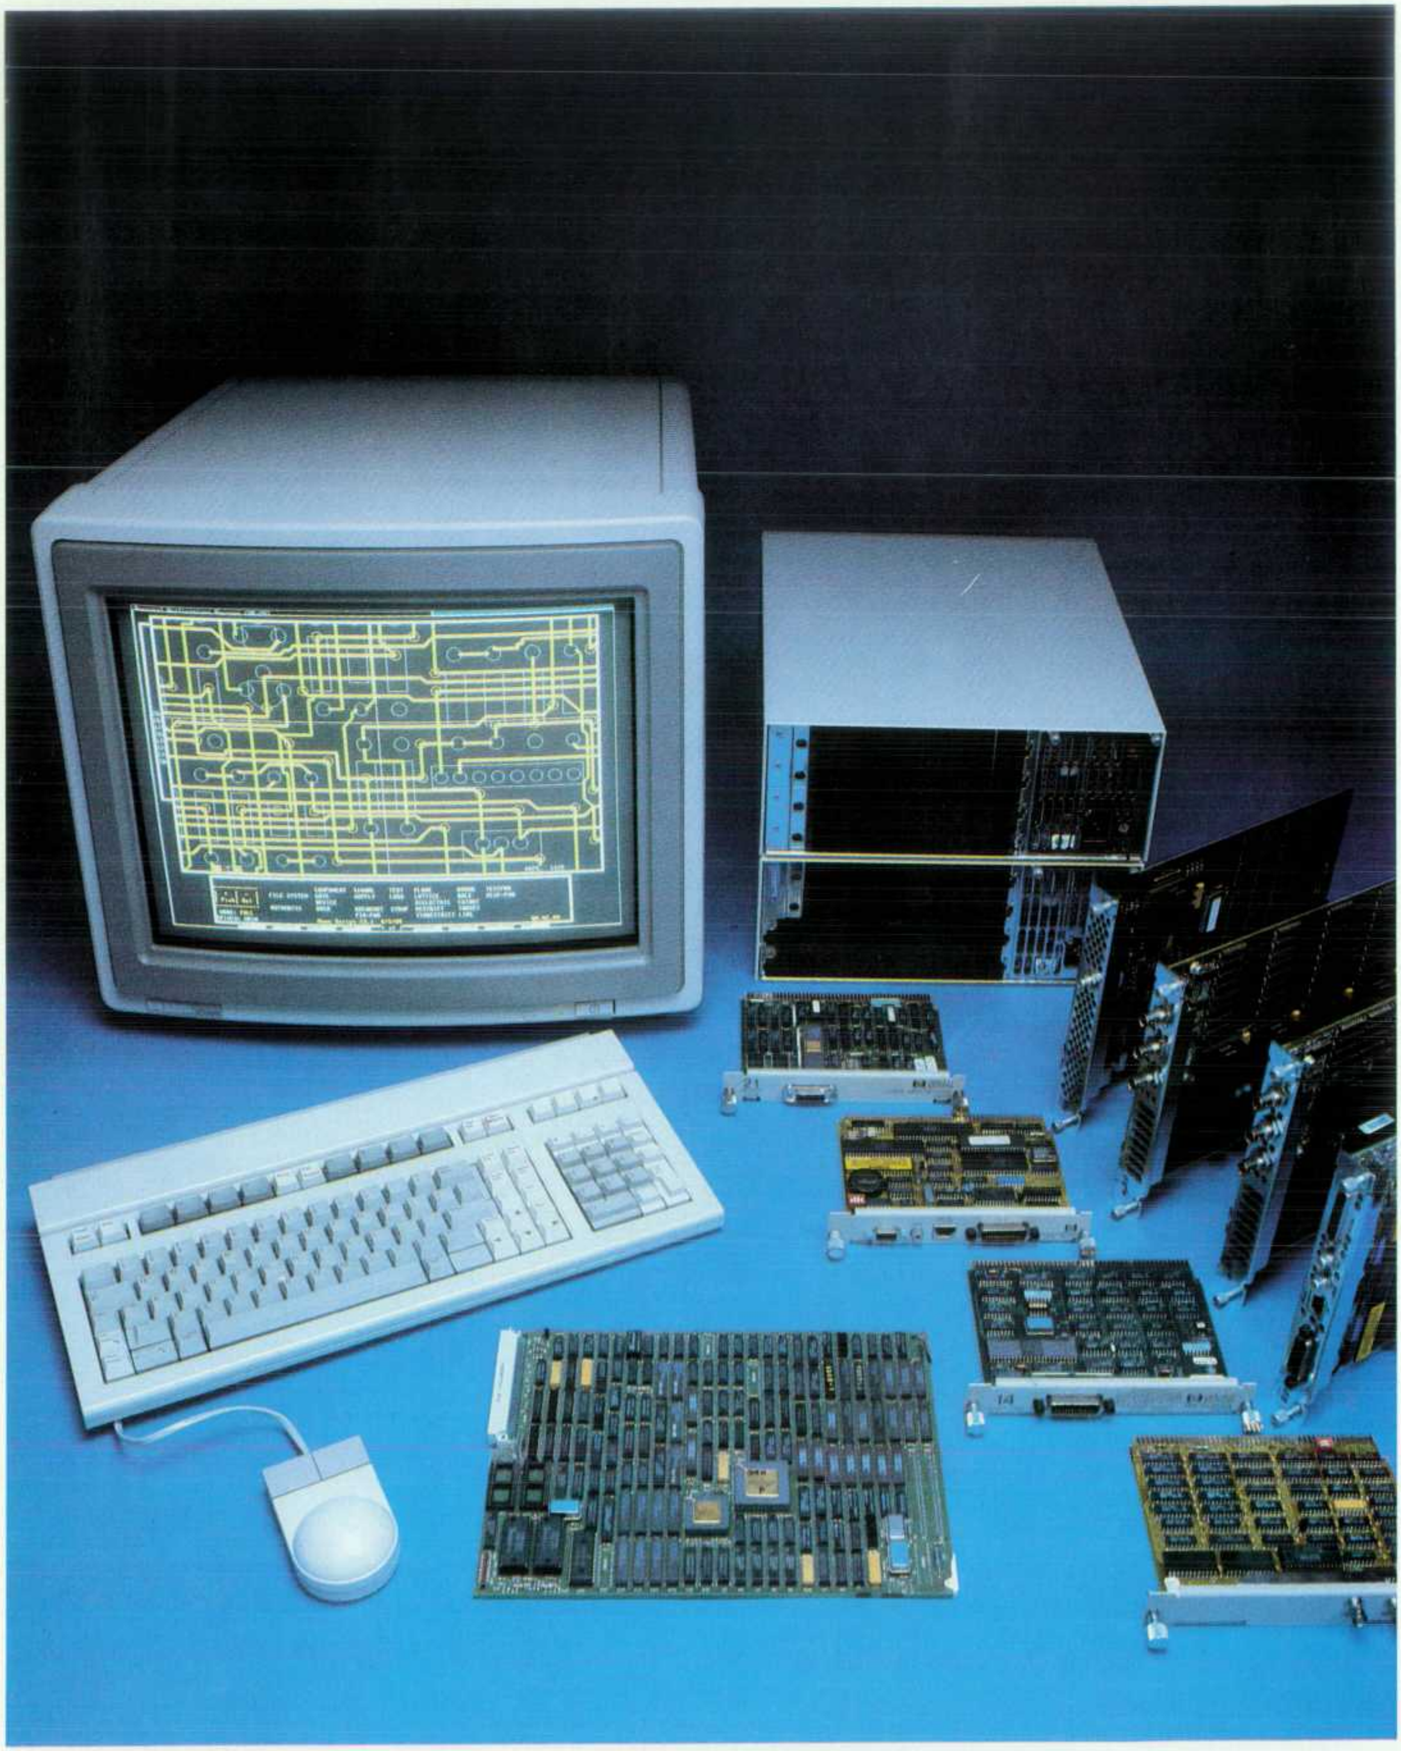

HEWLETT-PACKARD JOURNAL

September 1986 Volume 37 • Number 9

Articles

Departments

In this Issue

Engineering workstations are computers that are found on engineers' desks, but personal computers they are not, except for the superficial visual resemblance. A case in point is this month's cover subject, the HP 9000 Series 300 family of modular workstations. As the cover photo suggests, the Series 300 offers a range of options that would bewilder the typical PC buyer. Designing the Series 300 Computers and the production process so that not only does each customer get a correctly configured computer, but also any option can be installed later in the field, was a significant engineering challenge that's discussed in the article on page 4. The same article tells how the Series 300 came to its present definition, starting out as a proposed low-cost member of the Series 200 family and evolving into a modular replacement for the entire Series 200 that offers performance far beyond any Series 200 Computer. The performance of the Series 300 Computers begins with the Motorola 68010 and 68020 microprocessors and the 68881 floating-point coprocessor for the 68020. Model 310, the lower-performance member of the family, uses the 68010, which has 16-bit address and data buses and does 32-bit operations. The design of the Model 310 processor board is the subject of the article on page 9. Model 320, the higher-performance Series 300 family member, uses the 68020 and the 68881, which are full 32-bit processors. Model 320's design is described in the article on page 12. The issue of porting Series 200 software to the Series 300 was a major design concern that is addressed in the article on page 22. Powerful graphics capabilities, which are mandatory for the CAD/CAE applications that are an engineering workstation's daily fare (the display in the cover photo shows a printed circuit board design application), are provided by a bit-mapped display based on two custom VLSI chips. In the article on page 17, the chips' designers tell us how they dealt with requirements to support both monochrome and color monitors, each with either medium or high resolution, while maintaining compatibility with the Series 200, which doesn't have a bit-mapped display.

Electronic mail, along with word processing, spreadsheets, and graphics, is one of the basic services expected of an office automation system. Here at the *Hewlett-Packard Journal*, we began to use HP's own electronic mail product, HP DeskManager, early in 1983. Today, we'd be hard-pressed to get along without it. It has cut the time it takes for written communication with Europe or Japan from about a week to overnight. It is the way we prefer to receive article manuscripts, because we can copy them directly to HP's word processing product, HP Word, for editing, and then either mail them right back to the authors for review or dump them to our typesetting system, saving the multiple retypings that we used to have to do. HP Desk (we prefer the shorter name) is integrated with word processing, spreadsheet, and graphics software so that files created using those products are easily mailed electronically to any of the more than 50,000 HP Desk users worldwide. According to published reports, that number of users makes HP's network comparable in size to the larger commercial electronic mail networks. The paper on pages 30 to 48 is a comprehensive discussion of the lessons learned in implementing the HP network successfully. Topics discussed include strategy, specific tactical advice, and potential pitfalls. The paper should be valuable to any organization implementing, planning, or thinking about an electronic mail system.

-R. P. Dolan

What's Ahead

Computer networking will be the theme of the October issue. Topics will include HP AdvanceNet, which is Hewlett-Packard's overall networking strategy, and various networking products and services for HP 1000, HP 3000, and HP 9000 Computers.

The HP Journal encourages technical discussion of the topics presented in recent articles and will publish letters expected to be of interest to our readers. Letters must be brief and are subject to editing. Letters should be addressed to: Editor, Hewlett-Packard Journal, 3200 Hillview Avenue, Palo Alto, CA 94304, U.S.A.

Advanced Modular Engineering Workstations

This workstation system allows the user to choose the processor, display system, memory, interface cards, peripherals, and operating system most appropriate for the application.

by Gilbert I. Sandberg, Daryl E. Knoblock, John C. Keith, Michael K. Bowen, and Ronald P. Dean

THE HP 9000 SERIES 300 is a modular, high-performance, technical workstation family (Fig. 1) that can be configured to meet the needs of a wide range of technical applications. An engineer or scientist can choose from two SPUs, six displays, six I/O slots, and a wide range of input devices to meet exact needs, and can later upgrade the workstation in any of the options with only a few minutes of assembly effort. This article discusses the impact of such a large choice of options on the Series 300's development.

Development History

The development of a new low-cost member of the HP 9000 Series 200 Computer family was first proposed in 1981. Using a high degree of both physical and logical integration, this machine was to sell for only $2000 (U.S.). It was proposed that this computer consist of only two printed circuit boards enclosed within a video monitor: a monitor electronics board and a single-board SPU (system processing unit) tapping its power from the monitor board. Adding a keyboard and a disc drive would produce a complete system capable of running the HP-UX operating system, HP's enhanced version of AT&T Bell Laboratories' UNIX™ operating system. As the required logic was analyzed, two areas were singled out for proprietary LSI development: the address translation subsystem needed to run HP-UX, and the video display subsystem.

At that time, there was a growing awareness of the system advantages of a new type of display subsystem called a bit-mapped display. This type of display subsystem holds text character images in the same memory used for graphics pixels, rather than generating them on the fly from ASCII characters stored in a separate memory. Although there is far more flexibility in the generation and use of character fonts, much additional hardware is needed to maintain the performance of character placement and scrolling. Although this type of display subsystem seems inconsistent with a low-cost product, it is far more easily integrated into an LSI chip than its predecessors, because the memory for storing ASCII character values and the ROM needed to convert them to graphics images on the fly can be elimi-

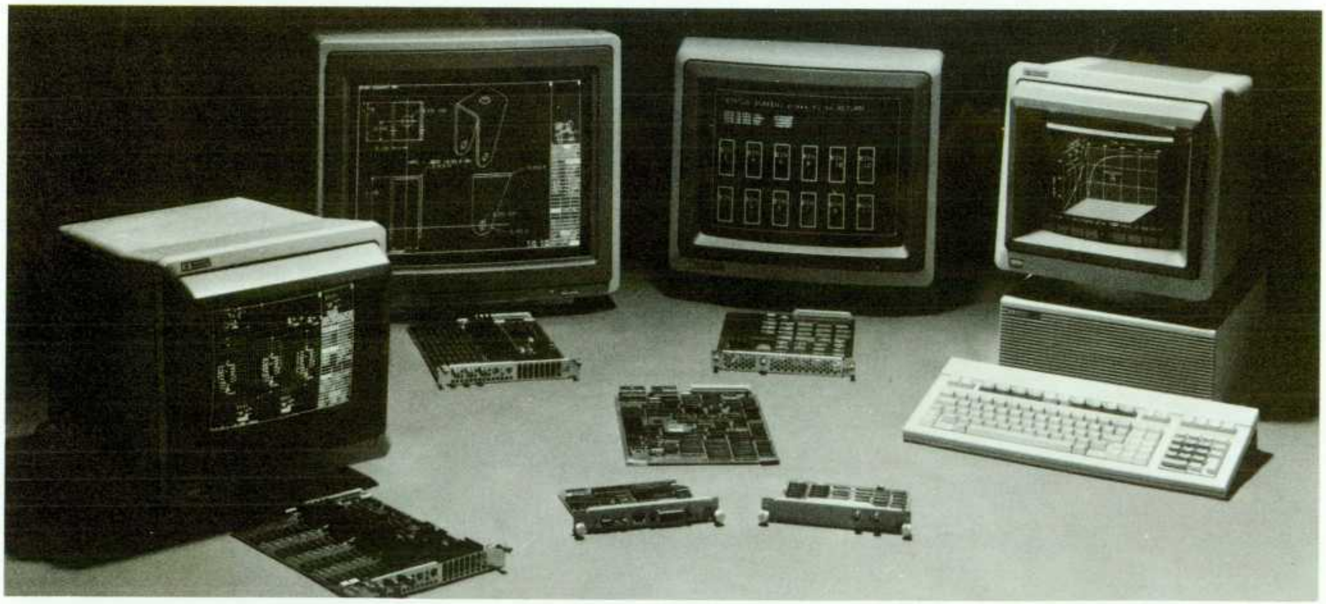

Fig. 1. *The HP 9000 Series 300 is a modular family of high-performance technical computer systems designed for instrument control and CAD/CAE applications. A variety of processors, graphics displays, I/O options, mass storage devices, and other peripherals are available to tailor a system for a particular application.*

nated. Therefore, it was decided to develop such a chip, based on one that had already been developed at HP's plant in Corvallis, Oregon, but capable of higher resolution, color graphics, and much faster character manipulation.

The address translation subsystem (called a memory management unit, or MMU) also began development at this time, and these two LSI efforts essentially drove the schedule for the first two years. But time has a way of changing the factors that define a project, and projects that contain large LSI efforts are particularly vulnerable because of their long duration. Two such factors had a dramatic effect on the Series 300 definition: advances in monitor technology and the development of a high-performance 32-bit microprocessor.

By 1984, Hewlett-Packard was developing a standard, uniform packaging strategy, known as the ITF family, for use within most of its computer systems. The main benefit of this concept is the ability to leverage modular components designed and manufactured at other HP facilities and still produce a system that looks as if it were designed by a single team (which in fact it was, although the team was spread over a large geographic area). This family includes low-cost color and monochrome monitors that match the needs of the Series 300 product, and because of their high production volumes, result in a lower total system cost than the original integrated product definition. As a result, the Series 300 was redefined to be a modular rather than an integrated product.

With this change to a modular definition, striking similarities showed up between the Series 300 and a parallel effort to define a new high-performance member of the Series 200 family based on Motorola's proposed 68020 microprocessor. In May of 1984 it was proposed to merge these two projects into a single project which offered in one product a choice of two processor boards: a single-board computer capable of running the UNIX operating system based on the Series 300 LSI technology and Motorola's 68010 processor (see article, page 9), and a high-performance three-board computer based on the full 32-bit 68020 processor (see article, page 12). This product proposal, which included a choice of four displays and I/O expansion from zero to twelve I/O slots, was capable of being configured to match the characteristics of every member of the existing Series 200 family.

With a single product, HP could consolidate its manufacturing of technical computers, offer better price/performance ratios in each market segment served by each member of the Series 200, and allow customers to tune a machine to their exact needs. In addition, marketing techniques and design features were proposed to allow customers to upgrade their processor, display, or I/O subsystems at any time after their original purchase of a Series 300 model.

Achieving Completeness

As the concept for the HP 9000 Series 300 Computer family evolved into the goal of completely replacing the earlier Series 200 family of computers, the task of providing a complete solution grew rapidly. HP's Fort Collins Systems Division (FSD) had completed many major successful programs in the past,[1,2] but never before had the division worked on a program to replace a complete family of computers. There were five different models in the Series 200 family that had to be replaced at one introduction event. The new family of computers had to be able to cover the complete range of functionality that the previous family provided, and, ideally, do it at a lower cost and with improved performance. At the same time, our market direction was being expanded to include design automation (CAE/CAD) as well as the traditional customers served by the Series 200. This placed additional requirements on providing a complete solution.

The most noticeable requirement imposed by the CAE market is in the display offering. Although the high performance needed for CAE applications is available from the 68020 processor, the medium-resolution color and monochrome displays selected at that time to replace existing Series 200 capability did not provide enough resolution for the detailed graphics images required for CAE displays. Therefore, high-resolution color and monochrome displays were added to the family definition.

Since the medium-resolution monochrome and color displays were developed at HP's Roseville Terminals Division, a major area of responsibility and needless diversion of resources was avoided at FSD. Only one part-time engineer was needed to act as a liaison with the monitor design group to ensure that these monitors would interface satisfactorily with the Series 300.

FSD already had displays that met the high-resolution requirements of the CAE marketplace. All that was required to use these displays for the Series 300 was to design the appropriate hardware driver boards and the custom integrated circuits (see article on page 17) that were necessary to reduce the physical size of these boards so they could fit in the product package.

The keyboard design was done in a similar fashion, this time working with the personal computer group in HP's Singapore facility. Keyboard design and manufacturing activities have been centralized in Singapore to maximize volume and make efficient use of resources for tooling and manufacturing robots. FSD uses the ITF keyboard from Singapore[3] on the earlier Series 200 Model 217 Computer. This same keyboard is used on the Series 300 without changes. Again, all that was required in the form of resources from FSD was someone to act as a liaison with Singapore, since although the keyboard was already in production, design changes continue to be made in the keyboard to improve reliability and customer satisfaction.

The main system processor unit needed to continue the concept of extensibility and flexibility. To achieve this, two CPU boards were designed to fit into the same system slot. With the 68010 single-board computer described in the article on page 9, many of the configurations of the earlier Series 200 family could be replaced with a unit that had higher performance and lower cost. However, by substituting the 68020 processor board with high-speed cache memory and a full 32-bit-wide memory bus described in the article on page 12, twice the performance is provided for complex applications.

The ability to incorporate extensive I/O capability was important for the measurement automation applications that are served by the Series 200, yet posed significant cost burdens at the low end. The maximum number of I/O slots of the SPU package had to overlap the number in the earlier Series 200 family, but also needed to match the low cost of the least-expensive Series 200 model. Since the 68010 processor board has most of the I/O capability required by many customers, this problem was solved by making an I/O card cage that is removable from the basic system package. Then, for those applications that require the high-performance 68020 SPU, which has no I/O capability and also frequently requires large amounts of memory, a passive I/O expander was designed. This expander fits on top of the 68020 SPU, and has mechanical and electrical connections that can be made by the customer with no special tools.

ID Module

Although only the major system needs in providing a complete family have been discussed, there were several other needs that, while smaller in scope and complexity, are just as important in terms of completing the functional replacement of the earlier Series 200 family. One such need is the ability of the Series 200 processors to return, upon command by the software operating system, a unique serial number. This feature was provided so that application programs that had the capability of checking this variable could be made to restrict their operation to only those machines that had had their serial numbers encoded in that software. This is important to prevent unauthorized use of application programs that have been supplied by third-party or independent software vendors.

As the development of the two processor boards proceeded, it became obvious that with the complexity and the physical size limitations of those two boards it would be impossible to include this self-identification feature on the Series 300 processor boards. This created a problem, since there are no other boards that always reside in the mainframe that could be called upon to host this function. What evolved to solve this problem is a small package known as the ID Module, which interfaces with the machine via the same mechanism as the keyboard. This ID Module also has the capability to return a unique number (actually its own serial number) to the software operating system for the same use as mentioned earlier. Having this capability in a separate package from the SPU, rather than integrated with it, provides two benefits. First, being portable, the ID Module can be moved from one machine to another. Thus a customer who legitimately has authorization to use a certain software package is not limited to using it on just one machine. If there is need to travel to an off-site location and run this same software on another Series 300 Computer, all the user needs to carry is the ID Module and that software. Second, if a machine were to fail, a user can, if more than one Series 300 Computer exists at the user's site, simply move to another computer and not be prevented from using authorized software.

Series 200 Display Compatibility

Compatibility with the Series 200 method of driving the CRT display was another requirement. The Series 200 uses separate memory planes for the alphanumeric and the graphics displays, while the Series 300 uses a single bit-mapped memory plane for both. Running software programs on the Series 300 that are designed for the Series 200 and make use of both the graphics and alphanumeric planes simultaneously could cause serious problems with the resulting display on the Series 300 screen. To resolve this compatibility issue, the same boards that are used in the Series 200 to provide the alphanumeric and graphics planes of memory are supported as an option to the Series 300, but with the addition of a software-controllable switch that allows either the bit-mapped driver or the two separate planes to be connected to the CRT display. The Series 300 software operating systems also incorporate features to make application programs written for the Series 200 compatible with the Series 300 (see article on page 22).

Design Verification

The Series 300 had to succeed in four areas: marketplace contributions, quality, schedule and factory cost. It was decided to test the Series 300's definition early in the marketplace to ensure that its contributions were adequate. Focus panels, step studies, and completeness surveys were done both inside and outside HP. The product was reviewed in detail with key customers.

To ensure early availability of software and applications on the Series 300, fifty hardware prototypes were developed and distributed inside and outside of Hewlett-Packard, worldwide. These units gave valuable feedback to the designers and allowed software development and application porting to Series 300 to begin months before introduction.

Product Design

The Series 300 represents the second generation of 325-mm-wide computers in the HP 9000 family. The primary objective of the product design was to make a Series 300 system easy to expand, manufacture, and service.

Expandability is the ability of the computer to support numerous configurations. At introduction there were two processor boards and four display boards. Each board has its own segment of a rear panel and conforms to the same dimensions. Captive screws that can be turned by hand are used to secure all of the cover plates.

The Series 300 uses the Series 200 DIO cards. The four-slot card cage can be deleted and added later. An eight-slot expander can also be added to the four slots in the computer. These expansions require the removal of the top cover held on by two captive screws. Installing the four-slot card cage requires a screwdriver for two flathead screws, but installing the expander does not require a tool. Fig. 2 shows a typical step-by-step expansion of the Series 300.

The box goes together very quickly, as illustrated in the exploded view of the product, Fig. 3. The fans and motherboard are fastened to the fan plate, which then is assembled to the chassis. Front and rear feet are attached to the chassis and the power supply slides in. The DIO card cage and the printed circuit boards are installed and the unit is tested. To close the box, the top cover and power supply door are added and the front panel is snapped on. The unit is then packaged and shipped to the customer.

The total assembly time is 30 to 45 minutes, depending on the configuration. Assembling the I/O expander takes only 30 minutes and requires almost no test time. In contrast, a configuration with a 68020 processor board and a high-resolution color monitor board requires 45 minutes (15 minutes of test time). This assumes that the printed circuit boards have already been built and tested. The printed circuit boards are aged at the board level, not at the instrument level.

