


HEWLETT-PACKARD JOURNAL

APRIL 1977



Silicon-on-Sapphire Technology Produces High-Speed Single-Chip Processor

This new integrated-circuit processor is a static CMOS/SOS 16-bit parallel device. Its architecture is optimized for controller applications. Instruction execution times are 0.5 to 1.5 microseconds at an 8-MHz clock rate.

by Bert E. Forbes

SILICON-ON-SAPPHIRE technology offers a combination of low power consumption, high speed, high circuit density, and static operation unmatched by other integrated circuit technologies. Seeking system-level performance advantages, Hewlett-Packard has been working on this technology for several years.

The first integrated circuit to be fabricated using HP's silicon-on-sapphire (SOS) technology is called MC², for Micro-CPU Chip, a 16-bit central processing unit built with complementary metal-oxide-semiconductor (CMOS) logic. It is capable of performing a full 16-bit register-to-register addition in 875 nanoseconds and draws about 350 milliwatts of power. The actual size of this powerful CPU is . There are ten thousand transistors in the 34 mm² area.

In the coming months, articles will appear in these pages describing new products in which the MC² is an important component. This article is a description of this new circuit, its capabilities, and the system environment in which it is designed to operate.


The Micro-CPU Chip as a Controller

Unlike large computers, the micro-CPU is often used as a direct controller of other electronic and/or mechanical devices and as such is buried inside the object being controlled. Controller applications place heavy emphasis on data manipulation and decision making. Therefore, like any good system manager, the MC² has the ability to make decisions rapidly based on the data at hand. It also has good two-way communications with the rest of the organization. Because each job that an MC² handles will be different, it is designed to work easily with varied special-purpose hardware such as memories, direct memory access, slave processors, and existing input/output standards and interfaces.

Parallel Bus

Like most modern computers, a system based on the MC² is organized around a parallel bus structure (see Fig. 1). To increase the bus speed and decrease the number of external components necessary to build a bus interface, there is no multiplexing of data and address lines. All signals are valid for the duration of a bus access. Asynchronous memory or I/O operations require only 250 nanoseconds total system access time.

Because memory and I/O are basically different in



Cover: This month's cover subject is HP's new silicon-on-sapphire micro-CPU chip, MC², a 5.7×5.9-mm large-scale integrated circuit containing 10,000 transistors. A distinguishing feature of the SOS process is the translucency of the sapphire substrate, dramatized here with the help of a bundle of optical fibers, each 0.4 mm thick.

In this Issue:

Silicon-on-Sapphire Technology Produces High-Speed Single-Chip Processor, by Bert E. Forbes page 2

CMOS/SOS, by David Farrington, page 5.

Miniature Oscilloscope Probes for Measurements in Crowded Circuits, by Carolyn M. Finch, Marvin F. Estes, and Lawrence A. Gammill page 9

A Small, Solid-State Alphanumeric Display, by John T. Uebbing, Peter B. Ashkin, and Jack L. Hines page 13




Fig. 1. A system based on HP's new CMOS/SOS processor, the MC², is organized around a parallel bus structure. Data, address, and control buses are separate; multiplexing is not used.

speed, use, and access characteristics, the MC² and the bus have separate asynchronous handshake lines for the two functions. This extra versatility allows the designer to match the hardware being controlled to the most efficient access type and simultaneously makes it possible to simplify the decoding and timing networks. The MC² and the bus use 12-volt logic, which provides the high noise immunity often needed in controller applications.

A major feature of this system is the full visibility of what's happening: the address, data and control lines are all static and therefore easily connected to a logic analyzer like the HP 1600A. The entire system may run at any speed up to the 8-MHz maximum clock frequency, allowing the user to single clock the system and examine everything one cycle at a time, if desired. This is especially useful for debugging and troubleshooting the user's design. Static operation is feasible because CMOS/SOS simultaneously provides low power dissipation and high speed.

The control lines on the bus include memory and I/O handshakes, read-write, instruction fetch, run and idle, system clock (single phase), interrupt and interrupt acknowledge, and power-on-reset. All of these lines are MC² inputs or outputs. They are continuously available, a feature that eliminates decoding circuitry and latches.