Assembly Process

The assembly process uses computerized process control. A serial number plate and bar code sticker are printed for the unit. A monitor displays what option is to be built. The computer only lets shippable orders onto the production line. A shippable order must have all peripherals available for coordinated shipments. The computer determines

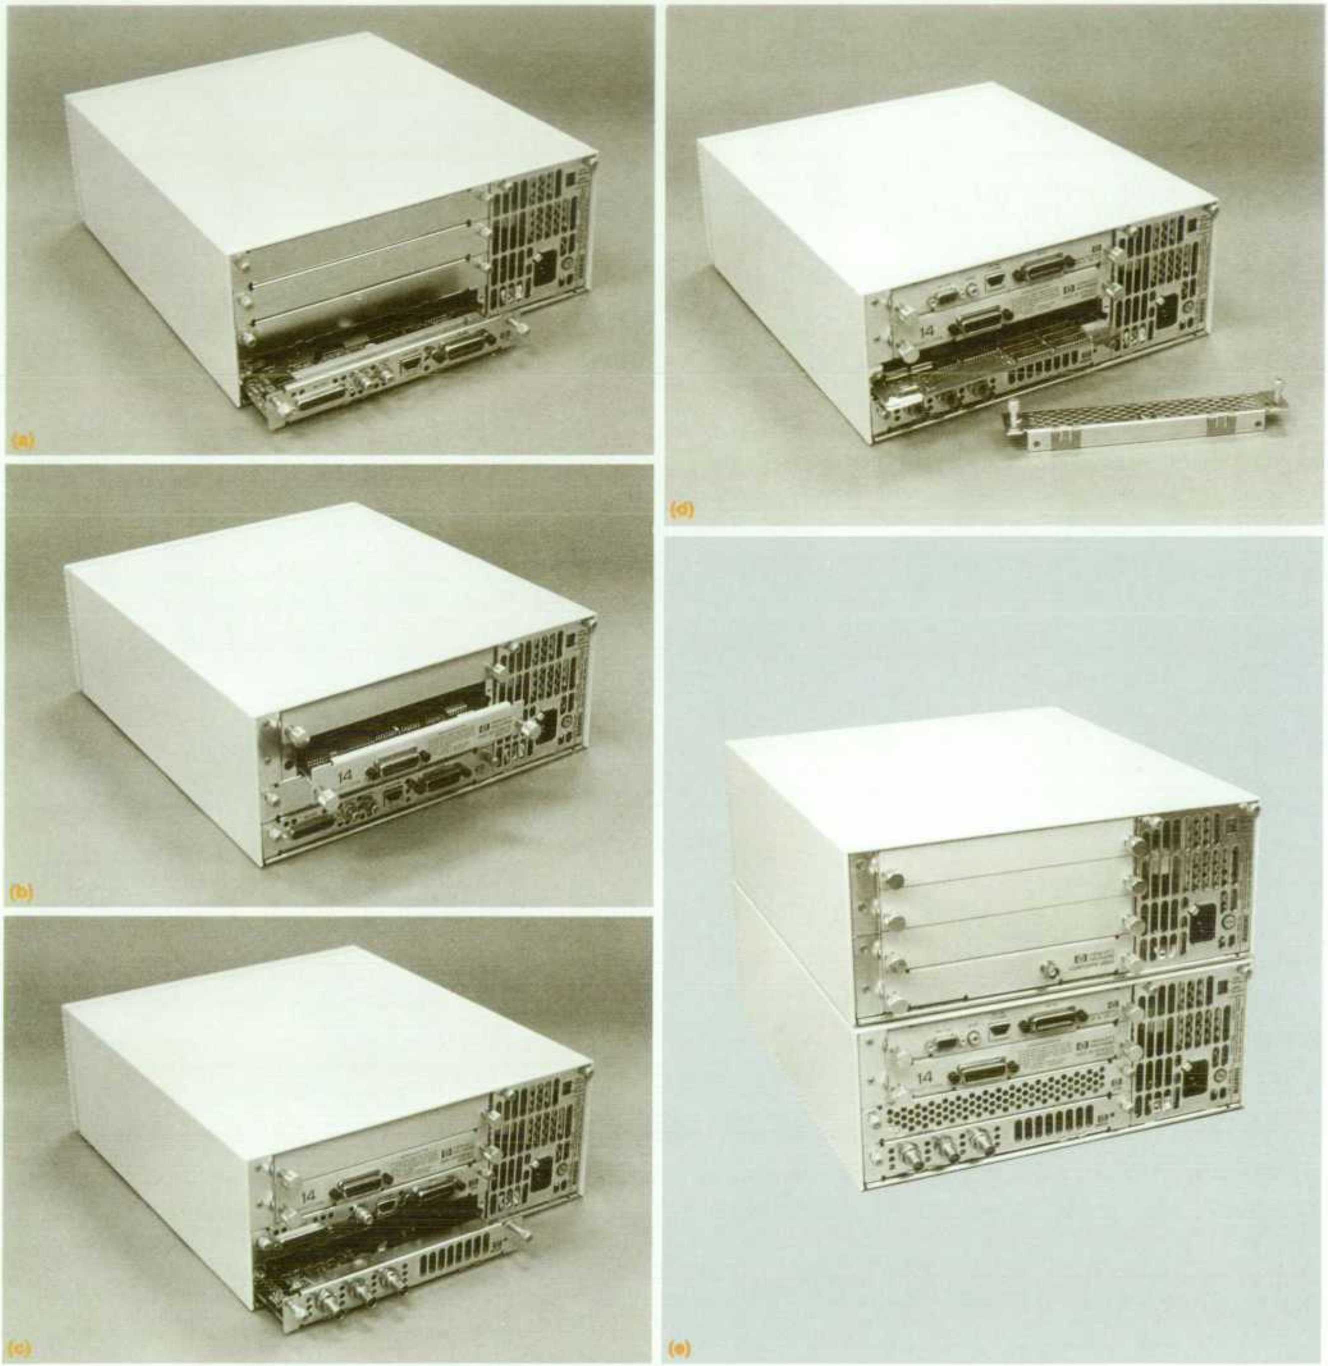

Fig. 2. *(a) The minimum Series 300 configuration is the single-board computer. (b) Expand the system by adding the DIO card cage. This allows the addition of more memory, I/O cards, or accessory cards. (c) Upgrade with any of four optional graphics interfaces—high- or medium-resolution, color or monochrome. (d) Upgrade to the 68020 processor. This requires the DIO card cage for the HP-HIL board and memory. (e) Expand again with the DIO expander for a total of twelve DIO slots.*

the sequence of configurations to be built to maximize throughput. Since testing is such a large proportion of the build time, a large backup would occur at the test station if units with long test times were built in sequence.

After the basic box is assembled, it is configured with the appropriate printed circuit boards. The operator reads the bar code into the computer and the display shows what boards to install. At the test station the test operator attaches the necessary cables, then reads in the bar code. This starts the test. The operator must verify the operation of the fans, the power-on indicator, the audio signal, and the video output.

The same operator then performs the final package closure. The bar code is read, and the computer prints the packing label (which also has a bar code on it). The top cover, front panel, door, and labels are installed. Finally, high-potential testing is completed.

At packaging, the bar code is read and the computer displays which cables need to be shipped with the unit. The bar coded cable kits are verified as they are packaged with the computer. The unit is then ready to go to the shipping department, where the bar coded shipping label helps generate a pick list of other products to include with the unit.

Package Design

The Series 300 computer and expander use the same design. Both boxes share the same tooling and parts, and were on the same development schedule. Since the DIO card cage is exactly the same height as two system boards, the expander just substitutes an additional four DIO slots for the two system boards. The same brackets and card guides are used twice in the expander. One innovative technique used to conserve space is that the connection to the backplanes is made on the back side of the board. Connectors are press-fitted in the spaces between the pins of

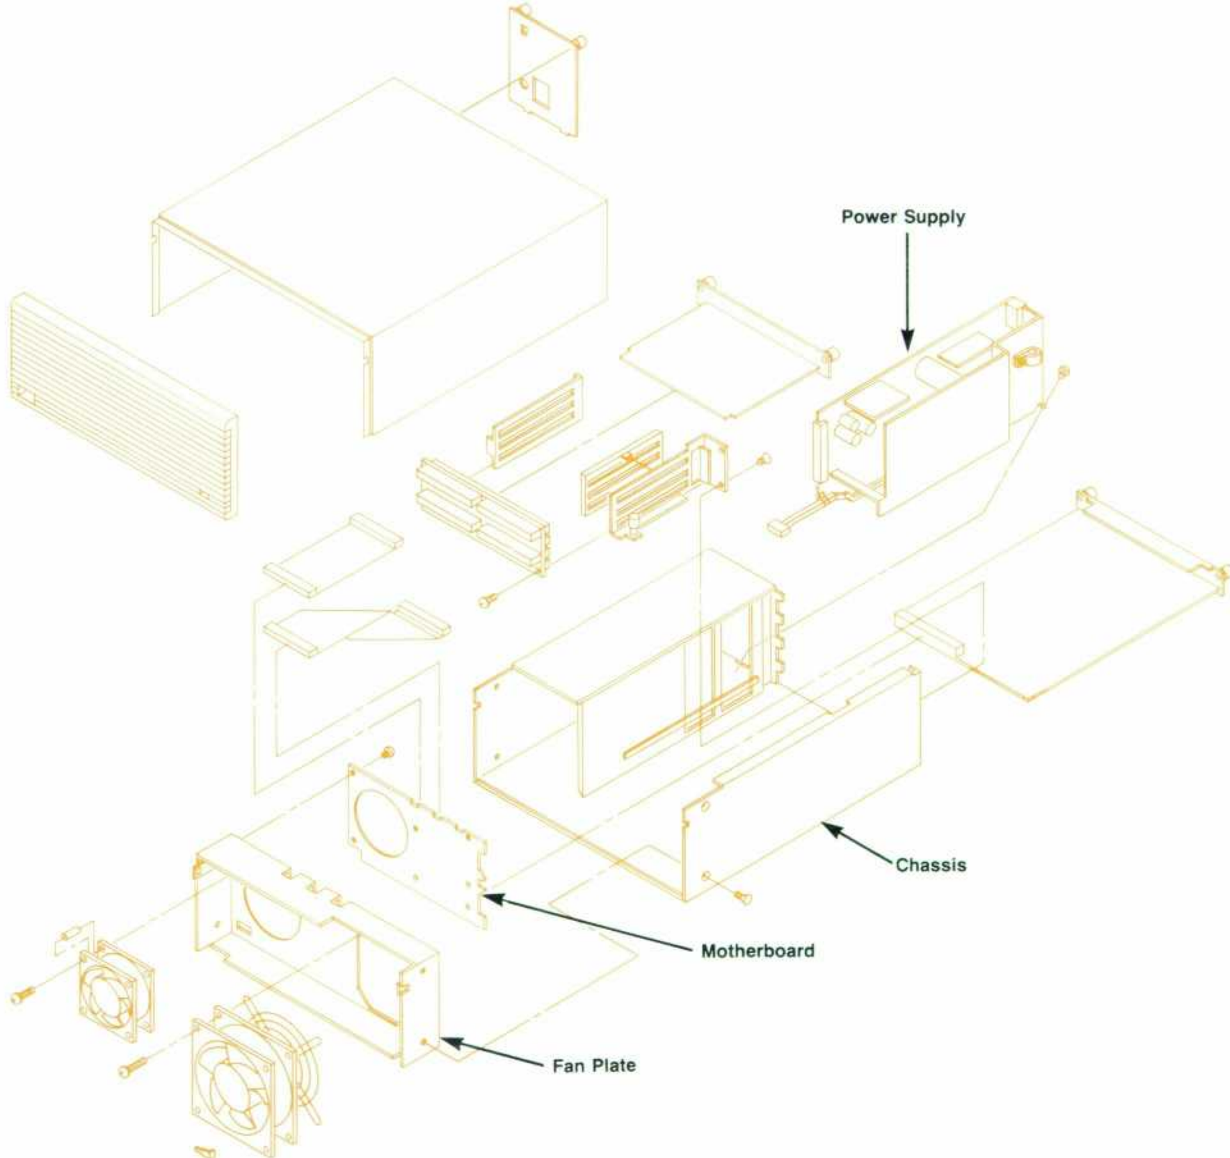

Fig. 3. *Exploded view of a Series 300 Computer.*

adjacent DIO card connectors.

The expander top cover and chassis have an added hole to bus signals from the Series 300 computer to the expander. The expander cover gets swapped with the computer during the installation and sliding clips hold the boxes together. The motherboards use the same board blank, but to reduce cost, the unused connectors are not loaded. The expander motherboard also differs because there is a power cable to the backplane. The rest of the sheet metal, all plastic parts, and the power supply are identical.

Other measures of manufacturability are the type and quantity of screws used in the unit. The assembly of the basic box requires four flathead screws to hold the sheet metal pieces together, seven panhead screws to hold the motherboard, and one screw to hold the power supply in place. Each fan requires two screws, and the DIO card cage, if present, requires six. The front panel snaps onto some spring clips and the interior side of the DIO card cage uses only tabs, slots, and hooks for attachment.

The sheet metal enclosure reduces RFI (radio-frequency interference) without the use of extensive gasketing. Each system board has two custom clips that contact the cover plate on the board below, or connect directly to the chassis. The power supply contacts some different clips on top and bottom. Those clips are the only items added specifically to suppress RFI. The extensive use of overlaps in the sheet metal parts allows this minimal addition of special parts. This avoids the addition of gasket material, which is often a concern to the manufacturing organization, being time-consuming and at times having a high scrap rate.

The Series 300 computer and expander are easy to manufacture and consequently they are easy to service. All of the system boards and DIO cards slide in and out from the rear of the unit. The power supply also slides out from the rear and there are no cables to disconnect. The fans and the LED power-on indicator can be replaced from the front of the box, even if the unit is underneath a stack of other boxes. The front panel is easily removed, giving the service person access to the cables and all of the necessary screws. However, the motherboard is very difficult to remove, requiring that the box be completely disassembled. Fortunately, the motherboard has only connectors, resistors, and diodes which have very low failure rates.

Acknowledgments

Doug Buhler, Series 300 Program coordinator, played a key role in that the program held schedule the last year of development. Danny Darr of the Fort Collins Engineering Operation provided valuable CAD/CAE inputs and Sandy Chumbley of the Colorado Networks Operation provided networking guidance.

References

1. HP 9845 Computer, complete issue, *Hewlett-Packard Journal*, Vol. 31, no. 12, December 1980.
2. Series 500 Computers, complete issue, *Hewlett-Packard Journal*, Vol. 35, no. 5, May 1984.
3. L. Dunn and M.R. Perkins, "A Standard Keyboard Family for HP Computer Products," *Hewlett-Packard Journal*, Vol. 35, no. 8, August 1984.

Modular Computer Low-End Processor Board Design

by Martin L. Speer and Nicholas P. Mati

THE HEART OF THE HP 9000 Model 310 system processing unit (SPU) is the processor board. With the exception of the power supply, no other major electrical subsystems need exist within the Model 310 SPU box. By adding a medium-resolution monochrome video monitor, an HP-HIL keyboard, and mass storage, a complete and useful workstation capable of running Pascal, BASIC, or the HP-UX operating system can be constructed.

The Model 310 processor board is a complete single-board computer with the following features:

- 10-MHz MC68010 microprocessor
- Up to 1M bytes of high-speed RAM
- Model-320-compatible memory management unit (MMU)
- HP-IB (IEEE 488/IEC 625) interface
- RS-232-C/V.24 interface
- HP-HIL keyboard interface
- Programmable sound generator (beeper)
- Battery-backed real-time clock
- Bit-mapped monochrome display electronics with 1024-dot-by-400-line resolution
- Programmable timer module (used by HP-UX operating system)
- Up to 128K bytes of boot ROM (64K bytes are currently being used).

Putting all these features on one board makes the Model 310 system more manufacturable, reduces the cost, and allows increased performance for both the RAM and the memory management unit. Yet, a system designed around the Model 310 processor board is still expandable by means of the Model 310's second system board slot, a four-slot DIO backplane, and an eight-slot DIO expander. (DIO is an asynchronous bus based on the 8-MHz MC68000 microprocessor.)

Several physical design challenges were posed by the Model 310 processor board. First, all the features listed above had to fit within 93.4 square inches and meet component height restrictions that ranged from 0.325 inch to 0.775 inch. Second, restrictions were placed on component location because of thermal design constraints imposed by the system box and a need to minimize trace lengths for critical signals. Improved timing margins, reduced cross talk, reduced trace capacitance, and reduced EMI (electromagnetic interference) were the key benefits of complying with the component location and trace length restrictions.

Increasing the performance of the Model 310 processor board had its share of design challenges when coupled with board area limitations and cost goals. The main performance contributions were made by adding a 10-MHz 68010 processor, a local RAM bus, and hardware to assist HP-IB parallel polls. The decision to use the 68010 microprocessor was made because earlier 8-MHz products were viewed as being too slow. To take advantage of the increased clock frequency of the 68010, a fast local RAM bus and RAM controller were designed. The local RAM bus supports faster memory accesses than are possible on DIO, whose timing is based on 8-MHz operation. RAM accesses to memory on the Model 310 board add no wait states while RAM accesses to memory on the DIO bus require two wait states.

Hardware to assist HP-IB parallel polls was added to improve the performance of the HP-UX operating system when the internal HP-IB interface is used as the disc interface. Without this additional hardware, the HP-UX system conducts software parallel polls periodically. Once a parallel poll is detected, the operating system must synchronize the pending disc transfer with the disc drive. This is a time-consuming task necessitated by the disc drive's requiring data within a short time after the parallel poll response. The hardware designed for the Model 310 processor board conducts a parallel poll and generates an interrupt when an HP-IB device responds. Interrupts can be serviced within the response time required by disc drives and the resynchronization software cycle becomes unnecessary. A performance increase for disc accesses of approximately 50% is observed because of this circuitry.

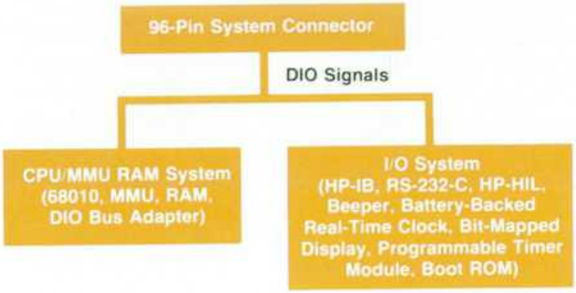

Fig. 1. *Basic system partitioning of the Model 310 processor board.*

A major design challenge was that of maintaining software compatibility with earlier HP 9000 Series 200 systems. With two exceptions, the final design of the Model 310 SPU board maintains a high degree of hardware architectural compatibility with Series 200 machines. The two exceptions are the bit-mapped display electronics and a Model-320-compatible memory management unit (MMU). All

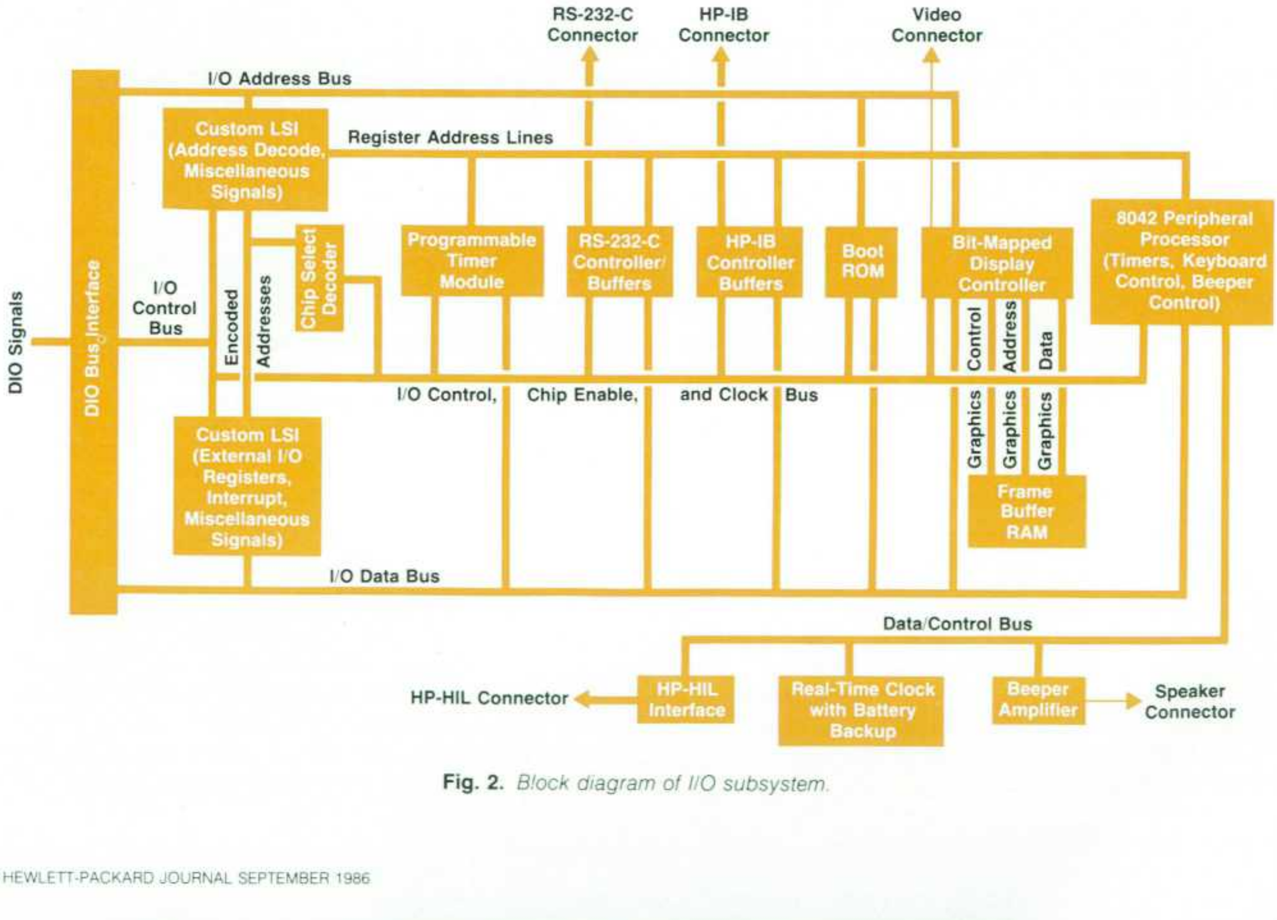

Fig. 2. *Block diagram of I/O subsystem.*

other hardware architecture changes are transparent to the operating systems and thus compatible. Cost and performance were the reasons for choosing the bit-mapped display architecture that is present on the Model 310 board. Performance, improved capability, and compatibility within the Series 300 family were reasons for the move to the Model 320 MMU.

Cost reduction and manufacturability were high priorities during the design of the Model 310 processor board. The use of custom LSI circuits, careful system partitioning, and careful component type and package selection were three important actions that contributed to cost reduction and manufacturability. By carefully partitioning the system, most of the support logic could be incorporated into four custom LSI circuits. The design of each custom IC was done with the total system needs in mind, thus further reducing the external components needed. Plastic DIP components were selected when available. These components are typically less expensive than other package types and they can be loaded onto printed circuit boards by existing production machinery.

The complexity of the design effort was greatly reduced by partitioning the Model 310 SPU board into two independent subsystems, the I/O subsystem and the CPU/MMU/RAM subsystem (Fig. 1). These two subsystems are connected only at their DIO interfaces, but share some commonly generated clocks. To ensure Series 300 family compatibility, the Model 310's I/O subsystem design was leveraged in developing the human interface card that is used in Model 320 systems.

I/O Subsystem

The I/O subsystem (Fig. 2) consists of the following sections on the processor board:

- HP-IB, RS-232-C, and HP-HIL interfaces
- Programmable sound generator (beeper)
- Battery-backed real-time clock
- Bit-mapped monochrome display electronics
- Programmable timer module
- Up to 128K bytes of boot ROM.

Two of the sections listed above have architectural changes that were added to get additional performance—the HP-IB section and the bit-mapped monochrome display section. As mentioned earlier, the HP-IB section on the Model 310 processor board has additional hardware for performing an HP-IB parallel poll. The article on page 17 provides information about the bit-mapped display controller and what hardware features were incorporated in it to improve its performance.

To keep the device count down and provide all the I/O features listed above, two custom LSI ICs were designed using Texas Instruments' standard cell technology. These two ICs generate all the chip enable signals, DIO handshake signals, DIO buffer control signals, and HP-IB DMA (direct memory access) support signals needed. In addition, all the architected DIO registers needed by the different I/O sections and the HP-IB parallel poll hardware are part of these two custom ICs. These two ICs make it possible to put all the I/O functionality mentioned above on the Model 310 board while still leaving room for the CPU, MMU, and 1M bytes of RAM.

Although the Model 310 is the low-end SPU of the Series 300 family, it has up to 85% of the performance of the earlier high-end Series 200 machine at a cost less than the low end of the Series 200. Unlike the Model 320, the Model 310 does not have a cache memory to improve performance. The addition of a cache would have been prohibitive in terms of expense and board area. Instead, the Model 310 is highly tuned to operate with its 1M bytes of on-board RAM. The 10-MHz 68010 processor runs no-wait-state memory cycles (maximum 68010 performance) out of the on-board RAM even when memory mapping is enabled. The Model 310 can also access RAM over the DIO bus, but these accesses take approximately 1.5 times longer and cause the 68010 to insert wait states in its memory cycle.

MMU/RAM Subsystem

Fig. 3 shows a block diagram of the MMU/RAM subsystem. All 68010 bus cycles to either DIO or high-speed RAM must pass through the MMU where the translation of addresses takes place. The architecture of the MMU is functionally identical to the discrete 32-bit MMU on the Model 320, which is described in the article on page 12.

The MMU/RAM controller was integrated into a Motorola MCA2800ALS ECL gate array because of space restrictions. The MCA2800ALS gate array provides 120 TTL-compatible I/O pins and high-speed operation. The gate array design is highly self-contained, requiring only TLB (translation lookaside buffer) RAMs and a few TTL support chips to form the complete MMU/RAM controller. Generation of system functions such as bus error timeout and bus master arbitration, as well as control of the DIO bus interface and execution of DIO bus cycles are performed by the gate array.

Full support for 256K-bit 120-ns dynamic RAM is provided by the gate array. It is designed to accommodate the 512K-byte or the 1M-byte loading options of the Model 310 board by sensing whether a pull-up or pull-down resistor is present at each RAS line during power-up. The presence

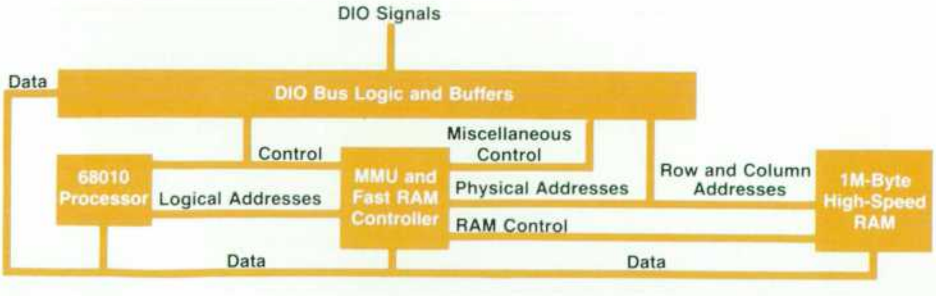

Fig. 3. *Block diagram of MMU/RAM subsystem.*

of loaded parity RAM is sensed by a transition on the parity data lines. Only after a transition occurs can parity fault checking begin.

I/O pins are a precious commodity on the gate array. Rather than allocate five pins to define the starting address of the on-board RAM, the RAM is instead automatically located in a fashion similar to the RAM on some Series 200 processor boards. By scanning down through the memory space, the boot ROM locates the starting address of the on-board RAM by finding a location where no device or RAM responds and a bus time-out occurs. On the first time-out after power-up, the MMU asserts HALT along with BUS ERROR, forcing the 68010 to rerun the bus cycle. At the same time, the RAM's address select decoders are latched at the vacant address. When the cycle is rerun, the on-board RAM is now located at the formerly vacant address and responds to the 68010.

Since there is no cache memory on the Model 310 board, a great deal of attention was paid to optimizing on-board RAM performance. At the beginning of a mapped 68010 bus cycle, the nine low-order address bits from the processor are allowed to ripple through the MMU directly into the RAM array to set up the row address. The upper bits of the logical address simultaneously pass through the TLB RAM. If the translated address selects the on-board RAM, then the appropriate RAS line is immediately asserted to the RAM array since the row addresses have met the setup time. The memory cycle to the on-board RAM is under way even before the DIO bus cycle has commenced. Once the DIO cycle has begun and addresses are latched in the DIO buffers, the column address for the on-board RAM is multiplexed onto the nine low-order address bits out of the MMU and into the RAM array. CAS is then asserted and the memory location accessed.

The dynamic RAM refresh controller uses the eight lower address bits out of the MMU to run a RAS-only refresh cycle to the on-board RAM. As a first priority, the RAM controller tries to initiate a refresh cycle during a DIO cycle. If it cannot bury a refresh during a DIO cycle within 4 μs, then it holds off further MMU activity while a RAS-only refresh cycle is run.

The position of the various control signals to DIO and on-board RAM can be precisely tuned for optimum performance since the gate array's state machines run at the system clock rate of 60 MHz. This fine resolution allows events to occur with little dead time and eliminates the need for analog delay elements on the Model 310 processor board.

High-Performance SPU for a Modular Workstation Family

by Jonathan J. Rubinstein

THE HP 9000 MODEL 320 COMPUTER is the high-performance member of the Series 300 family. It is based on a 16.67-MHz MC68020 microprocessor and an MC68881 floating-point coprocessor. The processor board is a full 32-bit implementation that uses a 16K-byte high-speed cache memory to allow the processor to operate at full speed. A 32-bit memory management unit (MMU) provides up to four gigabytes of virtual address space.

The Model 320 processor board is fully compatible with the 68000 DIO bus architecture, allowing it to replace the lower-performance Model 310 processor board based on a 10-MHz MC68010 microprocessor without any change to the system. This compatibility also allows any HP 9000 Series 200 memory board or I/O card to be used in the Model 320.

The 68020 is the 32-bit implementation of Motorola's 68000 microprocessor architecture. In addition to new 32-bit instructions and addressing modes, the 68020 contains a 256-byte instruction cache (I-cache) and coprocessor support. To increase its performance, a three-stage instruction pipe and instruction overlap are included in the 68020.[1]

The 68881 is the floating-point coprocessor for the 68020. It provides an extension to the 68020 instruction set with full IEEE 80-bit floating-point support, including transcendental functions. Since the 68881 is a coprocessor, the programmer is unaware that it is separate from the 68020 and thus sees the pair as a single processor.

The latest-generation microprocessors have a performance level that matches that of mainframes of a few years ago. To use this performance fully, architectural features similar to those seen in mainframes must be used. A memory hierarchy and memory management are examples of the types of features that can provide higher performance without adding excessive costs or constraining the size of the main memory subsystem. In agreement with this trend, the Model 320 contains a high-speed external cache and memory management unit. (For a general discussion of cache architectures, see reference 2.)

Cache Architecture

After careful analysis of performance, price, and board area, we selected a 16K-byte cache. The cache is implemented with 2K×8 RAMs, which are readily available and do not consume an excessive amount of power because of their power-down capability. To keep the complexity down, a "write through" memory update policy is used. This policy implies that a write executed by the processor is written to the cache and the memory in parallel. A write access incurs a penalty equivalent to one or two memory accesses, depending on the size. The Model 320 cache buffers both instruction and data accesses.

The 68020 is capable of accessing memory in three clock cycles. However, it is difficult to implement a cache for the 68020 that can be accessed without additional clock cycles (wait states). If a physical cache is implemented, it is almost impossible with current technology not to introduce wait states. The trade-off is to add one or more wait states for address translation or use a logical cache. If a logical cache is used, the hit rate of the cache is lower because of the cache purges required. However, simulation shows that the hit rate of a logical cache is not lowered to the point of reducing the performance below that which would be obtained if one wait state were added to use a physical cache. If address translation adds two wait states, then the logical cache is the clear choice over a physical cache.

A typical argument against a logical cache is the added software complexity. However, with the 68020, which already contains a logical instruction cache, little additional operating system support is needed for external cache support.

Given these considerations, a logical cache implementation was chosen for the Model 320 system. To achieve no-wait-state access, 35-ns and 45-ns 2K×8 RAMs from Toshiba are used. To build the valid bits for each entry, 25-ns AMD9150 1K×4 RAMs were selected. These RAMs have an additional clear capability, allowing the cache to be purged in one bus cycle.

Once we chose a logical cache architecture, an effort was made to increase the cache hit rate. We found through simulation that if the supervisor or user entries in the cache could be purged separately, the hit rate of the logical cache is only slightly lower than the hit rate of a physical cache. The assumption is that the operating system only purges the portion of the cache that requires it. To implement this enhancement, we use two valid bits for each entry in the cache, either of which can be set or cleared. In other words, every entry in the cache can contain data or instructions from either supervisor space or user space. Under software control all of the user or supervisor entries can be purged.

We selected a line size of 32 bits. The line size is the width of a cache entry and a 32-bit line size allows the processor to access the cache with a single bus cycle. The choice of a set size was more difficult. A set size of either one or two could be easily implemented given the available RAM technology; hence the decision had to be based strictly on the price/performance trade-off. After comparing the increased performance provided by the higher hit rate with the cost of the additional RAM and comparators required for a set size of two, we decided to use a direct mapped (set size of one) approach.

Memory Management Architecture

The Series 300 uses an HP-defined MMU architecture, which provides four gigabytes of virtual memory for each process in the HP-UX operating system, HP's enhanced version of AT&T Bell Laboratories' UNIX™ operating system. The page size is 4K bytes and a two-level, 32-bit table entry, paged MMU is used. The MMU used on the Model 320 is completely compatible with the Model 310 implementation, allowing identical HP-UX kernels.

The table structure supported by the Series 300 is similar to that of the earlier Series 200 Computers. This similarity helps minimize the software effort required to port the operating systems. In addition, the Series 300 MMU is also a subset of the Motorola 68851 PMMU definition, ensuring a compatible growth path with Motorola products in the future. The two-level table walk is shown in Fig. 1.

The 32-bit logical address is divided into three offsets: bits LA22 to LA31 are the offset into the segment table, bits LA12 to LA21 are the offset into the page table, and bits LA0 to LA11 are the offset into the page. The user or supervisor root pointer is chosen by FC2, which is part of the logical address.

The root pointer contains the upper 12 bits of the starting address of the segment table. The segment table offset is concatenated with the selected root pointer to find the segment table entry. The page table offset is concatenated with the segment table address entry (page table pointer), which chooses a page table entry. The page table entry

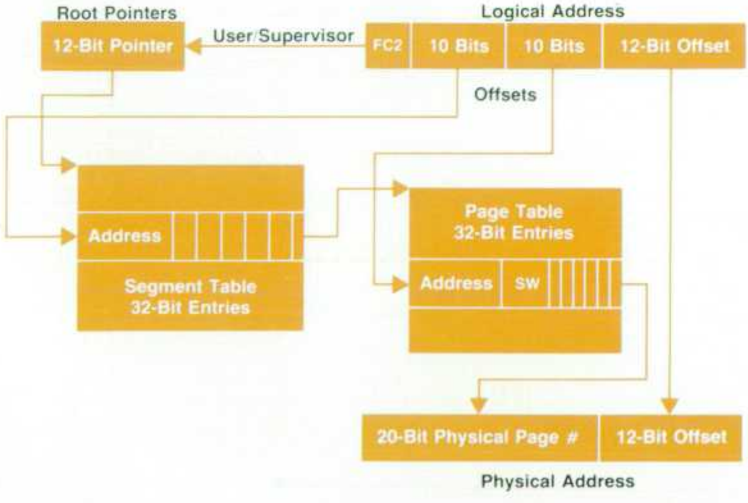

Fig. 1. *Two-level table walk.*

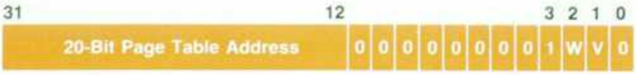

Fig. 2. *Segment table entry.*

contains a pointer to a physical page, and the page offset is used to access the correct byte address.

The segment table entry is defined in Fig. 2. The address entry is the upper 20 bits of a 32-bit page table physical address. The W bit is the write protect bit and the V bit is the valid bit.

The page table entry is defined in Fig. 3. The address portion of the entry is the upper 20 bits of the physical page address. The next four SW bits are ignored by the hardware and are allocated for software use. The D bit is the dirty bit and is set by the hardware when a page is written. The R bit is the referenced bit and is set by the hardware when a page is referenced. As in the segment table entry, the W bit is the write protect bit and the V bit is the valid bit.

A performance enhancement is made to the MMU by adding a cache inhibit bit to the page table entry. The cache inhibit bit CI blocks the cache from loading an entry from the associated page into the cache. This capability is used to keep I/O pages out of the cache and can be used to prevent pages used for DMA (direct memory access) activity from being cached. By not caching pages used with DMA (such as disc buffers), the need to purge the cache on completion of DMA is eliminated.

To speed up address translation, a translation lookaside buffer (TLB) is used to store recently generated translations. In the Model 320, the supervisor and user spaces each have a 1024-entry TLB. A TLB of this size allows eight megabytes of physical space to be mapped simultaneously and simulation shows that the miss rate of the TLB is less than 1%.

In the simplest form of logical cache, the MMU is not accessed until after the cache is accessed and a miss occurs. When this implementation is used, the cache must be purged more often because of the additional purges necessary whenever the operating system is executing MMU housekeeping activities. We avoided these additional purges by accessing the cache and TLB in parallel. The processor cannot use the data in the cache unless there is a valid entry in both the cache and TLB. Although it is impossible to guarantee that if the data is in the cache, the translation will be in the TLB, it is typically so. Thus, having both buffers accessed in parallel leads to little performance degradation caused by TLB misses and increases the hit rate of the cache by reducing the number of purges.

Cache Simulation Description

To design the Series 300 family, we developed a cache and TLB simulation that allowed us to make correct design trade-offs and characterize possible system configurations. To simulate the behavior of the 68020, actual 68000 data traces were obtained by monitoring the backplane of a Model 236 Computer. A special bus interface card was designed to monitor transactions across the bus and allow a second Model 236 Computer to store the data.

Once a 68000 address trace is collected, it is converted to a 68020 trace to simulate the external cache hit rate. We wrote a program to make this conversion and store a 68020 trace. This program simulates the instruction pipe and the I-cache of the 68020. To analyze the external cache, we wrote a cache simulation program that allows parameters such as cache size, line size, set size, replacement algorithm, and associativity to be varied. We later modified the cache program to allow similar simulations for TLB analysis.

Since the hit rate of a cache is not only dependent on the cache configuration, but also on the specific application, different traces are required for different applications. To characterize the Series 200 family, we traced the three operating systems available: HP stand-alone interpreted BASIC, HP Pascal Workstation, and HP-UX.

Thirteen traces were generated: one from BASIC, two from the Pascal Workstation, and ten from HP-UX. The trace from BASIC was a fast Fourier transform with graphics display. Since BASIC is interpreted and stands alone, the cache hit rate is less dependent on the application program than it is with other operating systems. The interpreter uses most of the system resources and thus a larger sample was not required. From the Pascal Workstation, two types of tasks were chosen: a large compile, which is representative of general processing, and the recalculation of a VisiCalc™ spreadsheet, which is computation bound.

To characterize the HP-UX operating system, thirty programs were selected. The HP-UX traces were generated using the following types of applications: floating-point and nonfloating-point intensive arithmetic programs such as B1D and Sieve, the Pascal, Fortran, and C compilers, the vi editor, several graphics intensive programs, disc intensive programs such as find and grep, and other system utilities such as nroff. With all thirteen traces, a total of 200 million accesses were analyzed.

An investigation simulating all possible cache organizations would have been overwhelming. Instead, we chose a few of the most promising cache organizations. Table I shows the cache hit rates for four organizations. The hit rate is calculated by dividing the number of hits by the total number of accesses coming from the 68020, and then multiplying by 100. For the larger caches, the hit rate is more dependent on the percent of writes since these accesses are always counted as a miss. Because of a 64% I-cache hit rate, the hit rates are lower than documented for many systems in the past.

Predicted and Actual Performance

To characterize the performance of systems with similar or identical processor architectures, we use the million-in-

Fig. 3. *Page table entry.*

Table I
External Cache Hit Rates

Cache size (bytes):	8K	16K	16K	32K
Set size:	1	1	2	1
BASIC system	76.7	84.0	86.7	85.0
Pascal Workstation	65.9	71.7	74.2	76.1
HP-UX: (average)	63.9	69.2	72.0	73.5
(standard deviation)	5.4	4.7	5.3	4.1
Worst hit rate	58.0	65.1	67.0	69.2
Best hit rate	76.7	84.0	86.7	85.0
Average of all traces	65.3	70.8	73.6	74.9
(standard deviation)	6.2	6.1	6.3	4.9

structions-per-second (MIPS) metric. To calculate a MIPS value for the various implementations, the average access time of the processor must be calculated. This access time is defined as the average interval from when the address is valid to when the data is sampled. This time is based on the minimum access time of the processor degraded by the access time to memory or cache. For the cache-based Model 320, the average access time can be calculated as follows:

$$AAT = E_{hr}C_{acc} + (1-E_{hr})\times[W_r(W_{wr}M_{acc16} + (1-W_{wr})\times(M_{acc32})] \qquad (1)$$

where E_{hr} is the external cache hit rate, C_{acc} is the cache access time (120 ns), M_{acc16} is the memory access time (540 ns), M_{acc32} is the 32-bit access time (1200 ns), W_r is the percent of misses that are writes, and W_{wr} is the percent of writes that are less than 32 bits.

From the 68020 simulations it was found that with a 64% I-cache hit rate, 17.4% of all 68020 accesses are writes, of which 37% are of byte or word size. Using the above data, Table II shows the possible average access times for the Model 320.

From these average access times and using an I-cache hit rate of 64%, we can calculate the MIPS value using the following equation:[3]

$$MIPS = 1.08(\text{clock frequency})\times [(\text{average clock pulses/instruction}) + (\text{average bus cycles/instruction})\times ((AAT/\text{clock period})-2))]^{-1} \qquad (2)$$

where the average clock pulses per instruction is 7.159 and the average bus cycles per instruction is 1.201. Table III shows the MIPS value for the system configurations specified above for a 16.67-MHz 68020 processor, using equation 2 and values from Table II for average access time.

Table IV compares measured 68020 performance with that of an 8-MHz 68000 system with one wait state and an 8-MHz 68010 system with 1.5 wait states (0.5 wait state for the MMU), all running identical benchmarks. The 68000

Table II
Average Access Time (ns)

Cache size:	None	8K	16K	16K	32K
Set size:		1	1	2	1
	540	452	392	362	348

benchmarks were run in Pascal on the Pascal Workstation. The 68010 benchmarks were run using C and Fortran on HP-UX. All benchmarks shown use 32-bit integers and 64-bit real numbers.

In choosing benchmarks to compare processor performance, computation-bound programs are preferable since I/O throughput tends to be disc, not processor, dependent. The integer benchmarks (Sieve, Acker, Puzzle, and Search) were collected by the University of California at Berkeley[4] and results have been published for several systems.[5] The UNIX benchmarks were published in *BYTE*[6] and the results are in real times rather than clock cycles. B1D is a standard Whetstone floating-point benchmark (10 million executions) and LFP is a very large, computation-bound, floating-point-intensive program.

The performance of an 8-MHz 68000 system with one wait state or an 8-MHz 68010 system with 1.5 wait states is about 0.5 MIPS. The calculated MIPS value for the Model 320 system is 1.5, which is about three times faster than the 68000/10 systems. From the benchmark comparisons, the measured average relative performance increase is 3.5 (Table V). Note that for the small integer benchmarks, the 68020 has higher performance than typically seen, because the program is loaded entirely into the internal and external caches.

As always, the choice of benchmarks can dictate how well a system fares compared to other systems. The intention of the benchmark data presented here is not to characterize the performance of the Model 320 system, but to correlate the theoretical results with the results from an actual system.

68020/68881 Qualification

To qualify a new system requires many forms of testing. This testing becomes more complicated if a new processor is being used. Not only must the system be proved reliable, but it also must be shown that the new processor executes correctly.

For the Series 300 project, we spent considerable time checking the 68020/68881 processors for correct operation. We also spent many months doing stress/life (strife) testing, and RAM/DMA tests. These tests were in addition to the standard HP Class B environmental tests.

Table III
16.67-MHz 68020 MIPS Values

Cache size:	None	8K	16K	16K	32K
Set size:		1	1	2	1
	0.80	1.30	1.43	1.50	1.54

Table IV
Benchmark Execution Time (milliseconds)

Benchmark Name	Operating System	8-MHz 68000/10	16-MHz 68020	16-MHz 68020 and 68881
Sieve	Pascal	679	158	
Acker	Pascal	5660	1730	
Puzzle	Pascal	23,410	5450	
Search	Pascal	3.4	0.8	
B1D	Pascal	376,480	105,000	14,700
Pipes	HP-UX	9000	4200	
SCall	HP-UX	11,900	4700	
FCall	HP-UX	1300	400	
Loop	HP-UX	11,500	2700	
B1D Fortran	HP-UX	413,500	113,500	14,500
LFP Fortran	HP-UX	2,079,600	732,900	345,600

To qualify the 68020, it was necessary to prove that the 68010 portion of the 68020 was correct before the new instructions and addressing modes could be checked. We used the test code and operating systems available for the Series 200, making only changes required for operation. Once the test code, BASIC system, and Pascal Workstation were operating, the HP-UX operating system was modified to run on the Series 300.

HP-UX tests not only instruction integrity, but also the virtual capability of the processor (i.e., instruction continuation). Any problems found were immediately verified with Motorola so that the problem could be fixed on the next revision of the part. During this testing, the new 68020 instructions and addressing modes were added to the test code and the operating systems were compiled with the newer 68020 compilers. After this process was repeated a few times, a 68020 was available for customer shipments that correctly ran all the Series 300 operating systems and had no known problems that could affect operation.

To help Motorola increase the speed and reliability of the 68020, we also did margin testing for each new mask revision. This testing consisted of low-voltage, high-temperature, and high-frequency tests. The frequency would be increased until the part failed. The failure was then analyzed and reported to Motorola. Using this process, we were able to attain 68020 parts capable of full 16-MHz operation much sooner. This same process was used to qualify the 68881. However, all new code had to be written since it was an entirely new part.

Table V
Benchmark Relative Performance

Benchmark Name	Operating System	8-MHz 68000/10	16-MHz 68020	16-MHz 68020 and 68881
Sieve	Pascal	1	4.3	
Acker	Pascal	1	3.3	
Puzzle	Pascal	1	4.3	
Search	Pascal	1	4.3	
B1D	Pascal	1	3.6	25.6
Pipes	HP-UX	1	2.1	
SCall	HP-UX	1	2.5	
FCall	HP-UX	1	3.3	
Loop	HP-UX	1	4.3	
B1D Fortran	HP-UX	1	3.6	28.5
LFP Fortran	HP-UX	1	2.8	6.0
Average performance:		1	3.49	
HP-UX average:		1	3.10	

Strife and RAM/DMA Testing

Strife testing consists of temperature cycling, power cycling, and vibration testing. The units are placed in a chamber that is cycled from −30° to 65°C. These are eight-hour cycles with three hours at both high and low temperature. During the temperature cycle, the power to the unit is cycled at various times to ensure the maximum temperature swing in the unit. When power is applied to the unit, a self-test program is executed that reports failures to a monitoring system outside the chamber. In addition to temperature cycling, the units are vibrated at 2g random acceleration for 10 minutes every few days.

Strife testing was initiated on the first 16 prototype units. These units were tested for over 80 cycles. When a failure occurred, the testing was terminated until the failure was analyzed and a fix was implemented on all the units. After the 80 cycles were completed, the first 16 production units were tested for 30 temperature cycles with no repeated problems. Naturally, any new failures found in the second strife test were analyzed and promptly corrected.

In addition to strife testing, the integrity of the Series 300 was verified using both long-term RAM error-rate tests and DMA tests. The RAM tests consist of various fill patterns and checking. The DMA tests use two HP-IB cards connected together and various test patterns are transferred back and forth from memory.

These tests were run on the strife units when the testing was on hold or completed. We tested the units at −25°C, room temperature, and 65°C. A total of 35,000 system hours of RAM/DMA tests were run with the failure rate in the range expected for soft RAM errors and no other significant failures were observed.

References

1. D. MacGregor, D.S. Mothersole, and B. Moyer, "The Motorola MC68020," *IEEE Micro*, Vol. 4, no. 4, August 1984.
2. A. Smith, "Cache Memories," *ACM Computing Surveys*, Vol. 14, no. 3, September 1982.
3. D. MacGregor and J. Rubinstein, "A Performance Analysis of MC68020 Based Systems," *IEEE Micro*, Vol. 5, no. 4, December 1985.
4. P. Hansen, et al, "A Performance Evaluation of the Intel iAPX432," *Computer Architecture News*, Vol. 10, no. 4, June 1982.
5. A. Gupta and H. Toong, "An Architectural Comparison of 32-bit Microprocessors," *IEEE Micro*, February 1983.
6. D.F. Hinnant, "Benchmarking UNIX Systems," *BYTE*, Vol. 9, no. 8, August 1984.

Custom VLSI Circuits for Series 300 Graphics

by James A. Brokish, David J. Hodge, and Richard E. Warner

THE DESIGN OF THE DISPLAY SUBSYSTEM for the modular HP 9000 Series 300 Computers was driven by the need for new levels of performance and flexibility. The overall design approach for this new family of workstations dictated that the display subsystem not only support both monochrome and color monitors, but also support both medium-resolution (512×390 pixels) and high-resolution (1024×768 pixels) displays. Another goal was to make the medium-resolution monochrome system as inexpensive as possible. Compatibility, both within the new family and with the earlier Series 200 products, is important. It was necessary to reduce the component count to improve reliability and make the single-board 68010 processor subsystem possible. To achieve these goals, we decided to implement a bit-mapped system with a custom display controller chip. A second custom chip provides the color map and video digital-to-analog converter (DAC) functions.

A bit-mapped system displays alpha characters on the screen using only graphics display techniques. This is less expensive and much more versatile than the traditional approach of having separate hardware for text and graphics images. Having every pixel on the display directly accessible by the CPU gives the ability to mix text and line drawings on the same screen, and to have multiple character fonts of different sizes and shapes. It also provides the raster support needed by window-oriented human interface programs.

The color display subsystem block diagram in Fig. 1 shows the architecture of the Series 300 display subsystems. The CPU address and data buses are tied to each display controller chip, which in turn moves data to and from the frame buffers. The frame buffer looks like a section of memory to the CPU, where each byte corresponds to a single pixel on the screen. This architecture allows displays from one plane up to eight planes (current implementations provide four or six planes). In a monochrome display only one bit of the pixel byte is relevant and a typical color display has four planes. In color displays, the parallel video from the display controller chips is run to the color map/video DAC chip. The function of the color map/video DAC chip is to map the data from the frame buffers into a specific color on the screen. The display ROM serves two purposes. First, it supplies display characteristics such as initialization constants and the number of display planes to the system software. Second, it supplies default character fonts appropriate for the display.

Display Controller Chip

The display controller chip is a custom integrated circuit

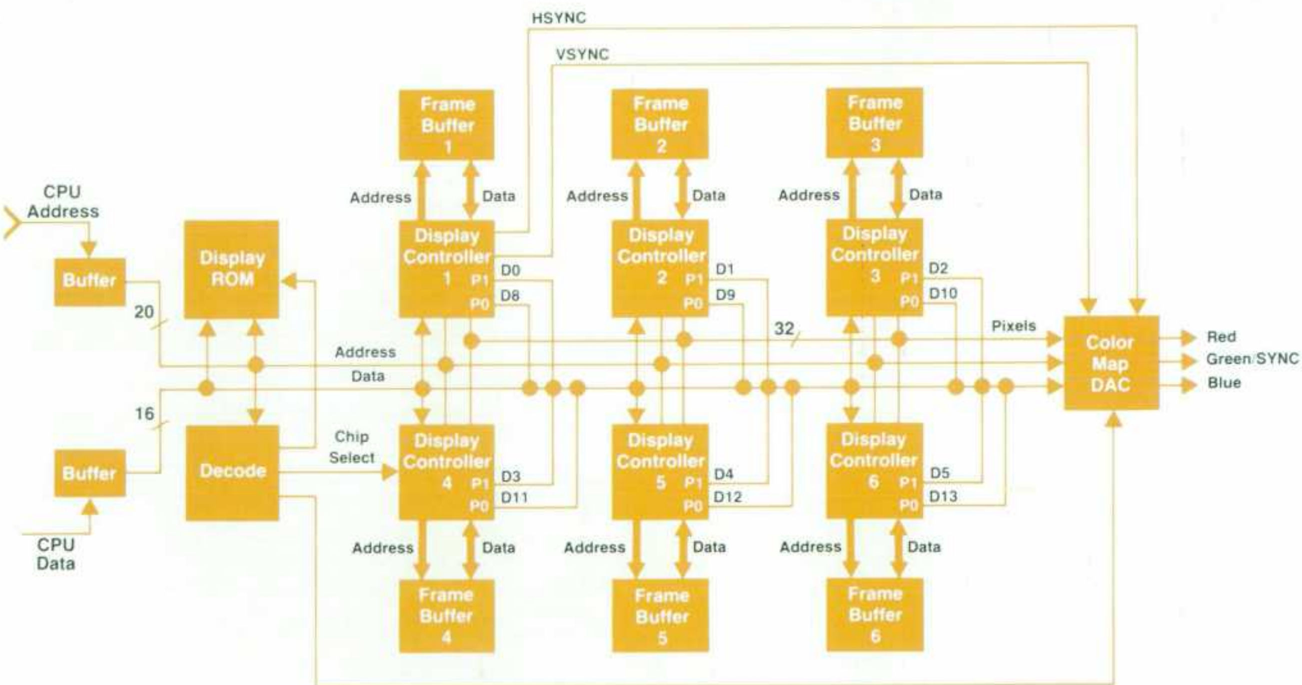

Fig. 1. *Block diagram of color display subsystem for the HP 9000 Series 300 Computers.*

built with HP's NMOS-IIIB process technology. It provides CRT control, frame buffer management, cursor, and bit-BLT (bit block transfer) functions for a bit-mapped display. A single display controller chip can control a monochrome display and multiple chips can be used for color displays.

Each of the display controller chips has two ports to the processor data bus. The first is a typical 16-bit-word data port used for most register accesses. The second is a 2-bit port called the pixel port (see the P0 and P1 pins in Fig. 1). Each display controller's pixel port is tied to two different I/O lines on the processor data bus for sensing and control. This configuration allows the Series 300 system processor to read or write two pixels to up to eight planes with only one access. For writes to the frame buffer, each controller has a pixel logic section to perform logical operations with the incoming data and the existing data in the frame buffer. Random pixels can be written to the display at 0.5 megapixels/s. Pixels written in address order will approach a rate of 2 megapixels/s.

Each display controller chip controls its own frame buffer independently, but in synchronization with the other controller chips. Accesses to the frame buffer consist of two interleaved cycles. The first cycle is a read that is used to refresh the display. The data from the read goes to an internal video shift register which sends video to the color map/video DAC chip. All other accesses to the frame buffer occur during the second cycle. Although a new generation of RAMs gets around this interleaving by integrating the shift register into the RAM, they were not used in this design, primarily because of cost.

The pixel port also serves as a port to single-bit registers inside each display controller chip. This allows the system processor to read or write the same single-bit register of each controller chip in only one access. One of the single-bit registers in the controller chip is the frame buffer write enable bit. To write pixel values only to certain planes, we set the frame buffer write enable bit only on those planes.

Since the word-wide controller chip registers are at the same address for all controllers, we needed a way to determine which controller accepts and/or drives the data on the data bus. To write to the word-wide registers only on certain planes, we set the register write enable bit on those planes. This provides the software a very powerful feature that allows different controller chips to perform different functions at the same time. For example, two different window moves can take place on two different planes at the same time.

On a word register read, only one display controller chip should be allowed to drive the 16 bits of the data bus since the same registers in different controllers may have different data. This is accomplished using the register read enable bit. The software sets only one of the the read enable registers to drive the bus. If the software attempts to enable more than one controller for 16-bit reads, a population detect circuit in each controller senses that more than one of the bits on the data bus is high. This error condition is resolved by returning all controllers to their previous state.

A potential problem with bit-mapped systems is slow character scroll speed. A hardware bit-BLT was implemented to correct this problem. A specified portion of the display is copied to another location on the screen. The bit-BLT feature is useful not only for scrolling up large sections of the screen, but also for adding new characters to the screen. This is accomplished by storing the character font in the areas of the frame buffer not displayed on the screen. The characters are then quickly moved from the off-screen memory to the displayed memory. The bit-BLT feature also acts as a powerful graphics primitive, which can be useful for area fill, line drawing, icon and cursor tracking, and animation.

Word-wide registers inside the controller chips are set up with the source window address, destination window address, window height, and window width. We then write to a window move enable register through the pixel ports to begin the window move on the selected planes. The data moved into the destination window is a function of the replacement rule that specifies a Boolean logical function between the source and destination windows. The replacement rule can be source only, NOT source, NOT destination, source XOR destination, or another Boolean operation. The completion of a window move is flagged by a DONE bit or by an interrupt to the system processor. The controller chip performs window moves at 30 megapixels/s.