Memory

The memory for an MC² system can be of any type: dynamic or static random-access memories (RAMs), read-only memories (ROMs), or programmable read-only memories (PROMs), of any access time compatible with the desired system speed. ROMs are used for program storage, fixed data, conversion tables, bootstrap loaders, and so on. Normally an MC² system uses fast-access ROMs, since they are accessed

at least once during every instruction. They should also be bus-organized to facilitate connection to the system with minimum overhead in both time and space.

RAMs are used for data storage, buffers, variables, programs, control/subroutine stack, and temporary storage. A minimum MC² system will have 256 words of random-access memory while the upper limit is a function of the job being performed (with a maximum limit of 64K words). HP has under development both ROMs and RAMs that exploit the high speed and low power of HP's CMOS/SOS integrated circuit process.

Input/Output

Since I/O circuitry is typically developed and optimized for each project, the MC² provides a flexible I/O interface capability. In general there are three categories of information handled in an interface: data (in and out), control (out), and status (in). Most interfaces have all three types, although simple ones may use implicit coding techniques to save hardware. We can visualize each of the types as one or more mailboxes where information is sent or picked up; what is done with the information is a function of the CPU program and the I/O hardware. The MC² uses this mail-




Fig. 2. The MC² is optimized for controller applications, and treats I/O devices as banks of registers. This simple paper tape reader interface takes advantage of the processor's static and asynchronous properties and direct I/O manipulation.

box concept in the form of external registers that may be read or written in the same manner and by the same instructions as the internal registers in the CPU. This allows tremendous flexibility in manipulating data while keeping the length of instructions to a minimum. External registers are logically grouped into banks of eight, reducing the number of address bits needed in an instruction to three. There may be up to 256 different banks; a simple device might have only one bank, while a complex device may use several banks of register addresses. This concept of treating I/O as banks of registers facilitates development of special large-scale integrated interface circuits.

Fig. 2 shows an example of a minimal paper tape interface. The "Do It" handshake of the MC² is used as the read command to move the tape; it enables the reader's three-state output buffer onto the low byte of the data bus. The feedhole signal from the tape is the asynchronous response handshake back to the MC² ("Done"). In this mode of operation the instruction that reads the input (MOVE, ADD, SUB, AND, etc.) takes one paper tape reader character time to execute (i.e., the MC² waits as long as necessary to read the external register). While this example would not be suitable for systems with high-speed real-time I/O constraints, it serves to show the ease of using a static asynchronous processor and direct I/O manipulation.

CPU

The heart of the system is the 16-bit CMOS/SOS microprocessor: the MC². This CPU's design emphasizes control and decision making to enable it to function quickly and effectively as a controller. It also has arithmetic capabilities equal to most low-end mini-computers (e.g., a register-to-register ADD or SUB takes less than one microsecond). When new features or more complex instructions (e.g., floating point arithmetic) are needed, the MC² has provision for easily enhancing its instruction set. The instruction set emphasizes condition code handling, subroutines, computed branches, bit testing, masking, priority encoding, and other processes that simplify the software needed to implement a controller. In addition, the 16-bit-wide data path manipulates addresses as easily and as fast as any other piece of data; this is especially important in maintaining speed of execution and ease of programming.

All of the 34 classes of instructions are 16 bits wide with a second 16 bits optionally used for a full word of immediate data or for full-range addresses (see Fig. 3). There is a broad register handling capability to facilitate both I/O and decision processes. A typical scenario is to obtain some information (from I/O), make a decision (an instruction) based on previous situations (data stored in registers and memory), and

cause some action (to I/O). In some simple cases this can be done with as little as one instruction in a program because the MC² instructions are designed to match that action sequence. As a slightly more complex example, consider the problem of interrupts: upon receiving one, the CPU must decide who interrupted, save the environment (or at least save STATUS), and go to a subroutine to initiate specific action. This takes the following sequence of instructions:

- Hardware saves the return address and disables interrupts
- PUSH STATUS onto stack (may also PUSH other registers if needed)
- MOVE EXTERNAL REGISTER TO INTPOLL REG (performs priority encoding)
- GO TO INTRROUTINE, INTPOLL REG (indexed branch through table).

In just three instructions the MC² has saved the status register, read a byte containing one bit per interrupting device, priority encoded the bits, and jumped through a branch table to the highest-priority routine. The controller-oriented architecture of the MC² keeps the total elapsed time, including the recognition of the interrupt, to at most 5 microseconds.