A hardware cursor is supported by the display controller chip. The software sets up the position and length of the cursor in selected registers of the controller. Correct substitution of the cursor for the normal video at the rate at which the video leaves the controller is a difficult problem. To work around this problem, the frame buffer data on each side of the cursor is prefetched during vertical retrace. Operations are then performed to combine the cursor with the prefetched data. When the time comes to shift out the cursor, the manipulated data is substituted for the normal data from the first half of the interleaved cycle.

The display controller chip also provides the signals for CRT control. Horizontal and vertical sync signals and a

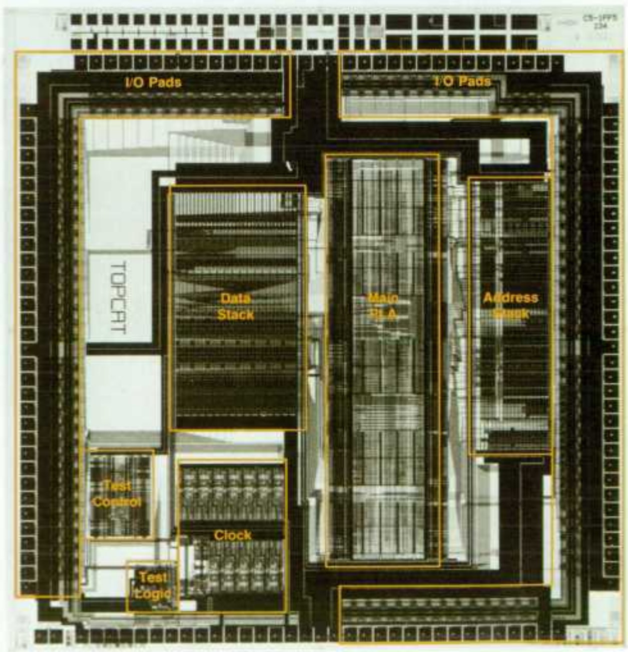

Fig. 2. *Photograph of custom display controller chip.*

composite blanking signal can be programmed to drive virtually any monitor.

Display Controller Architecture

A photograph of the display controller chip is shown in Fig. 2. The key components are a 16-bit address stack, a 32-bit data stack, the main PLA (programmable logic array), and the test PLA. The 16-bit address stack is used to calculate the addresses for frame buffer reads and writes. It also contains the majority of the registers directly accessible by the system processor.

The 32-bit data stack primarily contains data going to or coming from the frame buffer. The display refresh cycle described above dumps the data into a shift register on the data manipulation stack. A cyclic redundancy check (CRC) register on this stack performs a CRC on the video data as it is loaded into the shift register. This register can be read by the system processor to give very high confidence that the display subsystem is functioning properly. Also in the data stack are a 32-bit barrel shifter and a logic unit used for window move operations.

The main PLA is the largest section of the chip and controls all normal operation of the display controller chip. The test PLA and diagnostic interface port allow a chip tester to halt and single-step the chip operations. They also allow the tester to preset the chip to a certain state.

Color Map and Video DAC

The color map/video DAC chip is fabricated using HP's NMOS-IIIA process technology. It combines a 256-entry color map and three eight-bit DACs onto a single integrated circuit. It can drive a 1024×768-pixel, 60-Hz display with video data at approximately 64 MHz. Its RAM can supply data at rates up to eight times faster than conventional NMOS-III RAM designs. The keys to this eight-fold increase in speed are the chip architecture and RAM design.

The chip contains five major functional blocks: the processor interface and control, the video data pipeline, the RAM, the DACs, and the test logic. The processor interface allows the system processor to read and write the color map RAM, inquire about chip and display status, and perform signature analysis on a variety of signals within the chip. The chip provides a set of registers that are mapped into the processor's memory address space. The two basic types of operations are register accesses, which allow the system processor to read and write the address and data registers within the chip, and control accesses, which cause the chip to perform an internal operation such as reading or writing the RAM. The control accesses are the same as register writes from the point of view of the system processor, except that the data has no significance. The status register can be read to provide information about the internal state of the chip. It has bits that indicate the current states of several display timing signals and whether the RAM is busy performing a read or write operation.

Video data is supplied to the chip four pixels at a time. This simplifies the display subsystem by reducing the external clock rate. The chip has 32 video data pads, so each of the four pixels can have up to eight bits. Video data from the pads is routed (at approximately 16 MHz in high-resolution systems) to the input multiplexing and masking logic. First, the video data is split into two data streams of two pixels each. One stream consists of the first and third (odd) pixels, and the other the second and fourth (even) pixels. The multiplexing logic then alternately selects the pixels in each data stream and passes them to the masking logic. This logic provides a programmable mask that can be used to select which of the eight bits per pixel are routed to the color map RAM. For displays with fewer than eight planes, this feature is used to mask off the data bits for planes that are not present. It can also be used to disable the display of selected frame buffer memory planes. At this point, there are two data paths, each passing video data at approximately 32 MHz to separate RAMs.

The chip contains two identical 256×24-bit dynamic RAMs, one for the even pixels (see Fig. 4) and one for the odd pixels. This effectively doubles the RAM bandwidth. The RAM structure and cells are based on the four-transistor NMOS-III RAM described by J.W. Wheeler, et al.[1] The RAM arrays are arranged as 32 rows by 8 columns for each of the 24 bits. The basic clock rate of the color map RAM is approximately twice that of the RAM described by Wheeler for the HP 9000 Series 500 Computers (32 MHz versus 18 MHz), which again doubles the effective RAM bandwidth.

The operation of the color map RAM is very similar to the operation of the memory array of the Series 500 RAM. The basic memory cycle is four clock phases (two clock cycles). The first phase precharges the RAM array and the sense amplifier. During the second phase the row and column selects are driven and the cell data is driven to the sense amplifier. During the third phase the sense amplifier is disconnected from the array and allowed to stabilize, and during the fourth phase the output data is driven from the sense amplifier. To increase the bandwidth, two sense amplifiers are used, and the RAM cycle is pipelined so

(continued on page 21)

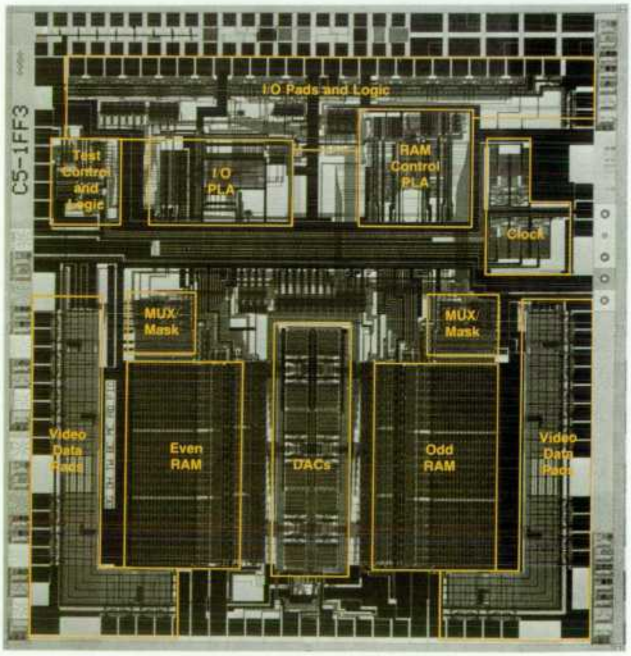

Fig. 3. *Photograph of the color map/video DAC chip.*

Display Custom IC Design Methodology

The display logic subsystem was targeted for significant price/performance improvements early in the definition of the Series 300 workstations. Previous HP workstations used two printed circuit boards with over 200 components to implement the display logic. Early estimates were that a custom display controller chip could provide the equivalent logic with less than 20 components for one quarter of the cost, while also reducing power consumption and improving reliability.

The architecture of the display controller chip dictated the design of several special logic circuits. These circuits include a logic unit, an adder, a bit field extracter, a comparator for greater-than-or-equal, and a video shift register. The high-speed requirements for these circuits led to the selection of HP's custom VLSI NMOS-III process.[1,2,3] The design methodology used was an enhanced version of that used for the design of the NMOS-III ICs for the HP 9000 Series 500 Computers. The tools used feature hierarchical design and layout, and run on HP 9000 Series 200 and Series 500 workstations with HP's Shared Resource Management (SRM) system. The methodology was developed to produce fully functional parts on the first pass, a goal that was achieved for the display controller chip.

A major feature of the design methodology is a transistor-level logic simulator which ultimately was used to verify the functionality of the entire chip. Proper logic functioning and sequencing were verified with the entire chip described at the transistor level. The simulator, called LISIM, simulates the transistors as unidirectional and bidirectional switches. LISIM uses two states (zero and one) and 63 strengths to resolve conflicts between multiple driving transistors. It solves the charge sharing problem on circuits driven by precharged buses by assuming that, when the strength is zero (no active drivers), ones have a higher strength than zeros. This results from our methodology in which zeros are always actively driven while buses may be precharged to one and allowed to float. LISIM was able to simulate the entire 77,000-transistor design on a Model 236CU Pascal workstation with two megabytes of memory at 19 seconds per input vector.

The LISIM transistor-level simulator allowed the design team to go directly from an algorithmic design to transistor-level schematics (for those cells not already in the NMOS-III library), bypassing the need to create and verify a logic-gate-level description. LISIM also made it possible to use the test vectors developed during logical design on the netlist that was extracted directly from the physical artwork.

The sequencing and logic operation of the display controller chip are controlled by a programmable logic array (PLA). The PLA was specified in a Pascal-like program. This program was then compiled using a program called QuickPLA that produced optimized logic equations. These PLA equations, along with test vectors, were used to simulate the logical operation of the chip. They were also used, along with a signal order list, as input for the PLA module generator. This program generated optimized artwork directly from logic equations in approximately two hours. Several iterations, including whole chip simulations, were required before the final PLA artwork was produced. Nevertheless, several months were saved by using the PLA generator.

The remainder of the chip is partitioned into blocks that the PLA controls. These blocks were designed using an in-house schematic capture package called SCIP. Beside schematic capture, SCIP features a schematic evaluator that outputs a FET list and a netlist, which it formats for both the HP Spice analog circuit simulator and for the LISIM digital logic simulator. Once each block's logic design functioned correctly within the chip simulation, the transistor sizes were chosen for speed, zero level, and power and then simulated using HP Spice.

Once the logic design was done, the physical design (artwork) was next. This was simplified by using stretchable cells and FET primitives. After the initial layout was done, the artwork was submitted to the hierarchical artwork system for design rule checking, netlist and component evaluation, and encapsulation. Encapsulation hides design details not relevant to higher-level blocks, so the designer sees the encapsulated block as simply an outline with bristles for its ports. This makes higher-level composition, design rule checking, and encapsulation a much faster process since the design rules need only be checked at block boundaries.

The chip floor plan was manually generated, taking into account optimal pad positions for signals as well as multiple power and ground pads. The clock and test pads were constrained by the methodology to occupy standard pad positions to ease testing. The power supply routing was modeled with current sources and resistors representing the current drawn by the cells and the size of the metal power lines routed between them. Then the routing model was simulated using HP Spice. The PLA and the routing channels between it and the register stacks on either side were automatically generated within the form factor constraints of the floor plan. While the PLA routing was not 100% complete, the router did save us many hours of manual editing.

Hierachical netlists described by a block description language were generated from both schematics and artwork. These netlists were automatically compared at each level of the hierarchy. These comparisons were sufficient to find all errors in intercell connections, logic implementation, and FET sizes within manually generated cells. Some of the cells with automatically generated artwork (such as the PLA) created simplified netlists for initial functional simulation that did not exactly match the extracted artwork. These cells were, however, verified by the final LISIM simulation using the actual artwork.

The LISIM simulator was rerun to verify the tools just before mask release on the netlist extracted from the artwork. In future designs this step could be omitted since all of the artwork has already been compared automatically with the corresponding logical schematic.

The custom Series 300 display controller chip does not use special circuitry for testing its data path cells. All cells with access to the main data buses are loaded from and dumped to a scannable register in the test section of the chip. The inputs and outputs of the PLA are made scannable so that the PLA itself can be tested and can control all of the register load and dump signals. This allows control and observation of the entire data path without additional circuitry. The pads were also made scannable so they can be completely tested from the serial test port without requiring an expensive high-speed tester. The pad scan path also made it possible to test the first-pass chips with the same vectors used during the design simulation. This was accomplished by scanning them into the pads serially and then single-stepping the chip.

References

1. J.M. Mikkelson, et al, "NMOS-III Process Technology," *Hewlett-Packard Journal*, Vol. 34, no. 8, August 1983.
2. J.P. Roland, et al, "Two-Layer Refractory Metal IC Process," *ibid.*
3. H.E. Abraham, et al, "NMOS-III Photolithography," *ibid.*

Bruce P. Bergmann
Member of the Technical Staff
Technical Workstation Operation
Kenneth P. Sandberg
Member of the Technical Staff
Fort Collins Systems Division

that during phase three, the array is again precharged to begin another RAM cycle using the second sense amplifier. Since screen refresh is performed continuously, this allows the RAM to provide a new data word each clock cycle (32 MHz). This increases RAM bandwidth by an additional factor of two, so the total color map RAM throughput is approximately eight times the throughput of the Series 500 RAM.

The outputs of the even and odd RAMs are alternately selected and routed to the DACs. The RAMs also have refresh counters and control logic. Refresh occurs automatically during retrace when no video is being generated. When the processor performs a read or write access on the color map, the address and data are latched until the next retrace, when the cycle is performed. Since an access can occur during each horizontal retrace, it is possible to update the entire color map in one screen refresh period.

The 24 bits of RAM output are distributed to three DACs in the Series 300. One DAC is used for each of the primary colors (red, green, and blue). Each DAC has eight bits of resolution, so the display has 16 million possible displayable colors. Each DAC has 255 current-steering transistor pairs for driving the video levels and 28 pairs for driving the blank level. The transistor pairs are built up in slices of eight for a single input bit, and the slices are combined into a stack with the slices arranged so that the transistor pairs for the higher-order bits are distributed within the array to minimize the effects of process variations. Reference current generators are distributed through the DAC array. The generators are calibrated by an external reference current supplied to the chip.

The chip DAC outputs are guaranteed to be monotonic, which means that for any two input values, the output current for the larger input will not be less than the current for the smaller input value. The linearity specification for each of the DACs requires the output current for a given input value to be within ±2.5% of the ideal current for that value. Tracking measures the difference between the output currents of two different DACs on the same chip when both have the same input, and is specified to be ±3.5% for any two DACs on the same chip.

Testing can be a problem for VLSI circuits, so the color map/video DAC chip incorporates some additional features to provide testability. One is a signature analysis capability with an input multiplexer that allows signatures to be taken on a number of different internal signals. The signals that can be tested include all of the input data bits and all of the RAM output bits. The input data bit signature can be used to verify correct operation of the input data path from the input pads through the mask and multiplexing logic, as well as the integrity of the data stream from the frame buffer. The RAM output signatures can be used to test each individual bit of the RAM.

The color map/video DAC chip also incorporates a serial diagnostic port similar to the one used in the custom CPU for the HP 9000 Series 500 Computers,[2] but uses only four pads on the chip (input data, output data, enable, and data strobe). This version of the diagnostic interface port was designed for the color map/video DAC chip and the same design and its descendants have been used in a number of other HP chips, including the display controller chip. The diagnostic interface port can be used to scan values into any of the internal I/O registers, and to execute commands that simulate all the normal chip functions and several purely diagnostic functions.

Acknowledgments

Art Dumont, Tony Walker, Bill Cherry, and Mark Coleman designed the color map/video DAC chip. Doug Buhler and Jim Jackson were part of the display controller chip design team. Brad Reak designed the display boards that

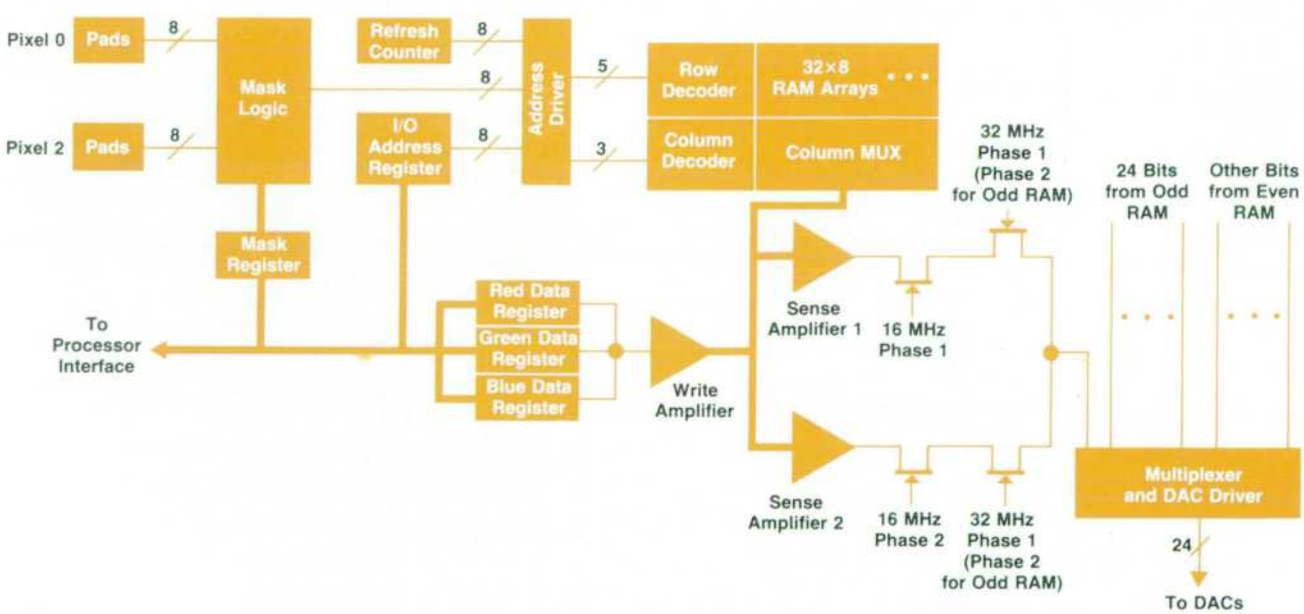

Fig. 4. *Block diagram of even-pixel RAM on color map/video DAC chip.*

use the chips. Dan Griffin managed the color map/video DAC chip project and Gary Taylor managed the display controller chip project. Many others in the Fort Collins IC Division's R&D lab and MDE organization made crucial contributions to the success of both chips.

References

1. J.W. Wheeler, et al, "128K-Bit NMOS Dynamic RAM with Redundancy," *Hewlett-Packard Journal*, Vol. 34, no. 8, August 1983.
2. K.P. Burkhart, et al, "An 18-MHz, 32-Bit VLSI Microprocessor," *ibid*.

Software Compatibility for Series 200 and Series 300 Computers

Several software obstacles exist for the Series 200 user who wants to move to HP's new family of modular workstations, the HP 9000 Series 300. This article identifies these obstacles and describes the features of BASIC 4.0 (the latest release of HP's enhanced version of the BASIC language system) designed to overcome them.

by Rosemarie Palombo

PRESERVING THEIR INVESTMENT in software is a primary concern of most computer users. They also want the flexibility to migrate to new more powerful computer systems and take their software with them. As the installed base of HP 9000 users grows, it is imperative that these needs be addressed.

Our main compatibility goal is to preserve the software investment of our customers by giving them the ability to run their existing software, without change, on state-of-the-art hardware. In addition, compatibility should be achieved without giving up any of the functionality of the new computer system and should extend throughout the entire family of workstations.

Hardware Differences

The Series 300 family of computers differs from its predecessor, the Series 200 family, primarily because the opportunity was taken during the development of the Series 300 to incorporate new technology into the existing line of HP 9000 Computers. The new hardware design resulted in some differences in machine characteristics between the two families of workstations. These differences are:

- Series 300 displays have bit-mapped planes with combined alpha and graphics. The Series 200 family has separate alpha and graphics planes.
- Some Series 200 alphanumeric highlights are missing from Series 300 displays. Gone are blinking mode (except for the alpha cursor) and half-bright intensity.
- The new Series 300 displays have different graphics resolutions. For example, on a Series 300, medium-resolution graphics displays have 512 horizontal by 400 vertical graphics pixels, whereas many of the Series 200 graphics displays have a resolution of 512 by 390 pixels.
- The new Series 300 displays have different alphanumeric and graphics hardware addresses.
- The Series 300 color map is different from that of the earlier Model 236C Computer, in that the Series 300 color map is used for alpha as well as graphics.
- Series 300 computers do not have a built-in ID PROM for software security. However, an equivalent feature is provided by an optional HP-HIL device—the HP 46084A ID Module (see article on page 4).
- Two new processor boards are available with the Series 300. One contains a 10-MHz, 32-bit MC68010 microprocessor and the other, a high-performance board, contains a 16-MHz, 32-bit MC68020 microprocessor with a new internal instruction cache, a different on-board external cache memory, and an MC68881 floating-point arithmetic coprocessor. This difference doesn't affect software portability, but some changes may be required to achieve specific performance goals.
- The serial RS-232-C/V.24 I/O interface on the Series 300 differs from the serial I/O interfaces of the Series 200 Computers in that the Series 300 interface has no hardware configuration switches.
- The HP-HIL keyboard differs from Series 200 keyboards, except for the Model 217 and Model 237 Computers.

BASIC Language

The Series 200 BASIC Language system was first introduced in 1981 concurrently with the HP 9826 Computer. It is a language well-suited for a wide range of instrument control applications, computer-aided design needs, and general computation.

Series 200 BASIC gained popularity for many reasons. Among them are its enhanced I/O capabilities, easy-to-use program development environment, structured programming features, and interactive, friendly human interface complete with knob (rotary pulse generator), softkeys, and softkey labels.

HP has continually added BASIC support for new members of the Series 200 family. HP's version of BASIC has been revised several times to support new hardware and provide additional software capabilities. In every case, support of older hardware was retained. Eventually, the gap created by the dissimilar hardware systems had widened enough (because of new state-of-the-art technology) to warrant the development of an entirely new family of HP 9000 Computers — the Series 300. Maintaining software compatibility presented a challenge for the developers of BASIC 4.0. This latest release of BASIC provides support for this new family of HP 9000 Computers and adds several new human interface capabilities, some designed especially with software compatibility as their goal.

Display

One area of change imposed by the Series 300 is the display technology. The integrated alpha and graphics displays of the Series 300 are produced by bit-mapping hardware and normally cannot be independently toggled. The Series 200 implementations have separate graphics and alpha screens where the graphics screen is bit-mapped and the alpha screen is produced by character-generation hardware. This makes it possible to turn the Series 200 alpha and graphics screens off and on independently. For many Series 200 users, the use of the separate ALPHA and GRAPHICS commands to turn the alpha or graphics displays on or off are fundamentally natural acts. A solution was needed that would allow Series 200 BASIC software that uses these features to run on the new Series 300 hardware.

One response to this need is the HP 98546A Display Compatibility Interface. This display interface provides Series 300 users with the separate alpha and graphics capability of Series 200 computers (except the Model 237, which has a bit-mapped display like the Series 300). The design of the HP 98546A is similar to that of the HP 98204B video board set for the Model 217 Computer, with the addition of an electronic video switch to allow switching between this Series 300 video board and perhaps another Series 300 video board. The switch is controlled by a register on the

Fig. 1. *The HP 98203B (top) and HP 46020A (HP-HIL, bottom) keyboard layouts.*

board and can be activated by a software write. This will shut down the current signal to the monitor and switch to the other video signal. The HP 98546A has the same alphanumeric and graphics characteristics as the HP 98204B. Therefore, any Series 200 program that is display-compatible with the HP 98204B will run on a Series 300 with the HP 98546A Interface and an appropriate monitor.

Any Series 300 Computer can be configured with the HP 98546A Interface. It can be the only video board present or it can be used in conjunction with any other Series 300 video board, medium- or high-resolution, monochrome or color. In the case where the Series 300 Computer contains both the HP 98546A and a high-resolution video board, two separate monitors are required. Use of the HP 98546A does not restrict the functionality of the Series 300 in any way, nor does it require any program changes. A limitation of this solution is that the HP 98546A only drives a monochrome display.

BASIC 4.0 supports the HP 98546A Interface and provides an easy way to select display boards whenever the HP 98546A and a Series 300 bit-mapped display board are present. The execution of a single CONTROL statement selects the specified display. The BASIC 4.0 system performs all the tasks required to switch displays and initialize alpha and graphics on the new display to the power-on defaults.

Some members of the BASIC 4.0 development team investigated another solution to the display compatibility problem. Was there any way that separate alpha and graphics could be provided on bit-mapped hardware by the system software itself? If so, this would eliminate the need for an HP 98546A Interface. The answer is that in some cases it is possible to simulate this behavior. Thus, another compatibility alternative for applications software that depends on the ability to manipulate separate alpha and graphics planes evolved and was designed into BASIC 4.0.

With multiplane displays it is possible to designate some planes for alpha and the remaining planes for graphics and subsequently perform alpha-only or graphics-only operations. The display controllers of the Series 300 color video boards allow selective plane read/write enable. There are plane control registers in which each bit controls its respective plane. By enabling and disabling planes properly, the functionality of separate alpha and graphics planes can be emulated.

Using this information, the BASIC 4.0 team decided to add a new feature that provides the ability to specify which planes to write-enable for alpha and which planes to write-enable for graphics. The graphics write-enable mask, accessed through the GESCAPE statement, indicates the frame buffer planes to be written to by graphics operations, and a CONTROL statement is used to designate the alpha planes (set the alpha write-enable mask).

With a four-plane color display we can designate planes 1, 2, and 3 for graphics and plane 4 for alpha. This provides only eight pure graphics colors, instead of 16, and a single alpha color. Restricting the number of planes that are write-enabled for alpha or graphics to less than the total number available will also restrict the number of color map pens available for use.

On bit-mapped color display hardware, this emulation gives many of the capabilities of a separate alpha and graphics system, including:

- Turning alpha and graphics off and on independently
- Dumping graphics without embedded alpha
- Independent scrolling of alpha and graphics.

This compatibility solution requires no source program changes since the appropriate masks can be set by a short configuration program.

Keyboard

When the Model 217 and Model 237 Computers were introduced, so was a new human interface, the HP-HIL (Hewlett-Packard Human Interface Loop), which included a new keyboard. This keyboard is used by the Series 300 and is different from the HP 98203A/B keyboards used by earlier Series 200 machines. The major differences between the HP 98203A/B keyboards and the HP-HIL keyboard are the number and layout of user and system function keys and the location of the rotary control knob. The number and size of the screen labels for typing aids are also different.