The MC² operates at clock speeds as fast as 125 ns per cycle and uses from 4 to 12 cycles per instruction (0.5 μ s to 1.5 μ s). These speeds are easily obtained using the SOS process. At full speed the MC² typically dissipates 350 mW, demonstrating the high speed and low power characteristics of the CMOS/SOS process. In fact, most of the power is dissipated in the custom logic arrays on the chip; these are




Fig. 3. MC² instruction formats. There are 34 classes of instructions. All are 16 bits wide. 16 additional bits for data or addresses are optional.

CMOS/SOS

by David Farrington

HP's complementary metal-oxide-semiconductor silicon-on-sapphire (CMOS/SOS) process has been developed as a large-scale integrated circuit technology. "Complementary" refers to the presence of both N-channel and P-channel MOS field-effect transistors on the same substrate.

The CMOS/SOS technology uses metal-gate transistors and requires both N^+ and P^+ diffusions for the sources and drains. The initial silicon material is N type and is converted to P type by ion-implantation in the gate regions where N-channel transistors are formed.

The major distinction between SOS and other integrated circuit technologies is that, instead of being formed in a wafer of bulk silicon, the circuit begins with a thin layer of silicon on a sapphire substrate. It is this feature that leads to the major performance advantages and the availability of some unusual circuit elements in this technology.


The accompanying diagram shows the structure of a complementary pair of transistors and an isolated diode in SOS. The thin layer of silicon (approx. $0.5 \mu\text{m}$ thick) consists of separate regions in which the isolated diode and transistors are formed. Electrical isolation is not by virtue of the usual PN junctions. Instead, the devices are separated completely by dielectrics—oxide and sapphire. For this reason, the isolation is extremely good. Moreover, the capacitance between adjacent structures is very small. The only significant capacitances are between the transistors' essential components—the sources, drains, and gates. The maximum attainable operating speed of

a circuit is determined by the ability of the available current to charge the capacitance associated with the circuit, and since that capacitance is minimized in the SOS technology, high-speed operation is possible. The logic implemented with this technology not only achieves high speed with low active power, but also consumes low standby power since the complementary circuits use significant power only during periods of actual switching.

The SOS process makes possible circuit designs demonstrating high speed and low power. For HP's purposes, the process must also be capable of producing large-scale integrated circuits with a high production yield. Large-scale integration is possible because the dielectrically isolated devices can be packed close together. And the process has been developed to the point where the incidence of defects is sufficiently low. The new micro-CPU chip, MC², described in the accompanying article has a chip size of 34 mm^2 and contains 10,000 transistors. It is produced with a yield of approximately 9%.

All the foregoing features are most attractive. The complementary static logic is easy to design and the resultant circuits operate at high speed with low power. The remaining important characteristic for a viable LSI technology is reliability of the components. This also is achieved. In 750,000 device hours we have observed a failure rate of less than 0.1% per 1000 hours at 65°C .

A range of circuits based on this technology will enhance future Hewlett-Packard products. The MC² is the first of these.






Fig. 4. MC² basic block diagram. The data path consists of four blocks of logic that may be connected to a common internal bus or operated independently. See Fig. 5.

implemented using NMOS logic instead of CMOS to save area at the expense of some power.

It will always be true that any CPU, no matter how carefully designed, can never be all things to all situations. With this in mind, the MC² is designed for maximum adaptability. The MC² incorporates a mechanism that allows a designer to easily add slave processors that can detect, execute, and change the flow of control of the MC² instruction stream. A signal (FETCH) is provided on the bus indicating that the data bus has an instruction; if the op code of the instruction is a hexadecimal F then the MC² goes idle for 250 ns to allow additional decoding time. If a slave processor (or front panel, or debug tool, or other device) decides that the remaining 12 bits of the instruction are one of its op codes, it forces the MC² to remain idle. The slave can then execute the instruction and has full read/write access to all of the registers in the MC² through two reserved banks of I/O addresses. This mechanism makes it simple to design a front




Fig. 6. Logic blocks in the data path are connected to the main bus by full complementary transmission gates. These gates are fast (<1.5 ns delay), small (540 μm²), and bidirectional, thereby simplifying complex logic implementation.

panel for such a system. This facility was also of great use in debugging the original MC² design.

MC² Implementation

Fig. 4 shows the basic diagram of the MC² and the connection to the data, address and control buses that have already been described. The data path (Fig. 5) consists of four blocks of logic that may be connected together on a common internal 16-bit bus or may be operated independently to provide faster instruction execution. The manipulation of data and the bus linkages is determined by the control section.