The HP-HIL keyboard has eight physical user function keys, labeled f1 through f8, while the HP 98203B keyboard has ten such keys, labeled k0 through k9 and the HP 98203A keyboard has only five function keys (see Fig. 1). Although the HP-HIL keyboard has fewer physical function keys, it

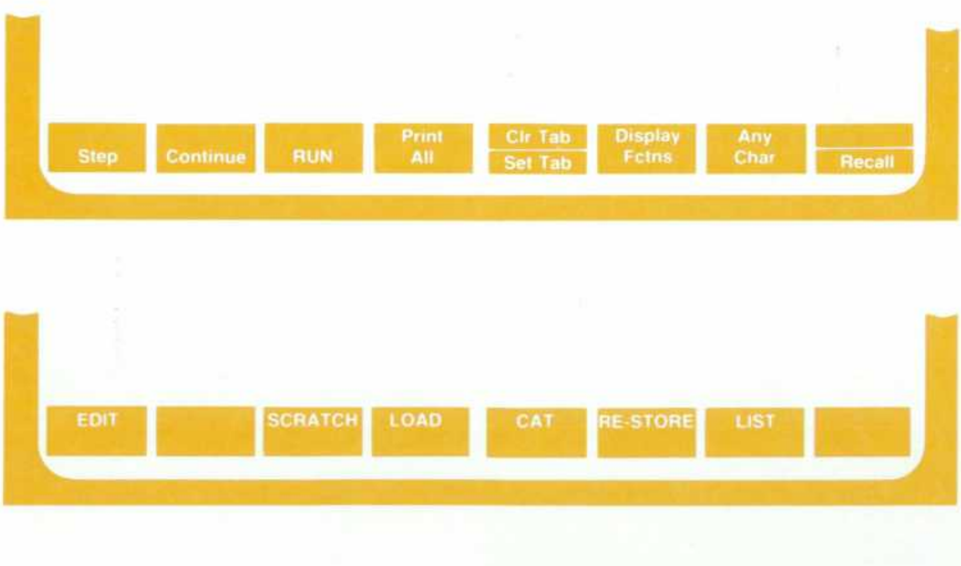

Fig. 2. *The HP-HIL system menu of keys (top) and default typing-aid labels (bottom).*

has more functionality than the HP 98203B keyboard. This is because BASIC 4.0 and BASIC 3.0 provide one menu of system keys and three menus of user definitions for the eight physical function keys, giving 24 user-definable keys compared to 20 such keys on the HP 98203B.

There are several keys on the HP 98203B keyboards, such as **STEP** and **CONTINUE**, that are not among the keycap labels of HP-HIL keyboards. These functions no longer have dedicated keys, but BASIC 4.0 makes these functions available in the system menu at all times (see Fig. 2). Other system key functions, such as ALPHA and RECALL, are available through dedicated but unlabeled keys (RECALL is also available in the system menu). A keyboard overlay was designed for BASIC users with HP-HIL keyboards to identify unlabeled keys, system menu keys, and some details of the HP 98203B keyboard compatibility mode (see Fig. 3).

A keyboard compatibility mode for BASIC 4.0 emulates some of the missing features of the HP 98203B keyboard when using an HP-HIL keyboard. This mode provides a convenient way of porting Series 200 programs to Series 300 machines without modifying the source program. In particular, it was designed to provide compatibility for Series 200 programs that were written for ten user function keys and their corresponding keylabel display. When a nonzero value is written to keyboard control register 15, keyboard compatibility mode is enabled. The HP-HIL function key row now acts as HP 98203B keys **k0** through **k9** with the HP-HIL **Menu** key acting as **k4** and the HP-HIL **System** key acting as **k5**. Similarly, the HP 98203B softkeys k10 through k19 are accessed by pressing the HP-HIL **Shift** key with the appropriate redefined function key. In this mode there is one row of keylabels for the display. Each label can contain a maximum of 14 characters and is formatted into two rows of seven characters each. If a label contains more than seven characters, it is wrapped around to the second row of characters (see Fig. 4). This softkey label format was chosen because it corresponds closely to the function key layout on the HP-HIL keyboard and it was the first choice of our human factors consultant, lab engineers, and customers who used it during development.

To emulate the HP 98203B keyboard and its softkey behavior fully, only three statements need to be executed to configure the keyboard and keylabel display. In this mode, the HP-HIL system menu of functions is available when the **System** key is pressed along with the **Extend char** key. This is documented on the new keyboard overlay.

However, keyboard compatibility mode is not as effective as we would like because:

- Displayed keylabels may be different since the length and format of the labels are not strictly identical to the HP 98203B version.
- Certain keycodes are not available on an HP-HIL keyboard, and therefore the corresponding keystrokes cannot be trapped.
- The HP 98203B system keys (for example, **RUN**) require two HP-HIL keystrokes, that is, **Extend char** with one of the function keys.
- The rotary control knob is not on the HP-HIL keyboard. However, it is available as a separate input device, the HP 46083A HP-HIL Knob.

Serial Interfaces

Another problem exists because the Series 300 features a built-in RS-232-C/V.24 serial interface that differs slightly from some of its Series 200 counterparts. Since the goal was to provide a low-cost interface in the Series 300, there are no hardware configuration switches for select code, interrupt level, baud rate, and line control parameters. If a Series 200 program depends on serial interface configurations as set by hardware switches, some software configuration is required to run the program on a Series 300.

To work around this difference, BASIC 4.0 sets default values for the baud rate and line control parameters and allows the user to change these defaults, thereby emulating the hardware switches. The select code and interrupt level are hard-wired to specific values and cannot be software controlled. During power-up, the BASIC 4.0 system sets defaults for these values. If a program expects values other than the defaults, they can be set by keyboard execution, a short configuration program, or the AUTOST routine by writing to the appropriate control registers (13 and 14). Only the cycling of power or writing specifically to these registers will change these values.

This compatibility solution requires no source code

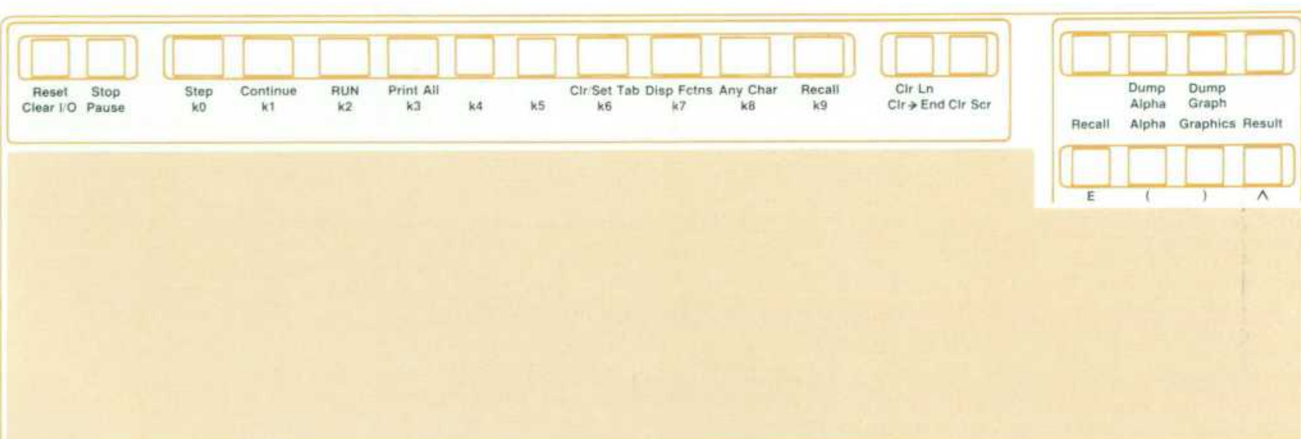

Fig. 3. *The BASIC HP-HIL keyboard overlays.*

changes. Another set of registers (3 and 4) is used to maintain the current values of these parameters, which can differ at any time from the default values. The current values can be changed by writing specifically to the appropriate control registers. However, if the interface is reset with a SCRATCH A statement, the values in these registers are restored to the default values stored.

Security

Many models of the Series 200 have a built-in ID PROM, which allows software security encoding. An equivalent feature is provided for Series 300 Computers by an optional HP-HIL device, the HP 46084A ID Module. This security module plugs into the HP-HIL interface card in every Series 300 Computer.

The ID Module returns a unique character string that contains its product number and serial number in a densely packed format. The ID Module is compatible with code that expects an ID PROM, but the character string returned by the SYSTEM$("SERIAL NUMBER") function is not in the same form as that returned by Series 200 ID PROMS. Therefore, some string manipulation is required to obtain a human readable form. This requires additional statements to be added to any Series 200 code that reads the ID Module for use on the Series 300. The required statements are included in the BASIC 4.0 documentation to make this an easy transition for users.

Limitations

As with any solution to a challenging problem, there are trade-offs to be made. Certain machine characteristics were impossible to duplicate. These incompatibilities are confined to the following areas:

- Series 200 Programs that strictly depend on the presence of an ID PROM will not work without some source code changes.
- BASIC 3.0X CSUBs (compiled subprograms) must be regenerated using the BASIC 4.0 Compiler Subprogram Utilities and the Pascal 3.0 or 3.1 language system, since the entry points to the BASIC system have changed. This has been true for every major revision of BASIC.
- Programs that depend on keycodes that cannot be generated by either the HP-HIL keyboard or the HP 98203B emulation mode (e.g., the HP 98203B EDIT and EXECUTE keys) will require changes.
- Programs that expect an I/O select code other than 9, or an interrupt level other than 5 for the built-in RS-232-C/V.24 serial interface will require changes.

The BASIC 4.0 development team reviewed these incompatibilities and found that a hardware or operating system solution was not feasible. The only viable solution for these problems was to document their likelihood for occurrence and, where appropriate, to specify the source changes required. Hence, a set of compatibility and porting documents was developed to assist customers in dealing with compatibility issues.

Pascal Workstation

BASIC is not the only HP 9000 language system that incorporated Series 200 compatibility support into its latest release. The Series 200 Pascal Workstation had identical overall goals and used similar techniques to meet the challenges of providing software compatibility to its customers. These include support of the HP 98546A Display Compatibility Interface with display selection at boot time, software emulation of the RS-232-C/V.24 hardware configuration switches, documentation, and verification.

Pascal 3.1 also provides support for the two new processor boards. Unlike BASIC, Pascal programs must work with the HP-HIL keyboard. There is no HP 98203B keyboard emulation mode available.

Verification

If the objective is compatibility, then success can be measured by determining how easy it is to run Series 200 programs on a Series 300 Computer. Consequently, we looked at the results of running Series 200 application programs. A suite of BASIC and Pascal application packages was as-

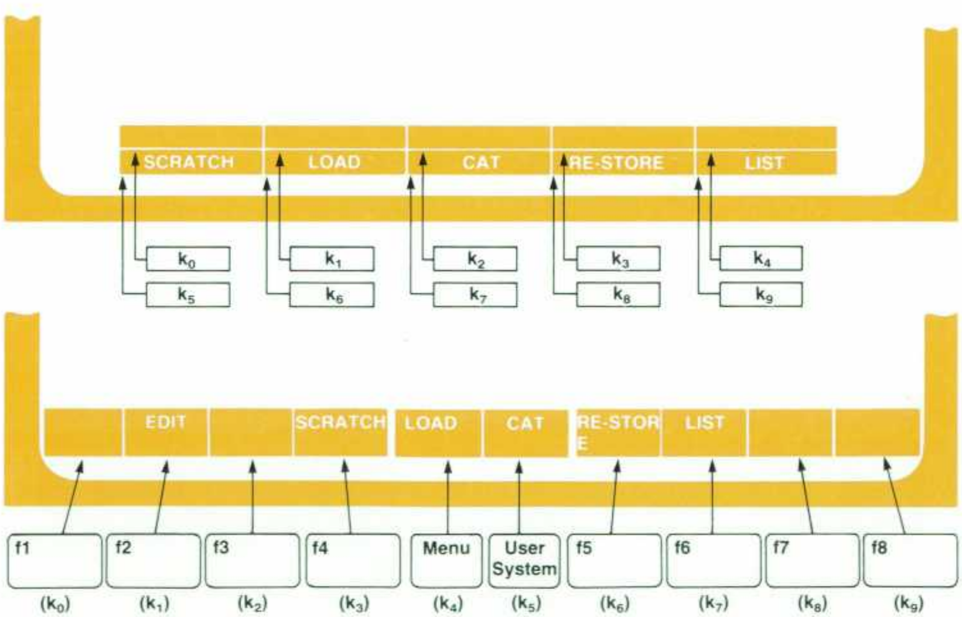

Fig. 4. *The HP 98203B (top) and HP-HIL HP 98203B keyboard emulation mode (bottom) displayed keylabel formats.*

sembled and used to test the various compatibility solutions. This group of programs included Hewlett-Packard software packages from several different divisions. These programs were augmented by the applications of third-party software developers who were interested in observing the performance of their software running in compatibility mode. The test suite was run on both the Model 310 and the Model 320.

In addition to testing functionality, the test suite also identified porting problems which were subsequently resolved by the team or documented for user correction. The HP 98546A Display Compatibility Interface provided the required functionality for all of the application programs with which it was tested. The HP 98203B keyboard emulation mode provided compatibility for the test suite with limitations as previously noted in the keyboard section of this article. Testing of the separate alpha and graphics emulation mode identified source code problems that were limited to incorrectly specified device selectors in PLOTTER IS... statements. Verification of the other compatibility features demonstrated that we had achieved the desired functionality.

Acknowledgments

Deserving of recognition is the BASIC 4.0 project team headed by Tom Christian. The team members were Ales Fiala, Nancy Madonna, Steve Taylor, Jim Tear, and Jim Whalen. The team never lost sight of the compatibility goals that were so important to the Series 300 project. Dave Dahms should be recognized for his work on the HP 98546A Display Compatibility Interface and Mark Archuleta for his compatibility and porting documentation work.

Bibliography

K.Y. Kwinn, R.M. Hallissy, and R.E. Ison, "The 9826A/9836A Language Systems," *Hewlett-Packard Journal*, Vol. 33, no. 5, May 1982.

Authors

September 1986

4 Engineering Workstations

Gilbert I. Sandberg

Gilbert Sandberg is an R&D section manager and has been the program manager for the HP 9000 Series 300 Computer. With HP since 1970, he has also contributed to the design of the HP 9817A Modular Computer, the HP 9816A Technical Computer, the HP 9920A Modular Computer, and the HP 9826A/S Technical Computer. He was born in Gettysburg, South Dakota and attended Brigham Young University. His BSEE degree was awarded in 1969 and his MEEE degree was awarded in 1970. A resident of Loveland, Colorado, Gilbert is married and has six children. He's a church leader and is active in local party politics. He enjoys basketball, skiing, and country and western music.

Ronald P. Dean

Ron Dean was born in Dearborn, Michigan and attended the University of Florida, earning a BSME degree in 1977. After coming to HP the same year he worked on a number of products in the HP 9000 Computer product line, including the HP 9000 Model 320 Computer. He's a coauthor of an *HP Journal* article on the HP 9000 Model 520 Computer. Ron and his wife and two children live in Fort Collins, Colorado. He enjoys chess, bridge, tennis, softball, bicycling, and skiing.

Daryl E. Knoblock

An R&D section manager at HP's Fort Collins Systems Division, Daryl Knoblock was responsible for CPU development for the HP 9000 Series 300 Computer. He also managed part of the development effort for the HP 3000 Series 30 and Series 44 Computers and for an HP-IB interface chip. Born in Detroit, Michigan, he attended the University of Michigan and received his BS degree in applied mathematics in 1969 and his MSE degree in computer engineering in 1971. He was a systems programmer at the University of Michigan before joining HP in 1972. A patent has resulted from his work on the HP-IB and he has written two articles on I/O channels and CMOS-on-sapphire technology. Daryl and his wife and three children are residents of Fort Collins, Colorado. He's active in his church and in community activities and enjoys skiing and camping.

John C. Keith

John Keith has been with HP since 1969 and is an R&D project manager at the Fort Collins Systems Division. He has worked on the HP 9815A and the HP 9845A Desktop Computers and has contributed to the design of HP 9000 Series 200 and Series 300 Computers. He's currently working on high-resolution display monitors. He is named coinventor on five patents related to desktop computers and is coauthor of a 1978 *HP Journal* article on the HP 9845A. Born in Des Moines, Iowa, John has a 1969 BSEE degree awarded by Kansas State University. He and his wife and four children live in Loveland, Colorado. He's active in his church and is an amateur radio operator (KD0GD).

Michael K. Bowen

With HP since 1979, Mike Bowen has a BSME degree awarded by the University of California at Davis in 1978. He is an R&D design engineer at HP's Technical Workstation Operation and has worked on the HP 9000 Series 500 and Series 300 Computers. His first HP assignment was in materials engineering. He's a coauthor of a 1984 *HP Journal* article on the Model 520 Computer. Born in England, Mike lives in Fort Collins, Colorado. He and his wife, who is also an HP engineer, have two daughters. He sails and races a Hobie cat, skis, and bicycles around Fort Collins.

9 — Processor Board Design

Nicholas P. Mati

Nick Mati is the designer of the MMU/RAM controller subsystem and the 2800 ALS gate array option for the HP 9000 Model 310 Computer. With HP since 1978, he also designed a RAM board for the HP 9845C Desktop Computer and the CPU board for the HP 9817A Modular Computer. He contributed to the design of the light pen for the HP 9845C and a portion of that work is the subject of a patent. Nick attended the University of Connecticut (BSEE 1976) and the University of Illinois (MSEE 1978). Born in Bridgeport, Connecticut, he now lives in Loveland, Colorado. He's an avid bicyclist and likes sailing.

Martin L. Speer

Born in Del Rio, Texas, Martin Speer studied electrical engineering at Texas A&M University. He completed work for his BS degree in 1977 and for his ME degree in 1979. After coming to HP the same year he contributed to the design of the HP 9816A Technical Computer and was responsible for the design of the processor for the HP 9000 Model 310 Computer. He's an IEEE member. Martin lives in Fort Collins, Colorado and likes photography, fishing, and camping.

12 — High-Performance SPU

Jonathan J. Rubinstein

With HP since 1979, Jon Rubinstein is an R&D engineer at HP's Fort Collins Systems Division. He was architect of the CPU for the HP 9000 Model 320 and Model 236 Computers. He's a member of the IEEE and the ACM and has written two papers related to the Motorola 68020 microprocessor and cache design. Jon is an alumnus of Cornell University. He completed work for his BSEE degree in 1978 and for his MEEE in 1979. He also has a 1985 MSCS degree awarded by Colorado State University.

17 — VLSI Circuits

David J. Hodge

Dave Hodge is a Colorado native who was born in Boulder and now lives in Loveland. He received his BSEE degree from the University of Colorado in 1975. With HP since 1979, he's a specialist in computer graphics and VLSI design. He has worked on graphics firmware and hardware for the HP 9845B Desktop Computer, on the video data pads and test circuits for the color map and DAC chip, and on the BIT-BLT algorithm and PLA for a display controller chip. Before coming to HP he designed milking machine controllers. Outside the office he's interested in model railroading, softball, and hiking. He recently enjoyed an extended trip to Australia and New Zealand.

Richard E. Warner

With HP since 1980, Richard Warner worked on the HP 9845B Desktop Computer before joining the original design team for the HP 9000 Series 300 Computer. He contributed to the product design, the IC for the display controller, and the graphics display boards. He's currently working on a future graphics product. Richard holds a BSME degree from the University of Colorado (1978) and an MSME degree from Stanford University (1980). Before coming to HP he designed precision machine tools and is coauthor of a research report on pneumatic isolation systems for such tools. He and his wife live in Fort Collins, Colorado and have one child. His outside interests include photography, bicycling, and history.

James A. Brokish

Jim Brokish has contributed to the design of several technical computers since joining HP in 1979. He designed the power supply for the HP 9915A Modular Computer and the alpha display subsystem for the HP 9836A Desktop Computer. He was also a production engineer for the HP 9836A, worked on a video display controller IC, and worked on the system bus design for the HP 9000 Series 300 Computer. He was born in Dickeyville, Wisconsin and attended the University of Wisconsin, receiving a BSEE degree in 1979. He's now working on an MSCS degree from Colorado State University through the HP Resident Fellowship Program. Jim lives in Fort Collins, Colorado, is married, and has one son. He and his son like spending time together in local parks. He's also a skier and white-water rafter.

22 — Software Compatibility

Rosemarie Palombo

With HP since 1985, Rose Palombo is a software engineer at the Fort Collins Systems Division. She has worked on BASIC for the HP 9000 Series 200 and 300 Computers and was responsible for the compatibility verification and documentation for the Series 300 Computer. She was born in Aliquippa, Pennsylvania and graduated in 1972 from Indiana University of Pennsylvania with a BS degree in mathematics. She also earned an MSCS degree from the University of Kentucky in 1985. She was a mathematics and data processing instructor before joining HP. A resident of Fort Collins, Colorado, she says she is a tennis fanatic, plays volleyball, and is learning to ski. She also enjoys travel, good films, literature, music, and computer games.

30 — Electronic Mail System

Robert A. Adams

A native of California, Bob Adams was born in Mountain View and attended California State University at Chico, receiving a BS degree in computer science in 1979. He worked for the National Aeronautics and Space Administration as a programmer/analyst during the same period. He joined HP in 1979 and has been a programmer and a project leader, and now manages messaging utilities for the Office Utilities Group. Bob lives in Mountain View with his wife, who is also a manager at HP. He enjoys chess and bridge.

Kristy Ward Swenson

With HP since 1982, Kristy Swenson has been the project leader for HP Desk implementation at HP's corporate headquarters, has coordinated an executive training program, and has been responsible for office automation user support, all within the Office Utilities

Group. She is currently a regional marketing engineer. Born in Portland, Oregon, she attended Scripps College, receiving a BA degree in Latin American studies in 1980. She also earned an MBA degree from the University of California at Los Angeles in 1982. Kristy lives in Half Moon Bay, California and is a volunteer both for Scripps College and for the American Field Service student exchange program. She's married and enjoys wine tasting, reading, cooking, and making beer.

Amy Tada Mueller

Amy Mueller was born in Honolulu, Hawaii and educated at Stanford University. She has a BS degree in mathematical sciences (1976) and an MS degree in computer science (1977). After joining HP in 1977, her first assignment was working on a manufacturing information system. She later moved to the Office Utilities Group and has been the manager responsible for the companywide implementation of HP Desk. Amy is married and lives in Sunnyvale, California. Outside of work she enjoys backpacking, cross-country skiing, softball, tennis, and golf.

Luis Hurtado-Sanchez

Luis Hurtado-Sanchez started at HP's corporate headquarters in 1973 as a programmer/analyst. He has contributed to the development, support, and management of office information systems, including HP Desk, and now manages the Office Utilities Group. He's the author of six articles and other publications and is a member of the Data Processing Management Association. He has a 1970 BS degree in chemistry from the University of California at Los Angeles and a 1973 MS degree in finance from the Massachusetts Institute of Technology. Luis was born in Guantanamo, Cuba and lives in Sunnyvale, California. His leisure interests include bicycling, bridge, films, and opera.

Rebecca A. Dahlberg

Born in Los Angeles, California, Rebecca Dahlberg studied Chinese history at the University of California at Berkeley (BA 1979) and management and marketing at Claremont College (MBA 1981). She joined HP's Neely Sales Region the same year, where she contributed to the implementation of an order processing system. After a transfer to corporate headquarters, she coordinated the companywide implementation of HP Desk and now supervises the HP Desk network support group. Rebecca lives in San Francisco, California and enjoys fly-fishing, bicycling, and travel.

Implementing a Worldwide Electronic Mail System

This paper reports Hewlett-Packard's experiences in the internal implementation of HP's own electronic mail system product, HP DeskManager. Prospective implementers of electronic mail systems can use this information to increase their likelihood of success.

by Luis Hurtado-Sanchez, Amy Tada Mueller, Robert A. Adams, Kristy Ward Swenson, and Rebecca A. Dahlberg

A BRIEF REVIEW of the current literature on electronic mail (EM) implementation shows a scarcity of material on this topic. In addition, many of the existing papers focus on pilot projects, small-scale implementation of EM systems, or both. This paper will concentrate on HP's large-scale (worldwide) implementation of its own EM system product, HP DeskManager (HP Desk), in which the initial or pilot project is but a small component.

The EM capabilities of HP Desk are its most important features, and this paper concentrates on their deployment throughout the Hewlett-Packard Company. However, HP Desk does much more than provide EM. It also provides the user with a set of tightly integrated fundamental office facilities. These include word processing, personal filing, and time management.

This paper aims to do three things. First, it sketches a generalized strategic framework for EM implementation suitable for use by most any type of organization, manufacturing or service, private or governmental, commercial or nonprofit. Second, within the limned framework, it provides direct and specific tactical advice to address the technical, operational, training, and support challenges that crop up in implementing an EM system. Third, it points out potential pitfalls and how to avoid them to ensure a successful project.

The paper is divided into five major sections. This first section addresses the history of messaging systems at HP and cost justification issues. It provides a background for understanding the development of the rest of the HP Desk implementation strategy within HP. It also illuminates cost-justifying projects, such as EM implementation, that do not have easily quantifiable favorable financial impacts on an organization.

The second section addresses developing and implementing a messaging strategy to cover the entire organization. It provides guidance on issues that affect the implementation of HP Desk across separate geographical entities in an organization, whether divisions or regions. These issues include: 1) data communications and networking choices, 2) project and milestone planning, and 3) management support. The next section addresses the interfacing of HP Desk with other EM systems and with diverse data communications networks. It is slightly more technical than the preceding two.

The fourth section discusses the types of training and support needed across the organization. These include both user training and support (how to use the product and take advantage of its many features) and technical training and support for personnel in charge of supporting the entities' local HP Desk networks (with regard to registered user directory updates, security guidelines, corrupt data base recovery, etc.). The fifth section addresses the day-to-day running of HP Desk at a particular entity. It offers guidance on such topics as intraentity networking, the choice between a distributed and a centralized data base, and standards and conventions to facilitate the smooth operation of an entity's local HP Desk network.

Earlier Messaging Capabilities within HP

Before undertaking the implementation of HP Desk on a massive scale, HP had a lot of experience in implementing and using other, much less sophisticated messaging systems. Knowledge of this history facilitates the understanding of the background processes that helped HP succeed with HP Desk.

Through the late 1960s, HP used almost entirely commercial TWX networks to transmit its business data, including administrative messages. Volume and cost considerations led to HP's designing and implementing its own private and internal (not available to customers) data communications network in the early 1970s, using the HP 21xx line of computers. This network, known as COMSYS/ROUTS, also included basic messaging and distribution list facilities, COMS89/MEMO, which were used in two major ways. First, some individual users had direct access to these facilities through a direct terminal connection to the COMSYS/ROUTS computer. Second, almost all entities established keyboarding departments which entered messages from written user copy into the computer for those users who did not have a direct terminal connection to it. Messages entered into the COMSYS/ROUTS network were printed at the receiving entity and then distributed through the mailroom, as was all outside mail.

During this time, HP personnel were getting accustomed to sending and receiving messages anywhere in HP with two to three days delivery time. Approximately 15% of the network volume was devoted to such administrative messages (a total volume of 10.4 billion characters was transmitted through the network in 1975, 54.4 billion in 1980). The concept of a simplified, but not simple, messaging system was born, and was well on its way to being accepted as part of the normal way of doing business.

During the late 1970s, the HP 3000 was becoming the main business data processing computer within HP. A number of small projects sprouted in Corporate Information Systems (CIS) to investigate the feasibility of using the company's large base of business computers to increase productivity through office automation (OA) utilities. As a result of the favorable results of these few first pioneer projects, the Office Utilities Group (OUG) was officially started in late 1977. (Located in HP's Corporate Offices in Palo Alto, California, both CIS and OUG served only internal needs.) One of the first projects of OUG was to put on the HP 3000 similar messaging and distribution list capabilities to what already existed on COMSYS/ROUTS, and this was achieved in 1979 through a utility called COMGRAMS/3000. In fact, COMGRAMS/3000 made use of the COMSYS/ROUTS network for distributing its messages throughout HP. Messages entered through it would still be printed at the receiving entity and delivered through the mailroom, as were messages entered through COMS89/MEMO.

COMGRAMS/3000 brought message entering and distribution list maintenance to thousands of users throughout HP. It made them comfortable with such activities. It allowed them to learn and experience first-hand the many benefits of even a simplified messaging system with extremely limited editing capabilities and lacking integration with other OA facilities, such as text processing, graphics, and spreadsheets.