The logic in the data path is designed to maximize the speed of the MC² while minimizing the chip area (the die is 5.7 mm × 5.9 mm). The registers are latches in a pseudo-RAM configuration; each block is independently readable and writable. The blocks are connected to the main bus by full complementary transmission gates (Fig. 6). These gates are flexible design components because of the low capacitance and small area characteristics of CMOS/SOS.




Fig. 5. MC² data path block diagram. External registers used for I/O operations are treated in the same way as the internal registers. This allows flexibility in manipulating data while keeping the length of instructions to a minimum.




Fig. 7. The ALU uses functional gating to save area and generate a 16-bit carry in 35 ns; only half the number of transistors are needed as for a conventional gate realization. Four-bit look-ahead is used in generating the 16-bit carry. There are four carry look-ahead circuits like the one shown here. The ripple method is used within each four-bit block; the combination of methods also helps save area.

This type of gate is also used extensively in the field extractor/shifter that preprocesses the data for the arithmetic logic unit. The ALU uses transmission gates and functional gating to save area and achieve a 16-bit carry in 35 ns (Fig. 7). Functional gating implements a Boolean expression directly in transistors instead of using NAND and NOR gates. The carry path consists of four carry/look-ahead functional gates with 18 transistors each (conventional gating would take 36 each). The carry ripples within a four-bit nibble to save area.

The ALU is capable of a wide variety of operations.

One specific operation is chosen by a hardware table for each instruction (e.g., the ADD instruction forces an ADD, and so on); this minimizes the control logic. Fig. 8 shows the control section, which consists mainly of two CLAs, or custom logic arrays. These are like PLAs, programmable logic arrays, but do not have a full output matrix. The CLAs are controlled by a five-bit state counter. The CLAs have very short delay times thanks to the low-capacitance SOS structure.

The internal flow of control is a function of the current state, the current instruction, the data path status




Fig. 8. MC^2 control section block diagram. Custom logic arrays (CLAs) are like programmable logic arrays (PLAs, a form of read-only memory) but do not have a full output matrix. The CLAs are controlled by a five-bit counter.




Fig. 9. The MC² is packaged on a 48-pad ceramic substrate. A special low-insertion-force socket is used.

indicators, and the external control lines. The control for the data path is from the second CLA, whose outputs are synchronized by a set of D-type flip-flops. This CLA's outputs are a function of the current state, the data path's status, and some external control lines. (The knowledgeable reader will note that this design is a combination of Mealy and Moore machines).¹

Physical Description

The MC² is packaged on a 48-pad leadless square ceramic substrate that measures approximately 30 mm on a side with a 2.54-mm pad-to-pad spacing (Fig. 9). This carrier fits into a special low-insertion-force socket, providing ease of exchange and aiding automatic board testing situations that require that the processor be removed. The package provides excellent heat dissipation characteristics; chip tempera-




Fig. 10. Photomicrograph of the MC².

ture rises only 10°C at the 350 mW power level.

Fig. 10 is a photomicrograph of the MC².

Acknowledgments

No IC is designed without the help of many people and the support of an organization like HP. My thanks to all of them. Two specific engineers did the majority of the design work and deserve special mention and thanks: Lam-Giang Dang and Richard Woolley. Peter Ashkin, John Figueroa, Pat Mulraney, Matt O'Brien, Ken Rothmuller, Jerry Stolle, and Randy Yee also contributed directly to the engineering design. 

Reference

1. Z. Kohavi, "Switching and Finite Automatic Theory," McGraw-Hill Computer Science Series, 1970, pp. 280-313.

Bert E. Forbes



MC² project manager Bert Forbes joined HP ten years ago as a computer design engineer. He was a member of the original design team for the HP 3000 Computer System, serving as CPU project manager. Later he spent a year and a half in Europe to help launch the HP 3000 there. Having completed the MC² project, Bert is working part time on the HP 3000 again, besides spending many hours starting his own small microprocessor interface company. Bert received his BSEE degree in 1966 from Massachusetts Institute of Technology and his MSEE in 1967 from Stanford University. A Texan by birth, he grew up in New Mexico and now lives in Cupertino, California. He's married, serves as counselor for a church youth group, and is "on the way to becoming a wine connoisseur and a father."

degree in 1966 from Massachusetts Institute of Technology and his MSEE in 1967 from Stanford University. A Texan by birth, he grew up in New Mexico and now lives in Cupertino, California. He's married, serves as counselor for a church youth group, and is "on the way to becoming a wine connoisseur and a father."