The success of COMGRAMS/3000 led OUG to work on NORMAN, a simplified EM system. Unlike the basic messaging systems (terminal to paper) already discussed (COMS89/MEMO and COMGRAMS/3000), NORMAN allowed the user both to send and to receive messages (terminal to terminal). It also offered other functions, such as saving and purging messages, all within an entity's local HP 3000 network in which all the computers were interconnected through hardwired links to the central or master computer where NORMAN's data base resided. Implementation of NORMAN began in 1981 in a small number of entities throughout HP. Although it proved to be short-lived, NORMAN brought the first taste of EM, though limited, to many users throughout HP.

NORMAN proved to be short-lived because in 1982 HP introduced HP Desk (then known as HP Mail[1]), a full-capability, store-and-forward EM system. HP Desk is an HP 3000-based product, available to customers, produced by the Office Productivity Division in Wokingham, England. It is a superset of NORMAN, both in terms of user functions (handling both messages and files as well as several levels of acknowledgments) and in its ability to handle multiple HP 3000s and destinations. Besides possessing capabilities superior to those of the other in-house systems, HP Desk is an HP product, and implementing and using it internally allowed for the transfer of the knowledge gained in the process to the product division (the "next-bench syndrome"), subsequently leading to a better product. The decision was thus made in mid-1982 to converge all current development and implementation efforts by OUG in the messaging and EM areas towards the implementation of HP Desk internally.

Fig. 1 illustrates the evolution of the use of the various messaging systems within HP.

Justifying Electronic Mail

Even more than for most other OA facilities (text processing, graphics, spreadsheets, etc.), it is hard to quantify the benefits of EM use enough to prepare such traditional financial analyses as rate of return on investment, payback period, etc. This makes it hard to cost-justify such facilities strictly on the basis of financial benefits. As an alternative, admittedly less quantifiable benefits are employed to justify EM implementation. These include less travel, fewer meetings, reduced telephone tag, and more timely and effective communications—some of these benefits less tangible than others.

The task of cost justification is made easier the greater the organization's experience with any kind of messaging system. Such was the case at HP. Previous use of COMS89/MEMO, COMGRAMS/3000, and NORMAN had resulted in a relatively large percentage of HP personnel, perhaps as large as 15%, who were accustomed to entering their own brief messages through workstations (personal office computers or terminals) for delivery on paper elsewhere within HP. Such broad previous experience meant that most of these users already had the necessary equipment to use HP Desk, equipment they had purchased for other reasons and for which messaging was a value-added capability. The hurdle of purchase of new equipment necessary for large-scale HP Desk implementation was minimal to nonexistent for HP. It also meant that resistance among the user community would be relatively low. Users were already familiar with the capabilities of simplified messaging and EM systems, accustomed to using them, and, from their experience, sold on their benefits. However, the dual challenge existed first of converting the large base of users to a new, more sophisticated system, more complex to administer, and then of building upon that base.

For most organizations, there is an evolutionary process to go through with regard to EM cost justification, similar

Late 1960s	Early 1970s	1979	1981	1982	1985
TWXs	COMS89/MEMO (COMSYS/ROUTS Facilities)	COMGRAMS/3000	NORMAN	HP Desk Introduced; Pilot Projects Started	44,000 HP Desk Users Worldwide

Fig. 1. *Evolution of HP's internal messaging capabilities.*

to what HP itself went through with HP Desk. At the outset, as one or more pilot projects get started, it is imperative that they be funded on an organization-wide basis, without making a particular group or individual pay for it through one department's budget. The project must be nurtured through its first stages as users are brought on board: trained, supported in their initial struggles, and allowed to experience personally the many benefits of EM. For organizations without any previous experience with messaging and EM systems, pilot projects are extremely important. They can, and must, provide the vital rockbed of knowledgeable, satisfied users upon which continued, successful, large-scale implementation depends.

Later, individual users can be charged for the workstation and computer resources they use, and entities might be allocated a flat or prorated (by volume) share of the concomitant data communications and support costs. Users who believe they benefit from EM would have to conclude, on the basis of their own experience, that they gain enough from using EM to justify replacing other items in their financial targets with the expenses related to EM use. If the implementation has been successful, there will be a moderate, or even minimal, drop in the number of HP Desk users and their message volume.

Current Status of HP's HP Desk Network

How successful has the HP Desk implementation effort been in HP?

The best measure of its success is perhaps the intensity with which it is used. In February 1986, nearly four years after the first HP Desk implementation efforts began, there were over 50,000 registered HP Desk users. Among entities, traffic reached over 3.1 million message records a month (a message record has slightly fewer than 2000 characters). Experience indicates that interentity traffic is about one third of an entity's total traffic. Thus, total HP Desk traffic is roughly 18 billion characters a month, about nine million screenfuls of information. The delivery standard of service for worldwide interentity traffic is next working day for regular messages and two hours for urgent mail. It is an order of magnitude better for local intraentity traffic. HP Desk interconnects 483 HP 3000 Computers in 31 countries.

Comparison with other public EM systems provides a yardstick for the order of magnitude of HP's internal HP Desk network. The 15 January 1985 issue of *Datamation* quotes Telemail as having 45,000 users and MCI Mail as having about 150,000 subscribers and a traffic volume of more than one million messages a month.

Subjectively, the success of HP Desk within HP can be measured by the uses to which it is being put. Initially, users availed themselves of the basic functions provided: read, send, delete, file, etc. Later, they began using the more complex functions, including integration with other OA facilities. Finally, users are extending the functions of the product in a way not originally foreseen. They are using it to give better service to their internal customers by setting up specially named HP Desk users to which other departments can channel their queries and receive prompt replies. They are using it to deliver expense reports to managers faster than bursting, collating, and hand delivering paper copies. They are using it to manage multientity projects faster and cheaper than before. They are using it to speed up the targeting cycle across geographical areas. In short, HP Desk has become HP's information distribution and management system, and users are integrating it into their daily jobs in a manner that is productive and personally meaningful to them.

II. Developing and Implementing a Messaging Strategy

Since 1982, OUG's messaging section has been responsible for coordinating the companywide implementation of HP Desk at HP. This section presents the messaging section's strategies, tactics, and experiences.

Since an important aspect of any EM plan is data communications, the first part of this section describes the

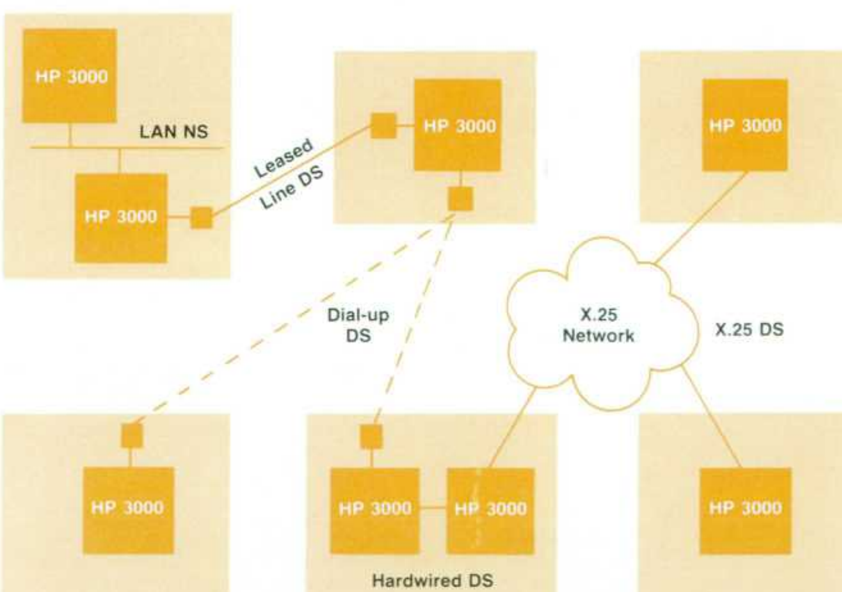

Fig. 2. *Some of the networking options available in an HP Desk network.*

available networking options. The second part chronicles the messaging section's experiences since 1982, which in retrospect, can be divided into four phases, each with its own set of challenges and tactics. Finally, the third part discusses several strategic issues to consider in the long-term planning process.

Although HP's specific strategies and tactics are based on EM implementation in a very large, distributed, highly decentralized company, the phases, issues, and challenges are applicable to EM implementation anywhere.

HP Desk Networking Options

To transfer messages between HP 3000 Computers, HP Desk uses HP's AdvanceNet data communications software, which supports a variety of communication links between computers, including DS (distributed systems) and NS (network services) links between HP 3000s (hardwired or dial-up, bisync or X.25, etc.).

Within a building or complex, *hardwired* lines have most often been used, giving a point-to-point connection with speeds up to 56 kilobits/second. Alternatively, these hardwired connections can be replaced by HP's IEEE 802.3 *local area network (LAN)*, with speeds of 10 megabits/second.

Several options are also available for machine connections between sites. A *dedicated phone line* connecting two sites can be leased. Average monthly costs for 9600 bit/second leased lines, which require modems on each end, are approximately US$600 between San Francisco and Los Angeles, US$1,500 between San Francisco and New York, and US$12,000 between San Francisco and Geneva. Since the cost is fixed, leased lines are best used for relatively constant, high-volume communication between two specific sites.

A second option is the use of *dial-up phone lines* with corresponding modems and autodial units, which usually run at 4800 bits/second. Although the autodialers are optional, the use of dial-up lines in an HP Desk network without autodialers is not recommended. Otherwise, operational and administrative problems are likely to lead to failure of the network. The biggest advantage of dial-up connections is that one computer can connect to numerous other computers (one at a time) using the same HP 3000 communication interface. Since cost is based on connect time, regardless of volume, dial-up is best used for low-to-medium-volume communication that can be queued up for transmission at scheduled intervals.

A third option for connection between sites is the use of a *public data network (PDN)*, using X.25 protocol. With PDNs, a customer pays a fixed monthly charge for each connection to the PDN in addition to packet charges based on data volume. This allows the computer to connect to any other computer that belongs to the same PDN. Like dial-up connections, X.25 PDN connections allow one computer to connect to numerous others using the same HP 3000 communication interface. In addition, the X.25 protocol allows multiple simultaneous connections to different computers through the same port. Since cost is based on volume regardless of connect time, X.25 PDNs are best used for low-to-medium-volume, "bursty," interactive communication.

Because each of these connections provides advantages, most HP Desk networks will use a combination of all of them. This is true for HP's HP Desk network.

Fig. 2 shows several of the networking options available in an HP Desk network.

In addition to supporting DS and NS connections, HP Desk also provides the ability to use other transmission media to move messages between HP 3000s (EFT, external file transfer) and also the ability to interface to other EM systems (FSC, foreign service connection). Both EFT and FSC were critical features in the implementation of HP Desk at HP, and their use is discussed in the next section of this paper.

Besides deciding how best to interconnect the HP 3000s physically, the logical paths through the network must also be planned. Since HP Desk is a store-and-forward system, the optimum network minimizes the number of intermediate machines that a message must pass through to reach its destination. On the other hand, if an HP Desk network consists of several hundred HP 3000s, a fully connected network where each machine directly communicates with every other machine is not administratively desirable. Some compromise between delivery speed and administrative control must be reached.

At HP, the messaging section did not have much control over the design of the physical network. Because of the decentralized nature of both management and information systems at HP, the structure of the network paralleled the structure of the company organization. The overall network is two-tiered, with the bottom layer consisting of the local HP 3000 network at each HP entity and the top layer consisting of one (sometimes more) *gateway* or *hub* HP 3000 per entity, through which all mail into and out of that entity passes. The information systems manager (ISM) at each entity is responsible for the design and performance of the *local* HP 3000 network, while the messaging section

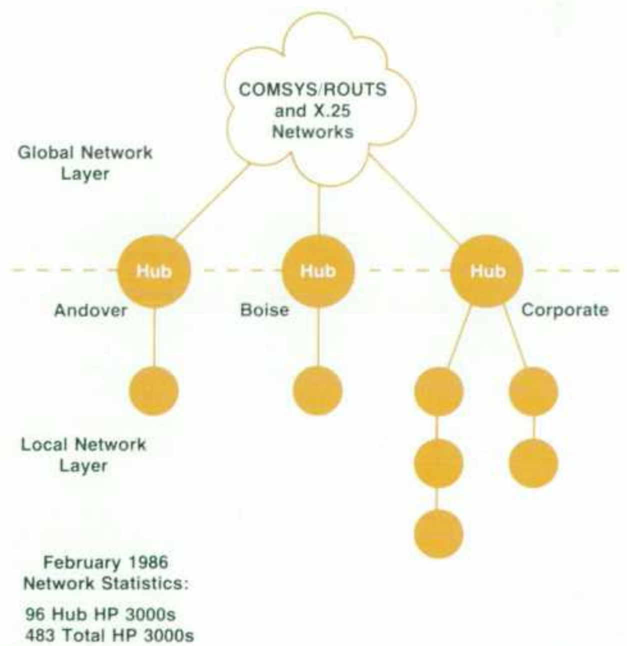

Fig. 3. *Structure of the global and local layers of HP's HP Desk network.*

is responsible for the design and performance of the *global* network connecting the HP 3000 gateways or hubs.

Fig. 3 illustrates the structure of the global and local layers of the network.

While this structure is effective in establishing clear lines of responsibility, it means that the network layout may change when the company organization changes. However, the impact of these changes can be mitigated because of the layered design philosophy behind HP Desk (see Fig. 4).

The top layer of HP Desk is the User Interface module, which is the only part that 99% of the users see. Through the User Interface, they can read, send, reply, file, etc. Messages are addressed to a name at an address called a *mailnode*. The middle layer of HP Desk is the Mailroom module, which takes care of delivering incoming mail to users on a particular HP 3000 and sorting and pigeon-holing outgoing mail. The bottom layer of HP Desk is the Transport module, which is responsible for transferring messages between HP 3000s. The Transport calls DS or NS or uses EFT/FSC. Each layer communicates, via a strictly defined interface, with the layers directly above and below it.

This layered design provides two important benefits. First, HP Desk automatically adapts to changes and enhancements to the data communications software. Thus, the HP Desk network can take advantage of new data communications features with no change to the HP Desk software itself.

Second, because of the well-defined interface between layers, one layer can be pulled out, completely rewritten, and reinserted, and the process is invisible to the other layers. This modularity enables the product to be enhanced to adapt to changing conditions on a timely basis.

Phases of EM Implementation

Since the messaging section started implementing HP Desk at HP, the HP Desk network has gone through four distinct but somewhat overlapping phases, each with its own set of issues. These issues should be anticipated and addressed ahead of time, as part of the strategic and tactical planning process.

Phase I: Startup

The objectives for this phase were to:

- Ensure that reliable networking was already in place
- Establish network-wide standards, conventions, and guidelines
- Appoint and train administrators at each remote location.

In February 1982, the messaging section started alpha-testing HP Desk (then called HP Mail) in three departments at HP's Corporate Offices in Palo Alto, California. Most of the users were sophisticated terminal and computer users and therefore needed minimal training. In the next few months, the messaging section learned a great deal about the product, especially what it took administratively to keep it running smoothly. HP Desk was officially introduced in April of the same year, and in May 1982, the messaging section was asked to head a project to implement an HP Desk network for the former Computer Group product organization. This involved about 20 entities spread throughout the U.S., with about 10 users at each entity. At that time, there were no direct DS connections between these entities.

Although the initial objective of this project was to provide fast HP Desk delivery between all Computer Group functional managers and secretaries, another objective was to experience firsthand what it took to implement a large, distributed HP Desk network. The messaging section worked for the next few months to build this network, working with contacts in the information systems department at each entity to install software, install and configure modems, set up delivery schedules, and configure HP Desk nodes, routes, and users. But because too much was done too soon, the project results were chaotic and nightmarish. There were problems with the modems, problems with the phone lines, problems with configurations, and problems getting cooperation from the diverse entities.

Eventually, the problems were straightened out and the network ran relatively smoothly. In retrospect, two major mistakes stood out from this initial large-scale pilot project. First, manpower expectations were not communicated ahead of time to each entity to get the needed commitment from the local information systems departments. These contacts did not have the time to become trained in HP Desk administration. Lesson number one was that a successful network needs committed, trained, HP Desk administrators at each remote location to resolve local problems and support local users. The second mistake was the concurrent attempts to build an HP 3000 DS network while trying to install and configure HP Desk. The data communications network that HP Desk uses needs to be built, tested, and fully managed before HP Desk starts to use it. In addition to these two mistakes, the initial product lacked certain administrative tools, and this hindered the successful administration of the system. Partly because of feedback to the product division from these experiences, these tools were incorporated into subsequent versions of the product. However, one thing was done right initially: to set up companywide address conventions and configuration standards.

Phase II: Expansion

In the second phase, the goals were to:

- Address the critical mass issue
- Provide good support to remote HP Desk administrators
- Provide local user training and support
- Set proper user expectations and establish good habits.

At about this point in time, the commitment was made

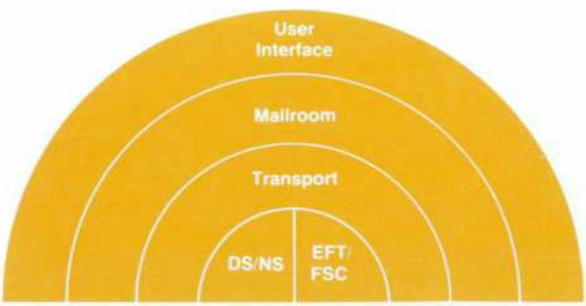

Fig. 4. *The layered design of HP Desk allows it to adapt quickly to changes in any layer.*

by top management to make HP Desk the primary messaging system for HP. Faced with the responsibility for coordinating this implementation, it was time to step back to reevaluate and revamp the existing strategy and tactics. Besides the problems mentioned above, the startup experiences brought up two additional lessons. First, it was beyond the scope of the messaging section, certainly in the time frame hoped for to implement EM, to build a companywide HP 3000 network connecting several hundred HP 3000s via DS. Yet, experiences indicated that the success of EM depends upon a reliable network. Second, like any other type of communication system, HP Desk needs a certain critical mass of users before it becomes extremely useful. A dilemma existed. No one wanted to commit the resources to implement HP Desk unless they could communicate with everyone else via HP Desk, but no one wanted to be first.

How did these problems get resolved? Both issues were resolved by taking advantage of HP's COMSYS/ROUTS data communications network and the EFT/FSC features of HP Desk. COMSYS/ROUTS was cost-effective, already connected every HP entity worldwide, was already managed (by another group), and most important, already existed. The plan was to provide interfaces between HP Desk and COMSYS/ROUTS using EFT and FSC. FSC was used for the HP Desk-to-COMGRAMS (HTC) interface, which allows an HP Desk user to send messages to both HP Desk and non-HP Desk users. Non-HP Desk users receive the message as a hard-copy message, just as before. HTC solved the critical mass problem and encouraged the use of one messaging system—HP Desk.

EFT was used for the HP Desk-via-ROUTS (HVR) interface, which uses the already existing COMSYS/ROUTS network to move HP Desk messages between entities on a next-working-day delivery basis. HVR solved the networking problem and enabled the messaging section to concentrate on implementing EM, rather than building a worldwide data communications network. (Section III of this paper covers the development of the EFT/FSC HP Desk-to-COMSYS/ROUTS connections in detail.)

In January 1983, a memo was distributed to the information systems manager (ISM) at each HP entity detailing the overall messaging strategy and tactics. Local implementation of HP Desk was encouraged and extensive implementation support services (see section IV) were offered. This established a local user base and gave local HP Desk administrators the time to learn how to keep their local HP Desk network running smoothly. Although the messaging section was responsible for the overall companywide implementation, the distributed nature of HP meant that each entity was responsible for its own local implementation and networking of HP Desk. Each entity was asked to appoint a local messaging coordinator (LMC) who would be responsible for the local implementation and administration and with whom the messaging section would work to carry out the companywide strategy and tactics. It is very important to take advantage of the LMC's initial contact with users to start them off with good habits and to set their expectations of message delivery speed properly. Once users develop a certain set of habits and expectations, it is very difficult to change these.

While each entity worked on local implementation of HP Desk (see section V), the messaging section worked on the HTC and HVR interfaces.

During the initial phase of HP Desk network expansion, a 20/60/20 rule was found to apply. Twenty percent of the entities would give HP Desk implementation high priority, do an excellent job of implementation, pioneer EM activities, and be genuinely excited about the implementation. The next 60% would be relatively cautious, wait to see what happened, then proceed slowly but effectively. The remaining 20% would always seem to be short of people, money, and equipment no matter what was suggested. Because of this spread, resources were focused on the first 20%, establishing some successes and providing good examples for the rest to follow. Thus, between January 1983 and July 1984, the HP Desk network went through its expansion phase.

Phase III: Optimization

The goals of this phase are to:

- Measure network statistics
- Improve reliability and delivery
- Provide alternate paths.

By July 1984, two years after the initial Computer Group network project, the HP Desk network numbered 15,000 users. The original Computer Group network had converged into the companywide network. HP Desk traffic within an entity flowed via DS, while most HP Desk traffic between entities flowed via COMSYS/ROUTS (HTC and HVR). HP Desk was being absorbed into the daily lives of many of its users and was beginning to replace other methods of communication.

However, about this time, severe delays began impacting certain parts of the network. Seventy-five percent of the time, messages were reaching recipients in 24 hours or less, but 25% of the time they were taking longer, as much as a week. After researching specific delayed messages, a few bottlenecks in the network were found. They were the result of 1) lack of recommendations for local network configurations and 2) no direct control of the size and number of messages being sent by users. Several months were spent implementing solutions to prevent these bottlenecks.

To avoid having to react constantly to emerging problems, resources were devoted to collect network performance data. Beginning in January 1985, a monthly report has been issued to all ISMs and LMCs worldwide. It contains detailed network-wide data that measures message delivery, HVR performance, and HVR outgoing volumes for each entity in the network. These monthly reports serve several purposes: 1) they are an objective measurement of the performance of the companywide network, 2) comparison of these monthly reports points out current and potential problem spots, 3) the publication of the reports results in a certain amount of peer pressure which encourages the entities with less than acceptable message delivery performance to improve, and 4) the objective data is used to set proper user expectations for typical message delivery.

Also in January 1985, the messaging section introduced a revised HP Desk strategy called *dual gateway networking*. The revised strategy is to use the existing COMSYS/ROUTS network with HVR for delivery of *normal* messages and to use X.25 networking for delivery of *urgent* messages and in selected cases for delivery of *all* messages. The establishment of dual routes between all entity gateways provides a good cost/speed trade-off in addition to backup paths in case of problems in either network.

HP's HP Desk network has continued to grow, and it is currently in the third phase, when resources are concentrated on improving the reliability and performance of the network, providing alternate paths for backup, and optimizing transmission based on volume.

Phase IV: Integration

Activities in the integration phase are to:

- Establish EM etiquette
- Set policies for external access to the internal EM system
- Provide the ability to send mail from information systems applications
- Consult with departments who have special uses for EM.

Although presently in the optimization phase, the network is starting to move towards a new, possibly final phase, when EM, like the telephone, becomes an essential tool and an integral part of everyone's job. In this new phase, employees find unique ways to use EM to perform their jobs and make themselves more productive. As the network has evolved, the typical user questions have evolved from "How do I send a message?" in the startup phase to "How do I send a message to X?" in the expansion phase to "How can I get this urgent message to X in two hours?" in the optimization phase. In the next phase a typical question is "My order coordinator thinks that we can save money, time, and frustration by using HP Desk to resolve problems with orders instead of the current method of playing telephone tag with the divisional order coordinators. How can we set this up?"

Fig. 5 traces the growth of the worldwide HP Desk network over three years.

Strategic Issues—Keys to Success

Based on the messaging section's experiences in implementing HP Desk at HP, several strategic issues must be considered for successful EM implementation.

Obtain management support and commitment. Top-level management support is obviously important in any project. However, it is especially important in EM implementation because 1) EM needs to be implemented companywide to be really useful, 2) the productivity benefits that it provides are often more intangible than tangible, and 3) although planning is usually done at a centralized corporate entity, actual implementation, administration, and troubleshooting are done at remote entities throughout the company. Another management issue is one of mandatory versus optional implementation. Because HP's initial Computer Group network project was mandatory for the participating entities, the important critical mass was established with this core group of entities and users in a relatively short amount of time. The rest of the entities in the company then joined the HP Desk network when they were ready and willing, which optimized the time and efforts of the messaging section.

Address network-wide resources. In retrospect, most of the problems stemmed from the fact that too many remote entities did not commit the necessary resources for a successful implementation: people, money, and equipment. On the other hand, one could not ask for too much at the beginning. Because HP Desk has proven to be an invaluable tool, most managers are now willing to devote the people, time, and money necessary to optimize the network, which they might never have agreed to do a few years ago.

Aim for early success. At the start, it is best to work with the 20% of the prospective user groups who have enthusiasm, will cooperate, and have the highest probability of a successful implementation. Once able to establish early success, the overall project will be viewed favorably, and subsequent groups will be much more receptive to EM implementation.

Establish good relationships with remote contacts. In a distributed network, one must rely on the remote contacts (LMCs) to keep each entity running properly. Since a network is only as good as its weakest link, it is important that the LMCs be provided with the proper training and

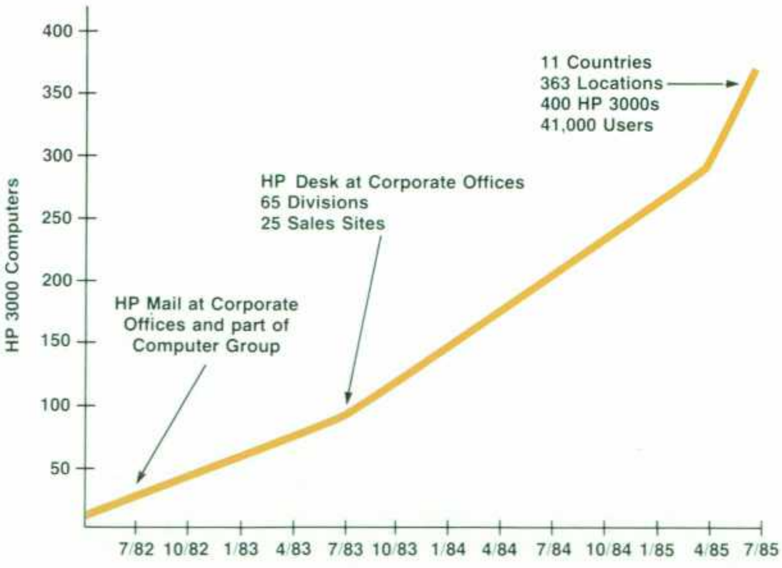

Fig. 5. *Growth of the worldwide HP Desk network over three years.*

level of support that they need to get their job done well.

Follow-up, fine-tune, forever. Work is never done. The network goes through a continual cycle of follow-up and fine-tuning as the product, user sophistication, message volume and flow, and communication technologies change. A successful EM network is one that has the proper structure in place—people, software, and hardware—to adapt as quickly and painlessly as possible to the changes forced upon it.

III. EFT and FSC

In 1983, a very significant enhancement was made to HP Desk to provide the ability to interface with other transmission media and foreign mail systems. This feature has been extensively used within HP's HP Desk network. Because of the size and worldwide nature of the network, this feature has proven essential to the network's success. The EFT (external file transfer) and FSC (foreign service connection) facilities of HP Desk are used on entity gateways to interface with the COMSYS/ROUTS network. This section of the article discusses the technical details of this feature and how it is being applied within HP.

EFT/FSC and the HP Desk Transport Manager

The Transport module of HP Desk has the responsibility to forward all messages having an address that is not local to the HP 3000 on which it is operating. Because HP Desk is a store-and-forward system, a message may make multiple stops through intermediate HP 3000s in the network before reaching the message's recipient. Each organizational entity is encouraged to collect all nonlocal mail bound for other entities on one HP 3000, occasionally more than one. These HP 3000s, called gateways, are an entity's principal interface to the companywide HP Desk network. A gateway is typically an HP 3000 designated to perform the network gateway role, usually in addition to performing its primary role of providing local business systems support.

The local messaging coordinator (LMC) at each entity uses a menu-driven HP Desk utility to configure a route for each mailnode. The route is the destination computer name or computer gateway used by the Transport to deliver a message to the *next* computer on the message's path to its final destination. The route may be a specific data communications line, an X.25 address, or a special EFT/FSC computer name. Every 15 minutes the Transport examines the availability of a route associated with mail waiting to be forwarded. The availability of a route is configurable on a 24-hour clock in 15-minute windows. Messages can be prioritized by the sender as either normal or urgent. The Transport prioritizes all outstanding mail based on the priority and the number of messages destined for a specific mailnode.

When mail is scheduled for transmission to an EFT/FSC computer, all mail destined for a specific mailnode is written to a data file. One file is created for each mailnode. For example, if messages for 10 mailnodes are processed, the Transport produces ten files, each containing all messages for one mailnode. For each file created, a record is written to a special file which acts as a queue of all EFT/FSC files produced. This file is called an IPC (interprocess communication) file. IPC files can be read and written to simultaneously. IPC files also allow an application to issue a read to an IPC file and wait until data is written to the file. Use of this feature allows dependent events to occur in immediate succession. IPC files are used extensively by HP Desk for interprocess communication between and among its User Interface, Mailroom, and Transport modules. Fig. 6 illustrates the interface between HP Desk and the EFT and FSC facilities.

Each record of the IPC file designated for use by EFT/FSC contains the name of the corresponding data file, the format type of the file, and the mailnode associated with the data file. The four types of formats include HP Desk internal format (EFT) and three types of ARPA[2] formats (FSC). EFT can be used when the destination is an HP 3000 and the transmission medium is a nonproduct data communications network. FSC format includes three different ARPA

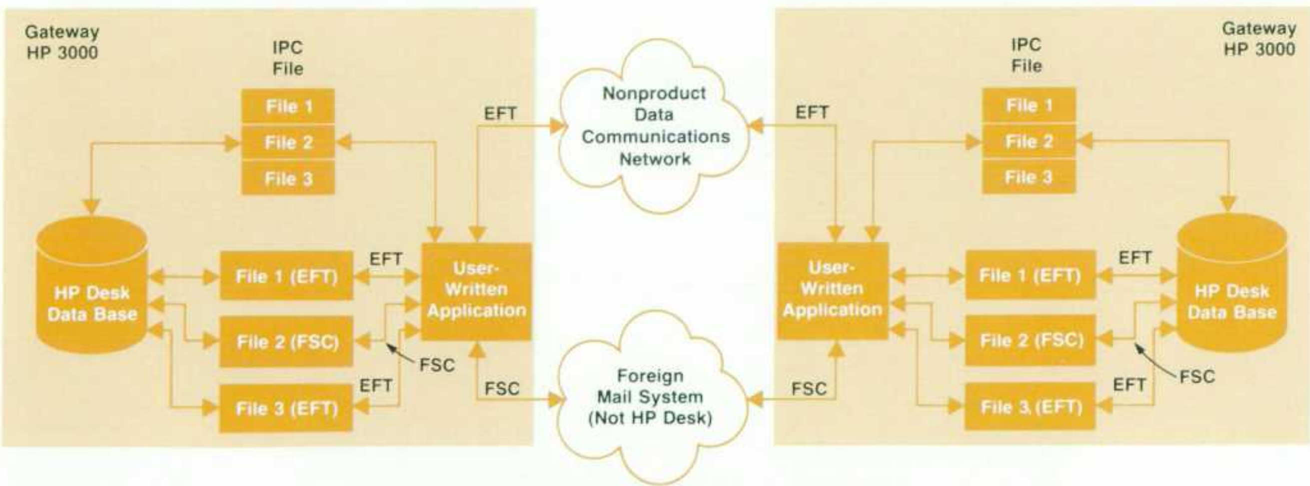

Fig. 6. *How HP Desk uses EFT (external file transfer) and FSC (foreign service connection). An IPC (interprocess communication) file acts as a queue of all EFT and FSC files containing messages for various mailnodes.*

standard formats. The three types of ARPA formats allow the user to transmit ASCII data only, both binary and ASCII data, or a compressed version of ASCII data only. ARPA is an internationally recognized standard format used for messaging. The type of format conversion done by HP Desk depends on how that address has been configured. For example, within HP, a specific mailnode per location is always configured by the LMC to be routed out the FSC gateway. EFT supports the transfer of messages in HP Desk's internal format across nonproduct data communications networks to remote HP Desk data bases. FSC supports the transfer of messages to foreign mail systems using ARPA standard formats. Both EFT and FSC require that a simple application program be written to provide the interface with HP Desk. HTC and HVR, discussed in the following sections, provide a good example of utilities using the EFT and FSC facilities. Fig. 7 illustrates the interface between EFT and FSC and such user-written applications.

HTC

In July 1983, HTC (HP Desk to COMGRAMS) was introduced and was eventually installed in over 90 locations within HP. Design and coding of HTC took six engineer-months.

HTC uses HP Desk's FSC facility to provide an interface to COMSYS/ROUTS for ASCII messages created in HP Desk. HTC reads the IPC file associated with the FSC gateway for outgoing HTC messages. This IPC file acts as a *directory* for all files containing ASCII messages. The data files associated with each IPC file record are then reprocessed into a format compatible with COMSYS/ROUTS and loaded into the COMSYS/ROUTS network via tape or a data communications line. At the receiving entity the messages are printed on a line printer and distributed as hard copy.

HTC represented a significant milestone in HP's EM network development, because by using the FSC facility of HP Desk in this fashion, users were enabled to send messages to anyone in HP. HTC allowed HP to leverage its existing data communications network and an already existing messaging system. This also allowed the phasing out of older messaging systems and the promotion of HP Desk as HP's companywide EM system.

HVR

In December of 1983, HVR (HP Desk via ROUTS) was released for use within HP. HVR was designed and coded in six engineer-months. By October 1985, HVR was installed in over 90 locations within HP. HVR incorporates the features of HTC and provides the additional capability of sending messages from HP Desk to HP Desk (terminal to terminal) and not just from terminal to paper (HTC).

The use of the EFT facility allows HVR to transmit any type of data contained in an HP Desk message through COMSYS/ROUTS. HP Desk allows the transmission of almost any HP 3000 file, including word processing, graphics, spreadsheet, and personal computer generated files, with the exception of certain privileged files, such as Image data bases.

HVR consists of modules for exporting messages to and from COMSYS/ROUTS and loading them back into the local HP Desk network. The outgoing module of HVR is similar to HTC's, with the exception that the IPC file used by HVR to locate the files queued by HP Desk contains the names of files that could be in either EFT or FSC format. The HVR outgoing module reprocesses these files and formats them for transmission through COMSYS/ROUTS.

The HVR incoming module mirrors the HVR outgoing process. The HVR incoming module processes messages transmitted from an EFT gateway and imports them back into HP Desk. The recipient of the message sees no difference between a message sent locally and one that flows in from outside using EFT and COMSYS/ROUTS. The HVR incoming module has the responsibility for importing the messages from COMSYS/ROUTS and breaking up the data into data files. Each data file contains one or more messages that have been sent from another entity to a local HP Desk address. After each file is created, a record is written to an IPC file which serves as a directory for HP Desk of all data files containing messages to be imported into HP Desk. Once all HVR processing has been completed, an HP Desk utility program is run which uses the IPC file to locate each data file and load it into the local HP Desk network. The EFT and FSC facilities are powerful extensions of HP Desk which allow interfacing to other transmission media and foreign mail systems.

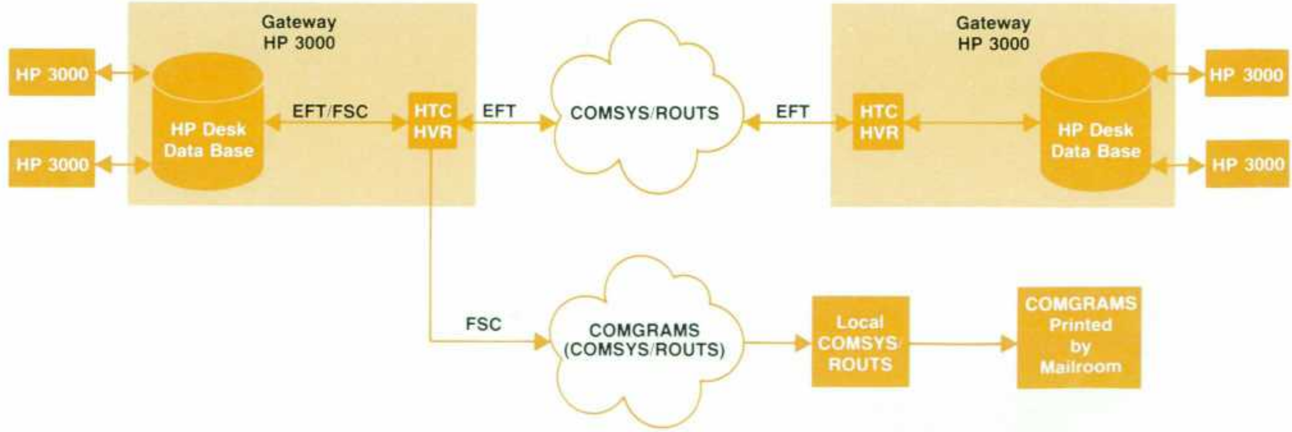

Fig. 7. *HTC (HP Desk to COMGRAMS) and HVR (HP Desk via ROUTS) are examples of user-written utilities that use the EFT and FSC facilities of HP Desk.*

Other EFT/FSC Applications

In January 1984, CTH (COMGRAMS to HP Desk) was introduced by OUG's messaging section for use within HP's Corporate Offices only. Before CTH, ASCII messages (comgrams) sent to the Corporate Offices were printed on a line printer and distributed as hard copy. This manual process often delayed message delivery one or more days. Using CTH, these messages are instead imported into HP Desk using the FSC incoming gateway. CTH takes incoming messages in the COMSYS/ROUTS comgram format, translates them into a documented ARPA standard format, and then writes a record to an IPC file used by HP Desk as a directory of files containing messages to be imported. Once the message files are queued, an HP Desk utility program is run to read the IPC file and import the messages into the local HP Desk network. CTH in many ways mirrors the incoming module of HVR.

The FSC feature of HP Desk is also available on a programmatic level. Programmers can import data into and export data from the HP Desk data base using standard HP 3000 files and IPC files in exactly the same manner as discussed earlier. Both HTC and CTH are good examples of such programmatic access.

HP Telex also uses the FSC gateway. HP Telex is an HP supported product that provides an interface to major U.S. and international Telex vendors. It is certified in over 16 countries and provides the ability to send Telex messages from either a program running on the HP 3000 or directly from HP Desk. The FSC gateway is used by HP Telex to transfer the message from HP Desk to an HP Telex program which has responsibility for managing the Telex lines and message flows.

FSC is also used in a project to interface UNIX™-operating-system-based EM systems within HP with the internal HP Desk network. The interface uses FSC for both importing and exporting messages, very much like HTC and CTH.

EFT/FSC Summary

EFT and FSC make up a powerful and versatile feature of the HP Desk product. Certainly, HP's HP Desk network would not have grown as quickly as it has if HP had not been able to leverage its previously existing data communications network (COMSYS/ROUTS). EFT and FSC also transparently allowed HP to converge its existing messaging systems into HP Desk. HTC, HVR, and CTH are all utilities that the messaging section wrote using the standard HP Desk product and documentation. They are intended for internal use only and are not available for use outside HP. Customers who wish to interface with their own data communications networks and messaging systems can take advantage of the EFT/FSC facilities to write their own customized utilities to meet their particular needs.

IV. HP Desk Support and Training

This section of the paper discusses the types of HP Desk support and training needed in an organization to be successful at growing and maintaining a large-scale, distributed HP Desk network such as HP's.

UNIX™ is a trademark of AT&T Bell Laboratories.

Support and training include both support and training for local messaging coordinators (LMCs), who are in charge of supporting the entities' local HP Desk networks, and support and training for end users, who need to know how to use the product and take advantage of its many features.

Why is support necessary? For an EM network to be successful, it must deliver *availability* (up to 24 hours a day, 7 days a week), so users can read and send mail on their own schedules. It must also deliver *reliability* (message delivery within a guaranteed time), so users have confidence in the system and therefore will use it. Lastly, it must deliver *usability* (not just ease of use but knowing how to use the EM product for maximum value), so users see a definite benefit in EM. Having the proper support network in place is the key to delivering these three ingredients and thus ensuring a higher probability of a successful network.

Defining the Support Issues

Developing and then documenting a support plan is essential to achieve a clear definition of the support issues. Although the support plan may change over time, a written plan forces the organization to address at least the obvious and critical issues. Support responsibilities must be addressed early in the plan. The main categories of support include technical, administrative, and user support.

This section of the paper takes a closer look at each of these support roles and their manpower requirements and investigates how HP has coordinated and enhanced these support roles over a worldwide network.

Critical Support Responsibilities

The area of *technical support* is critical because its mission is to resolve highly visible and complex technical problems such as corrupt data bases, program aborts, and other problems that impair or disable user access and mail delivery. The occurrence of technical problems is rare, particularly if the day-to-day administration tasks (discussed next) are performed regularly. However, technical problems are highly visible and, if not taken care of immediately, can lead to additional and more complicated problems. Therefore, each entity in the network needs to identify and have trained one primary person (or LMC) and one backup person for HP Desk technical support. The ideal person to fill the technical support role is a programmer/analyst with a good knowledge of the HP 3000, Image, Query, and DS and NS. In a large HP Desk network, the time requirements on this person vary directly with the *number of HP Desk data bases* at the entity, the *number of users per data base*, and the *quality of HP Desk and HP 3000 administration* at the entity.

The area of *administrative support* is also critical since this role addresses the daily required duties. Some of these duties include: 1) ensuring that the HP Desk production jobs are up and running at least during business hours, 2) ensuring daily maintenance job runs and checking the daily HP Desk reports for warnings and errors, 3) monitoring and forwarding, or printing, messages that have landed in General Delivery (these are messages that for one reason or another could not be delivered to the intended recipient), and 4) maintaining and updating local addressing directories. Typically a systems administrator is a good match for this role. In a large HP Desk network, the time required of this person varies directly with the *number of HP Desk data bases* at the entity and the number of *users per data base*.

Finally, *user support* is needed to ensure that all users know how to use HP Desk, use it on a regular basis, and use it both properly and effectively. Typical duties of an HP Desk user support person include conducting training classes, providing a call-in or write-in hotline for questions and suggestions, and consulting with departments for tailored uses of EM to help them do their jobs better. An office automation coordinator (OAC) is best suited for the user support role. Staffing needs for this role are directly related to the *total number of users* at the entity.

Fig. 8 details the relationship between the magnitude of the local HP Desk network and the support staffing requirements. It is important to stress that the actual total number of people required to support HP Desk at a single entity varies widely and is related to both the number of HP Desk data bases and the number of users (per data base and total) at an entity. The number of support people in place, the time they are able to devote to HP Desk, and thus their rate of success, also depend on local management's interest in and commitment to EM.

Many entities within HP have several hundred users distributed over several HP 3000s. In some cases, these entities appoint one full-time person to cover all three areas of HP Desk support responsibilities. Other entities have elected to distribute the responsibilities over several people, each of whom works part-time on HP Desk. For example, one successful entity with 500 users and three HP Desk data bases has a software support analyst to address HP Desk technical problems (this accounts for 10% of this person's time), a systems administrator to handle the administration (accounting for 20% of this person's job), and an OAC to cover training and user support (HP Desk amounts to 30% of this person's job).

What matters with respect to HP Desk support at an entity has little to do with the number of people involved and a great deal to do with the assurance that all three support roles are covered at least during business hours and every business day. Clearly defining these responsibilities, making them an official part of someone's job, and identifying official backups help achieve this goal. Hence, commitment breeds success.

Once the support plan has addressed and defined the support responsibilities, it should address how to coordinate these local responsibilities across the network.

Network-Wide Support Strategy

The staffing requirements in Fig. 8 are for entities' local HP Desk networks and they complement the decentralized support strategy adopted by HP. One other option for coordinating network-wide support would have been a single centralized support group. Because of HP's internal distributed information systems strategy and its highly decentralized management structure, the decentralized HP Desk support strategy had a higher probability of success and was easier to implement than a centralized support strategy.

With decentralized support in place, HP's network enjoys the advantages of more timely identification and correction of problems as well as availability of support contacts for the local users. On the other hand, decentralized support requires more people and offers less consistent operational control of the network.

Fig. 9 illustrates the network support structure within HP. Each entity has identified one primary local messaging coordinator (LMC). Sales regions and multientity sites often appoint a group messaging coordinator (GMC), whose responsibility it is to coordinate support efforts among the many LMCs within the region or site. Together, these HP Desk support people manage HP Desk for their entities and implement new strategies and EM projects as recommended by OUG's messaging section or as developed locally. In an effort to make the support network an effective one, the messaging section provides a variety of special support services for the LMCs' use.

Network Support Services

Network support services include a variety of projects intended to help LMCs carry out the responsibilities of their job most effectively and reduce duplication of effort. These services are provided only internally by OUG's messaging section and are made available exclusively to LMCs

HP Desk Support Responsibility	Time Requirements (by Magnitude of Local Network)		
	1-5 Data Bases 100-250 Users	5-15 Data Bases 250-500 Users	Over 15 Data Bases Over 500 Users
Technical (Programmer/Analyst)	10% (of one person's time)	15-20%	20%+
Administrative (Systems Administrator)	15-25%	50-75%	75%+
User (OA Coordinator)	15-25%	30-40%	40-50%+

Fig. 8. *HP Desk support requirements.*

on an on-going or as-needed basis.

LMC Technical Support Classes. To educate LMCs on the technical aspects of HP Desk and HVR data base and program internals, operations, troubleshooting techniques, and network philosophy, the messaging section provides technical support classes. This class has become a vital part of HP's network support strategy. It has proved to promote consistency and quality of support at all local entities. The classes are offered approximately every three months.

LMC Support Line. Incoming calls and messages to this phone number are monitored by support people in the messaging section. The phone number is for LMCs who need immediate help in solving an HP Desk problem.

LMC Bulletin. The messaging section sends this bulletin to all LMCs. It includes items such as performance tips, helpful hints on user, technical, or administrative support, "problem and solution of the month," etc. LMCs are encouraged to contribute their own articles to this bulletin.

LMC Support Groups. LMC support groups (made up of internal HP Desk support people) have sprung up in various parts of the world. They are self-organized and meet regularly for the purpose of sharing ideas and addressing common issues. Typically, meetings are held once a quarter for two to three hours.

Training Material. It is important that entities offer formal training classes for the EM users. Some users prefer formal classes to self-paced courses, and many times the only opportunity for a user to learn a new software package is to get away from the office and enroll in a class for a few hours. The messaging section has helped by providing LMCs with an HP Desk training package, which includes trainer's notes and outlines for beginning and advanced HP Desk classes, class materials, overhead slides, and lab exercises.

HP Desk Network Status Line. An LMC calling the HP Desk network status line hears a recorded message (updated twice daily) which informs the caller of existing or recent (as far back as one week) message delivery delays in the network. Before hanging up, the LMC can also leave status information regarding the LMC's own entity. The purpose of the status line is to help LMCs track the flow of a message to any entity in the network. By calling the status line, an LMC finds out if there have been any delays or data losses at the receiving entity. At the same time, the messaging section stays well-informed on the state of the network on a daily basis.

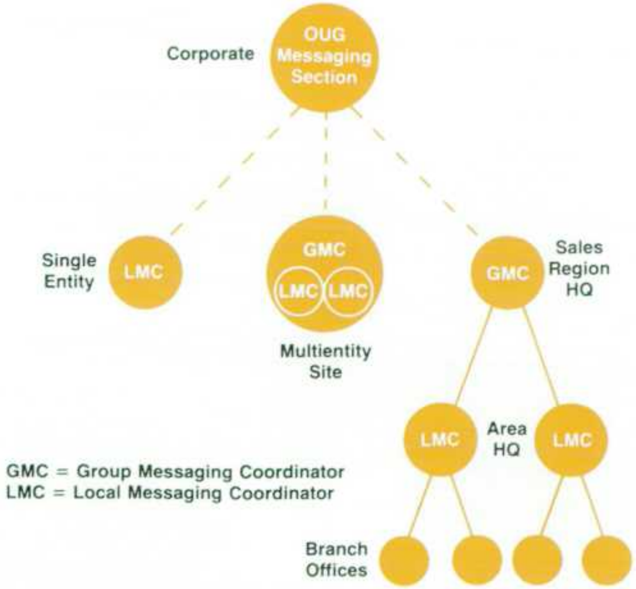

Fig. 9. *Hewlett-Packard's HP Desk support structure.*

HP Desk User Training

Since an EM network is of little value unless a critical mass of users is using it, user training is extremely important. Also, since EM is unlike any other OA tool in that it potentially involves nearly all individuals in an organization, a good user training and support program is crucial to its success.

Before HP Desk became the primary messaging tool for HP, many entities already had at least one part-time OAC whose job it was to train and support users on various OA utilities. HP Desk was added to the array of products for which OACs assumed user support. Entities that had not yet hired or appointed an OAC eventually found the introduction of HP Desk to be the ideal time to do so.

The mission of the OAC serving as HP Desk trainer is to ensure that the user community knows how to use HP Desk, actually uses it on a regular basis, uses it effectively, and obtains enough product knowledge over time to become creative in identifying new ways to increase productivity.

Larger entities faced somewhat of a problem with only one OAC, since it is potentially impossible for a single individual to offer good user training and support to more than a thousand users. So these larger entities expanded their training and support base beyond the single OAC by enlisting workgroup or departmental OA coordinators (DOACs). The ideal DOAC is one who has a good understanding of the department's business, visualizes advantages of using OA products within the department, and preferably has some influence over the direction of the department. In some cases, these DOACs are departmental supervisors; in other cases, they are department secretaries.

HP Desk Training Options

Since EM will potentially reach most, or even all, members of an organization, it is important to have a variety of training options for them, since different people have different training needs and respond to different training methods. What follows is a brief discussion of the most important training tools.

On-Line Interactive Training Facility. The HP Desk software includes an on-line interactive training facility (ITF) composed of several modules, each pertaining to a different area within HP Desk (for example, the In Tray, Out Tray, Pending Tray, etc.). Each module takes approximately ten to fifteen minutes to complete, and the user can do as many or as few as needed. This mode of training is ideal for users who are already familiar with the keyboard, do not require personal training, and prefer to learn on their own. It is also an excellent mode of training for those who cannot get away from their desk to attend a formal class or need to go through training during off-hours.

Formal Classes. Many entities have elected to offer formal training classes to users who want them. Experience indicates that a nominal fee should be charged those who enroll. Certain groups tend to benefit most from formal classes: department secretaries, who need to know the more advanced or complex features of the product, DOACs, who are to become the front-line user support persons for their departments, and those who need to get away from their desk to get the training.

Customized Training Classes. Special groups may need something other than the ITF or formal training classes, and customized classes work well in these cases. For example, a customized class may be offered to executives.

Help Facility and Reference Guides. These are good options for those who prefer to learn by using. The on-line help facility within HP Desk is excellent. At virtually any prompt within HP Desk the user can type HELP or ? or HELP <command name> and receive on-line instruction. Information systems personnel who are familiar with the HP 3000, or who otherwise prefer to explore a new product on their own, may want to train themselves via this option.

A good training plan puts together the available options in a way that meets an entity's individual needs and priorities. For example, one HP entity with over 300 target HP Desk users had but one part-time LMC who was responsible for user, technical, and administrative support on HP Desk as well as technical support on several other production applications at that entity. Having no time to conduct formal training classes on HP Desk, this LMC sent out an announcement to the 300 target users to announce HP Desk and explain how to sign on to the ITF. Users were asked to contact the LMC after they had run the ITF and, following a brief quiz, were set up as registered HP Desk users.

Getting Users to Use HP Desk on a Regular Basis

Users are accustomed to checking their paper in-baskets on a regular basis,and since EM senders depend on EM receivers to read their mail, it is critical that EM users assume responsibility for signing on to HP Desk on a regular basis to check their electronic in-basket. At a minimum, users should sign on three times a week. For many newcomers to EM, regularly signing on is a whole new behavior that must be learned, and it is up to the OAC to help instill this new, required behavior. The following actions may help achieve this objective.

Teaching Messages. One LMC created ten teaching messages, each describing a different feature of HP Desk. Every message asked the recipient either to reply or to perform some other action (such as file it in a newly created folder). Every time a new department of users was added to the local network, this LMC would send a new teaching message each day to each new user, for ten days. The teaching messages not only served to inform users of various HP Desk functions, but they also coaxed new users into the habit of signing on to read their mail at least once a day.

Monitoring Sign-On Frequency. The HP Desk maintenance report details a number of statistics, among which is the user sign-on rate. For each user local to a specific data base, the report gives the date and timestamp of the last sign-on, which tells the OAC who is signing on regularly and who is not. OACs can then personally follow up with users who have not signed on at all for one to two weeks.

Promoting Effective Use of HP Desk

While training users to use HP Desk is the first step, and getting users to sign on regularly is the second, the third step is to promote effective use of HP Desk. The following actions help achieve this objective.

Disc Space Use. One issue to be addressed with any EM product is disc space use. EM systems encourage electronic filing, which usually means increasing disc space requirements over time. The maintenance report of HP Desk details disc space use by user and therefore permits the LMC to monitor use. Many entities within HP have set a limit on disc space use and then monitor the user community based on that limit. At one entity, the set limit is 10,000 sectors per user, and most users there consume between 4,000 and 6,000 sectors. Users may need some initial help in setting up folders and understanding how to limit the number of items they hold electronically to what is necessary. Many users may not realize that everything within their HP Desk trays is saved until specifically deleted by the user. Training users early on how to make effective use of the Filing Cabinet and how to delete unneeded items will help avoid disc space problems.

Setting Up a Personal HP Desk Password. To ensure a certain degree of security for both the sender and the receiver, users should be taught how to set up a password and change it. It is easy to develop a utility that scans all local users in a data base and points out those who do not have passwords. The OAC should follow up with those users to correct possible security breaches.

Greater Benefits with Less Time and Effort. Once users have mastered the basics, the OAC should begin to implement special programs which can take full advantage of more advanced productivity gains. Interfacing batch jobs to HP Desk, letting HP Desk do manual chores, and setting up special HP Desk user names for specific tasks are some of the myriad possibilities.

V. Implementing and Running HP Desk at an Entity

This section of the paper discusses what is involved in implementing and running HP Desk at an entity. It offers guidance on such topics as starting a pilot project, intra-entity networking, and operational standards and conventions.

This section of the paper is largely based on OUG's experiences in implementing and running HP Desk at HP's Corporate Offices since February of 1982. The number of HP Desk users at the Corporate Offices has grown to over 1700 in February 1986. During the same period, the number of HP 3000s on which HP Desk is installed has grown to 27, and the number of messages sent outside the Corporate Offices has grown to over 59,000 a month.

Getting Started—Implementation Checklist

There are six major actions to take in HP Desk implementation at an entity. They are:

- Choosing the implementation and support team
- Choosing the pilot group
- Planning the local entity network
- Addressing operational issues
- Expanding the user community
- Promoting the use of HP Desk.

Choosing the Implementation and Support Team

Four major roles must be addressed by the HP Desk implementation and support team. They are: creation of the implementation plan, administrative support, technical support, and user support. This does not mean that four people are necessary. The size of the Corporate Offices team has ranged from one to three persons. A two-person team has worked best, with one person assigned technical and administrative support (the equivalent of the LMC), the other assigned user support (the equivalent of the OAC), and both members contributing to the implementation plan. In a smaller entity one person would have been sufficient. Briefly, the responsibilities of the implementation and support team include:

1. Implementation plan. Members of the team should create a plan addressing the six areas listed above (covered in more detail in this section of the paper). The plan should document the implementation stages, including the expected time to completion. As implementation proceeds, the plan should be updated to reflect new action items and goals.
2. Administrative support. This role has daily requirements, and as the entity's network grows, this role becomes more time-consuming and more critical. Specific goals include: ensuring that HP Desk is available to users, monitoring message delivery, monitoring General Delivery, checking administrative reports, keeping data bases in synchronization with the rest of the network, and preparing a disaster recovery plan.
3. Technical support. This role requires someone with a strong technical background including knowledge of the HP 3000, Image, Query, and DS and NS. This person will be called upon to solve problems, resolve data base corruptions, and install software upgrades as needed. Local technical expertise ensures a more reliable entity network.
4. User support. This person will train new users and answer user questions. The most important function is to educate the user community on how to use HP Desk to increase productivity.

Experience with HP Desk has shown that it is critical to commit the proper resources to supporting it. HP Desk is a production system and must be treated as such.

Choosing the Pilot Group

One component of the implementation plan should be a pilot phase, allowing support personnel to gain controlled experience with several small groups of users.

A suggested profile of a pilot user group is as follows:

- Strong communication ties
- Combination of secretaries, managers, and professionals
- Twenty to thirty users
- Access to terminals or personal office computers
- Interested and able to devote time to the project
- Minimal training needed.

The most important consideration in choosing members of the pilot groups is their communication paths. Groups should be chosen that need to communicate with each other. It is also important to look at hardware availability, enthusiasm, training needs, and the extent of management participation. Users generally do not like to go to shared workstations to read their mail. They much prefer to have their own workstations. While this may not always be possible, it should be kept in mind. It is also important to choose groups that want to participate. Choose groups that are excited about the product and that can act as "ambassadors" during the expansion of the user community. The pilot phase is a learning time for the support people as

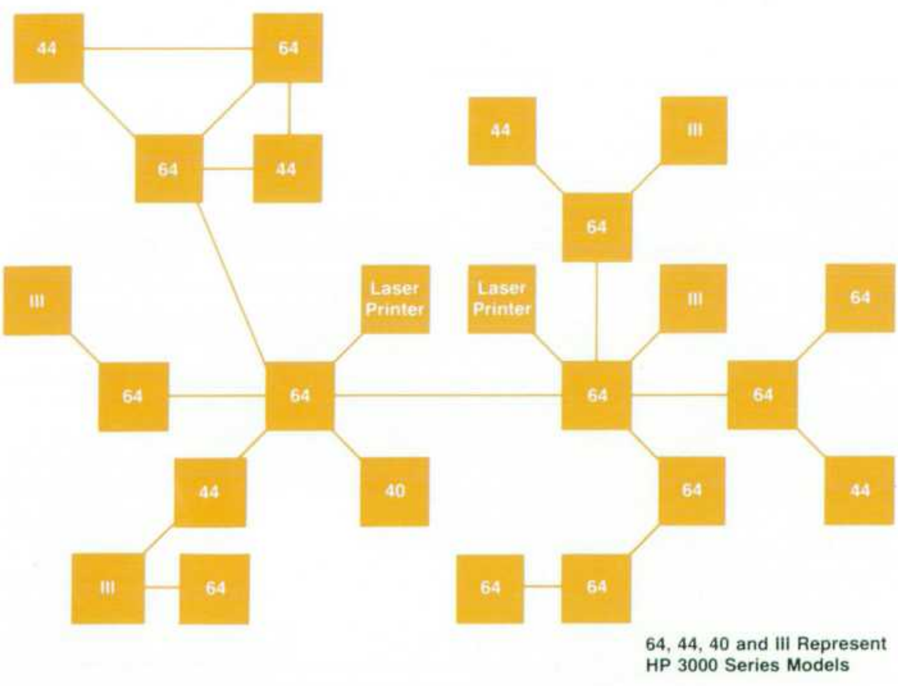

Fig. 10. *In the original computer network at HP's Corporate Offices, the two HP 3000s in the center acted as gateways to the companywide HP Desk network.*

well as for the users chosen. A group that is already familiar with using terminals, logging on to the HP 3000, and running programs makes training a much easier task. Finally, make sure the group's managers are involved. Their support is very important and helps ensure a smooth implementation. Their example as users is invaluable in getting other members of the group started.

The pilot phase should continue for a predefined period of time. Success at this stage will establish user community support for later phases. Build on that success and save the less receptive groups for later.

Planning the Local Entity Network

Before user community expansion can begin, members of the HP Desk team must look at networking options and know which departments use which HP 3000s. How mailnodes will be named must also be decided. At this stage it is essential to work closely with the EDP center. OUG is not part of the Corporate Computing Center (CCC), but it has worked with CCC and adjusted plans based on their input. Many decisions related to network topology and operational issues were made jointly or were based on established CCC policies and procedures.

If HP Desk is to be used effectively and fully, it is important to integrate it into each person's daily work. One of the most useful features of HP Desk is the ability to send any HP 3000 file, but to take best advantage of that feature, HP Desk must be installed on the user's home HP 3000. Users usually dislike having to use one HP 3000 for most of their work and another for HP Desk.

When HP Desk implementation began, there were many groups that did not use any HP 3000; HP Desk was to be their first application. The original plan was to put most users on the same HP 3000. The HP 3000 chosen at the Corporate Offices for new computer users quickly became overloaded. HP Desk response time and session slots were problems during peak periods such as early morning, right after lunch, and midafternoon. Everyone wanted to read messages at the same times. It soon became evident that HP Desk would have to be installed on most HP 3000s, with new computer users distributed among them.

Only one mailnode is necessary per HP 3000, but early in the implementation process, OUG decided that as each group was added, it would be given a unique mailnode. It was decided to assign two-digit numeric sublocations, with the first digit signifying the functional area the group reported to. (Mailnodes in HP Desk consist of six characters indicating a location—an entity, in HP's case—and two characters indicating a sublocation, i.e., XXXXXX/XX.) For example, sublocations in the 30s are in personnel, in the 50s in marketing, in the 70s in administration, and so on. As implementation proceeded, one could check against a list of all departments and judge the progress to date. Giving each group its own mailnode has also proved essential in the follow-up process. Each mailnode has a contact person that OUG communicates with frequently and regularly.

Until about mid-1984, there were virtually no network topology options. CCC had connected the HP 3000s in a fairly linear manner with hardwired links, as indicated in Fig. 10. When a new group was added, the HP 3000 they used had to be directly connected to another HP 3000 already running HP Desk. This product requirement, coupled with the network design, influenced the order in which some groups were added. In one case it was necessary to install HP Desk on an intermediate HP 3000, running the Transport module only, to connect HP 3000s farther back in the chain. No local mailnode was added there for several months.

In CCC's original network (Fig. 10), the two HP 3000s in the center, both laser print stations, became the center of the Corporate Offices HP Desk network. Both have connections to other parts of the network, outside the Corporate Offices, through HVR and X.25; they are the entity gateways. Originally, HP 3000s on one side of the network in Fig. 10 sent their messages out through one gateway, while those in the other half of the network sent their messages out through the other gateway. With this topology, minimal alternate routing was feasible. If an HP 3000 in the center of the network went down, one half of the network was cut off from the other half. Messages traveling from users on one HP 3000 had to make many intermediate stops to get to their destination and delivery times within the building were barely acceptable.

About mid-1984, CCC installed an X.25 switch, enabling each HP 3000 to connect directly to every other HP 3000, as indicated in Fig. 11. The routing on each HP 3000 was changed so that some HP 3000s use one gateway and the rest another, so the load between the gateways is reasonably balanced, and each Corporate Offices HP 3000 is directly connected to every other Corporate Offices HP 3000 for the delivery of local mail. The local entity network is currently evolving towards HP's IEEE 802.3 local area network, a better and more appropriate means for local networking.

Addressing Operational Issues

From the user's point of view, the success of HP Desk implementation depends on operational success. Their most important criterion, against which the HP Desk team is always judged, is HP Desk availability. The goal at the Corporate Offices is to have HP Desk available by seven o'clock in the morning. It is met about 90% of the time. HP Desk is then available until the HP 3000s are taken down at night for system backup.

Operational issues cover more than just availability. Each entity must plan for daily report checking, data base backups, security, and disaster recovery. When HP Desk was installed on a small number of HP 3000s, the support team logged on to each one every morning to check the current status. The present network size makes that impossible. Depending on departmental contacts to point out problems was not consistently reliable. The procedures now in place work very well. Before seven o'clock each morning, the operators on grave shift have a checklist to complete. They are to check whether or not HP Desk is running on each machine and make any necessary comments. Additionally, a job is run each morning that mails a message to the technical and administrative support members of the HP Desk team. When the message is received, it is known at what time HP Desk became available to users by checking the creation time of the message. It is a double-check against the operators but also an alert for network problems. If the operators have checked off that HP Desk is running but no message was received, it is a signal to investigate the link between the two HP 3000s in question.

To assure uptime during the day, the startup procedure of each HP 3000 runs a job that starts HP Desk with a minimum of human intervention. The first time the machines are brought up in the morning, each system runs an HP Desk maintenance job. This job stores the data base onto tape. This data base backup is essential in the event of unrecoverable problems. The data base also exists on the system backup tapes, but having both the system backup and the data base store tapes ensures that at least one set is accurately labeled and free of errors. It is also much faster and easier to restore from a single tape containing only files related to HP Desk rather than from multiple tapes containing copies of all system files. Each data base store tape is kept a minimum of seven days, according to the requirements of HP's internal auditing department. The maintenance job also checks the data base for corruptions and purges those items a user has previously marked for deletion. Running this job daily is the only way to ensure that the data base is kept error-free. Depending on the size of the data base, this job may take from one to several hours on a lightly loaded HP 3000. If a system should crash during the day, it is necessary to run HP Desk's recovery program to make sure no items in the data base have become corrupted. This job completes its checks and has HP Desk running again in under five minutes.

The shutdown procedure of each HP 3000 has a job that cleanly turns HP Desk off. This is used in the event of planned downtime as well as each evening, before the HP 3000s are taken down for system backups. The shutdown job creates a disc file with information about any messages that may not have been delivered yet. That disc file is copied into the message each system sends the support team in the morning. The morning message is created by jobstreams written to help automate the administration of the Corporate Offices network. MAILON commands have been added at the end of the maintenance job and recovery program so no intervention is required to get HP Desk running again once the jobstreams have completed. As more people want to access HP Desk from home and while traveling, uptime becomes a more important issue.

Each time the maintenance job runs, administrative reports are produced that must be checked on a daily basis; doing this takes five to ten minutes per HP 3000. What needs to be done includes:

1. Checking the data base item count summary for any corrupt or patched items. Any items that could not be handled by the maintenance job are looked at by the technical support person.
2. Checking the In Tray of General Delivery to see if any new messages have arrived. If so, log on and forward them to the proper recipients. A jobstream has been written to print and then delete General Delivery messages, but it is preferable that the recipients receive them in electronic form.
3. Ensuring that data sets have adequate capacity. HP Desk uses Image, and it functions best when data sets are under approximately 70% of capacity.

The five to ten minutes per HP 3000 per day is time well invested to ensure that all messages get properly delivered and that the data base has adequate room for new items.

Security

The HP Desk data base has passwords in it allowing each HP 3000 to log on to others to deliver messages. It also contains all data pertaining to the local users on that HP 3000. Hence security is an especially important issue. Each Corporate Offices HP 3000 has a unique set of passwords for the account, user MGR, and groups MAILJOB and MAILDB.

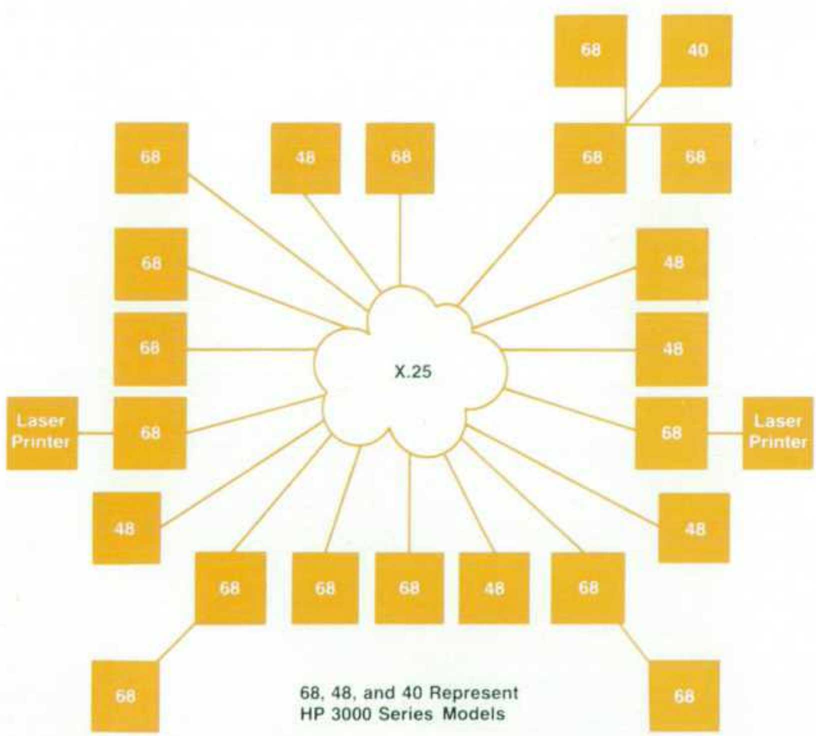

Fig. 11. *The present network at HP's Corporate Offices can connect every HP 3000 directly to every other HP 3000 in the local entity network. There are still two gateway HP 3000s to the companywide network.*

It also has lockwords on the configuration and utility programs. These passwords are changed regularly, and access to them is limited. HP Desk user passwords are another important element of security. Each user's password is kept encrypted in the data base. On a regular basis, a job is run that produces a list of all users who do not have passwords. The support team then follows up to make sure passwords are added and that users understand their role in protecting data and themselves.

One final area of operational concern is disaster recovery. In December 1984 the HP Desk team wrote a disaster recovery plan for HP Desk at the Corporate Offices. The plan assumes operational support and is integrated with the disaster recovery plan for the entity. It begins with a definition of an HP Desk disaster in terms of system downtime or backed-up messages. For each type of disaster (short-term, intermediate-term, or long-term), necessary and appropriate steps to recovery are included. There have been a few times when it has been necessary to activate the plan. Having specific steps listed with home and work phone numbers of the recovery team has saved much data and time. An important facet of the plan is user notification. When there is a disaster, a message is sent to each departmental contact person, who then disseminates the information.

Expanding the User Community

When time comes to expand the user base beyond the pilot users, there are several options. Everyone can be added at once (a support nightmare and not recommended). Expansion can address horizontal layers of the organization working from the top down or the bottom up. Expansion can also proceed through functional areas such as all of marketing and then all of administration. Or expansion can follow some combination of the above approaches. Corporate Offices followed a combination approach. Part of the pilot group had been top-level executives. Once they felt comfortable using HP Desk, implementation trickled down to the management levels below them. At the same time expansion was going from the top down, it also proceeded through one functional department at a time.

Full-scale expansion of the Corporate user community began in October 1982. After adding two groups beyond the pilot community, a memo was sent to all Corporate managers and supervisors describing HP's overall messaging strategy, HP Desk's most important features, and training options. A list of current HP Desk users was also included. The memo explained that the HP Desk implementation team would be contacting each group in turn, but that priorities were flexible to accommodate volunteer groups. As hoped, user demand drove additions. A few adventurous groups called immediately, and these people were added as registered users first. The objective was to begin by adding groups with the highest probability of success first. That included groups with the following desired (although not required) characteristics:

- Communication needs not being met with current alternatives
- High ratio of terminals to users (one to one is ideal)
- Some previous experience with other OA utilities
- Adequate printing resources
- Adequate disc space on the HP 3000
- Expressed interest in using HP Desk
- Managers and supervisors willing to use the product.

Before any group was added, its manager was asked to identify a departmental contact person. It was strongly recommended that the contact go through both the beginning and the advanced HP Desk classes before anyone in the group was added as a user. This was to help ensure a group expert and thus alleviate future support burdens. The group was evaluated and appropriate resource decisions made. Did users have enough terminals? Were they willing to share? Did they have access to printers? The group's training needs were also evaluated. Did users know how to log on to the HP 3000? Had they used any programs before? Did they need a customized class? Could they teach themselves with the help of the ITF? Another important piece of information was the list of other departments each group worked closely with. That was used to help guide implementation. It was important to add next the groups users needed to communicate with. Communication paths are as important during full-scale implementation as during the pilot phase.

Before the group was actually added to HP Desk, a

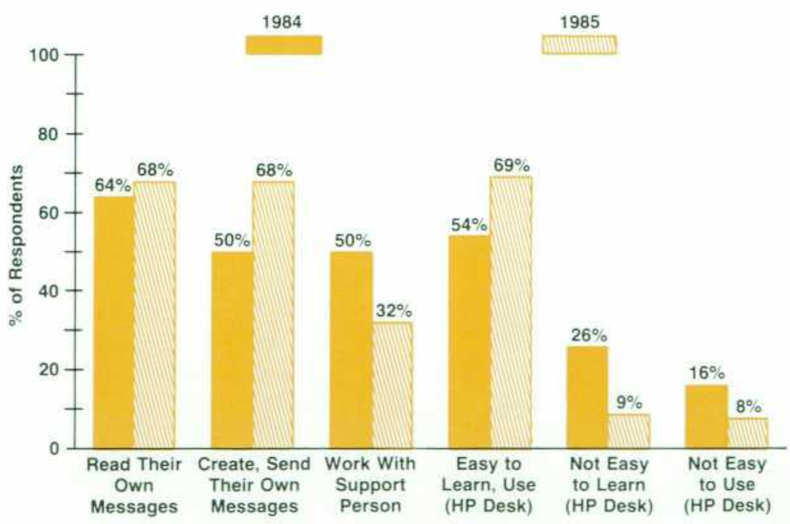

Fig. 12. *Changes in HP Desk use among HP Corporate Offices executives.*

member of the HP Desk team made a presentation to its members. The presentation covered HP Desk features, what each HP Desk tray is used for, examples of some of the screens, HP's overall messaging strategy, and the Corporate Offices implementation strategy. This allowed for answering any questions and introducing the support person. Once the group had been added, the HP Desk support team would send each member of the group a training message daily for 10 days.

Promoting the Use of HP Desk

Among the most important steps of the implementation plan is creating a follow-up plan. The real challenge in HP Desk implementation does not lie in adding more user groups but in getting existing users to sign on faithfully to read and reply to their messages. Using HP Desk means establishing new work habits. Groups should not be added and then ignored. After a specified period of time (usually four to eight weeks), it is important to check with the departmental contact and other group members to make sure they are using HP Desk, know the basic features, and can print messages to their favorite printer. Often, new users need help setting up user-defined commands (UDCs) to make the program easier to run. It is also necessary to let everyone know of newly added user groups.

From the beginning of implementation, the Corporate Offices' HP Desk team has followed up with each group. On a regular basis, information from the maintenance reports is logged. For each user, a check is made of disc space use, last sign-on date, and Calendar/Diary use. That information is shared with the departmental contact who can then pass it on to others. The follow-up is one way to measure use of the product.

The successful implementation of HP Desk at the Corporate Offices is exemplified by its widespread use among managers and executives. In April 1984, a survey was conducted among the Corporate Offices top-level managers and executives to assess such use. The survey was repeated one year later. Fig. 12 summarizes the results of the two surveys. Percentages in the graph are derived only from the questionnaires returned in each instance (slightly over half) and therefore are not strictly comparable.

HP Desk is also being used innovatively by several corporate departments. The accounting department distributes targeting and expense information to all corporate managers and supervisors. Formerly, stacks of reports were printed, collated, and distributed, costing several person-days of work per month. That process was automated with HP Desk. The reports now stay in their electronic form and are copied into HP Desk for mailing. Reports reach their destination faster than before and a number of person-days per month have been saved. While phasing in the new procedures, the accounting department worked closely with many managers, even giving them individualized HP Desk instruction. This project also motivated many groups to be added earlier than they might have been otherwise.

CCC has a special HP Desk user called CCC REGISTRAR. Anyone interested in the classes taught by CCC can send a message to the registrar for information concerning class content, scheduling, and fees. Telephone interruptions have dropped dramatically. There is still no electronic equivalent for a signature, so a paper form signed by the prospective student's supervisor is required before final enrollment.

The facilities site services department also has several special HP Desk users. Through one, it is possible to order refreshments for a class or meeting. The refreshment coordinator files requests in the Calendar/Diary and then prints the day's work orders from it. Paperwork has been virtually eliminated. To another special user, messages can be sent regarding everything from light bulb replacement to electrical wiring projects. Appropriate work orders are then created and scheduled.

In the San Francisco Bay Area, the treasury department

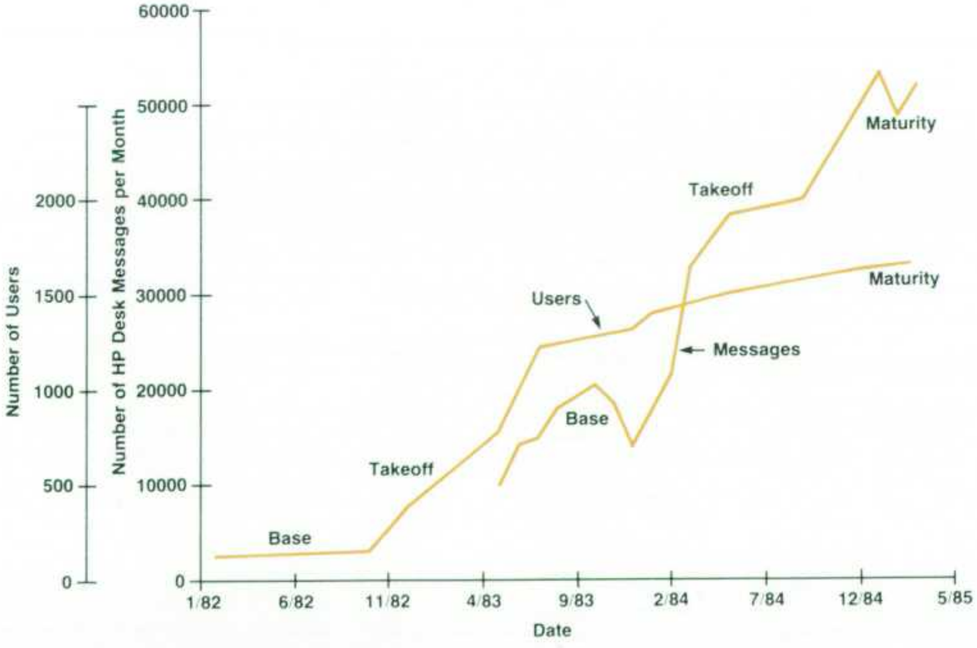

Fig. 13. *Growth in number of users and number of messages processed in the HP Corporate Offices HP Desk network. Growth occurs in three phases, labeled base, takeoff, and maturity.*

keeps a list of people whose checks have bounced at the petty cash department. Formerly, the list was typed monthly for each Bay Area division, then copied, stapled, and sent through interoffice mail. The information is now consolidated and distributed through HP Desk. Not only is it received much faster, but it is possible to track delivery through the use of acknowledgments. The information is now much more secure; only authorized individuals have access to an HP Desk mailbox. The department also collects worldwide investment data from foreign entities. Using jobstreams, files of outstanding investments are automatically sent to the Corporate Offices through HP Desk. Information is received in a timely manner and in a format ready for additional processing. Maintenance of these jobstreams is easy for the foreign entities because the users understand the HP Desk dialogue and can make any necessary modifications. Lastly, a member of HP's Geneva treasury department has developed a foreign currency exchange rate system using HP Desk and a spreadsheet software package on the HP 150 Personal Computer. The currency information originates at the Geneva treasury department and is mailed out to HP's offices in European countries and the Corporate Offices. It is automatically downloaded to an HP 150, where it updates a spreadsheet. This information can then be used to generate graphs automatically for the spot, accounting, or cross rates for European currencies. This eliminates the need for several country offices to track this information, manually produce charts, or reinput data.

Growth Phases

The growth in both the number of users and number of HP Desk messages processed proceeds through several distinct phases, as shown in Fig. 13. Fig. 13 goes through March 1985, when the user population began to stabilize at about 1700 users. Message volume has continued to increase since then, although not so sharply as before. In the first phase, a base of users and support experience is established. This is followed by a period of explosive growth. As new users are added and they become comfortable using HP Desk, the number of messages processed also grows, though this growth lags the addition of users by several months. Finally, there is a period of maturity. The addition of users slows, while there is continued and steady growth in message volume.

Myths

In closing, it would be useful to deny some HP Desk implementation myths. First, HP Desk is *not* a typical OA utility. It is not like HP Draw or HP Slate, utilities that can be installed and then forgotten until the next software upgrade. There are daily support responsibilities that must be addressed, and implementation is much more involved. Second, HP Desk should *not* necessarily be limited to one HP 3000 at each entity. HP Desk installation should be distributed among the HP 3000s that users are assigned to for all their other work. Third, HP Desk does *not* severely impact the performance of other programs. It takes job slots and disc space, but HP Desk users primarily impact their own response time, not that of other programs on the HP 3000. Fourth, all users should *not* be added at once. For the implementation team to handle the early support load, phased implementation is essential. This also ensures that the product is used appropriately. Fifth, the maintenance job should *not* be run weekly, but daily. A good backup for critical data is essential. Finally, secretaries should *not* necessarily read their manager's mail. It is feasible and desirable that managers manage their own messages.

Acknowledgments

The authors want to acknowledge the contribution of a number of people to the implementation of HP Desk at HP. First, Hank Taylor, Corporate Telecommunications and Office Systems manager, for his support and confidence in OUG's plans. Second, the Corporate Computing Center staff and their manager, Phil Wilson, for their help and cooperation in running the network. Third, Phil Way, for growing and managing the former Computer Group network. Fourth, the Pinewood HP Desk team, for making the changes necessary for successful large-scale implementation. And certainly not least, the local messaging coordinators, for their work in making the network run as smoothly as it does.

References

1. I.J. Fuller, "Electronic Mail for the Interactive Office," *Hewlett-Packard Journal*, Vol. 34, no. 2, February 1983.
2. D.H. Crocker, *Standard for the Format of ARPA Internet Text Messages, RFC #822*, Department of Electrical Engineering, University of Delaware, August 1982.

Hewlett-Packard Company, 3200 Hillview Avenue, Palo Alto, California 94304

Bulk Rate
U.S. Postage
Paid
Hewlett-Packard
Company

HEWLETT-PACKARD JOURNAL

September 1986 Volume 37 • Number 9

Technical Information from the Laboratories of Hewlett-Packard Company

Hewlett-Packard Company, 3200 Hillview Avenue
Palo Alto, California 94304 U.S.A.
Hewlett-Packard Central Mailing Department
P.O. Box 529, Startbaan 16
1180 AM Amstelveen, The Netherlands
Yokogawa-Hewlett-Packard Ltd., Suginami-Ku Tokyo 168 Japan
Hewlett-Packard (Canada) Ltd.
6877 Goreway Drive, Mississauga, Ontario L4V 1M8 Canada

CHANGE OF ADDRESS: To subscribe, change your address, or delete your name from our mailing list, send your request to Hewlett-Packard Journal, 3200 Hillview Avenue, Palo Alto, CA 94304 U.S.A. Include your old address label, if any. Allow 60 days.

5953-8552