David Farrington



David Farrington was born in Stockport, England, and received his BSc degree in physics from Nottingham University in 1958. During the early years of his career he did research on avalanche multiplication solid state devices. In 1967 he received his MSEE degree from Stanford University. With HP since 1969, he's done development work on various IC technologies and devices and helped develop HP's CMOS/SOS process. He's a member of IEEE. David has two children and lives

in Sunnyvale, California. An avid cricket player in northern California, he spent three weeks last year playing cricket in South America, and looks forward to going back some day. He also enjoys badminton and sailing.

Miniature Oscilloscope Probes for Measurements in Crowded Circuits

Resistive-divider probes only 2.4 mm (0.1) inch in diameter can access test points in densely populated circuits without shorting to adjacent leads. Grounding options preserve fast rise times.

by Carolyn M. Finch, Marvin F. Estes, and Lawrence A. Gammill

RESTRICTED ACCESS FOR TESTING is one of the annoying consequences of the ever-increasing complexity of electronic devices. Circuit integration, component miniaturization, new printed-circuit production techniques—all have contributed to denser packages and more closely spaced wiring systems. As the packages become more complex, it becomes more difficult to access them for test purposes, especially with the presently available probes and accessories that were originally designed for larger, simpler devices.

In response to this situation, Hewlett-Packard has developed a new family of miniature, high-impedance oscilloscope probes. The basic probe is a small (2.4 mm diameter, 25 mm long) cylinder with a needlelike tip (Fig. 1). It is used in conjunction with a



Fig. 1. New miniature oscilloscope probe (right) is shown here alongside a conventional probe (left) of recent design. The 2.4-mm (0.1-inch) diameter of the new probe permits probing where device leads are on 0.1-inch centers.



Fig. 2. New probe, shown at top with its insulating sleeve removed, is supplied with a spring-type ground lead, a slip-on hook tip, and a flexible ground lead with an alligator clip and it will also be supplied soon with a miniature pincer-type "grabber" (not shown). An IC clip and a probe-tip-to-BNC adapter are also available.




Fig. 3. The insulating tip retracts far enough to allow the spring ground tip to establish a ground-reference point at the end of the barrel for measurements on high-speed signals.

variety of insulating/interfacing accessories (Fig. 2) to meet a variety of testing situations. The narrow body lends itself to use in congested areas and, when combined with the proper fittings, enhances the user's ability to access difficult measurement points with little danger of accidental shorts to adjacent leads.

Despite their small size, these broadband, compensated, voltage-divider probes can be used with any oscilloscope or other instrument that has a BNC input connector and the correct input impedance. The first probe of the series (Model 10017A) has a 10:1 signal-division ratio and an input impedance of 1 M Ω shunted by less than 8 pF when used with instruments that have a 1-M Ω input shunted by 9-14 pF. Probes with differing length cables (1 to 3 m), division ratio (1:1), scope-compensating capability (to 30 pF), and input impedance (50 Ω) are also available.

Conventional Probing

An insulating sleeve added to the bare probe approximates a miniature version of the traditional oscilloscope probe (Fig. 1). In this configuration, the probe looks and handles like a 1/2-scale version of the traditional oscilloscope probe with one exception: the forward barrel insulator is retractable, making the traditional slip-on cover for the high-speed ground unnecessary. With the barrel insulator retracted and the ground spring in place, the probe is configured with a very short ground lead for high-frequency, point-to-point probing (Fig. 3).

With the barrel insulator in the forward position, the probe may be used with the 20-cm flexible ground




Fig. 4. With the slip-on hook tip and flexible ground lead in place, the miniature probe can be used like a conventional pincer probe for attachment to test points or component leads.

lead for point probing where this type of grounding gives adequate response fidelity. The pointed tip makes positive metallic contact to narrow conductors and penetrates commonly-used protective coatings, and the extended insulating sleeve prevents shorts to closely-spaced adjacent leads.

With the barrel insulator retracted and the flexible ground lead in place, the probe may be used with the slip-on hook tip or soon-to-be-available pincer tip for attaching to various component leads (Fig. 4).




Fig. 5. Miniature probe with insulating sleeve removed is held in place on an IC lead by the optional IC clip. The bare ground-reference pin, shown here in the right-hand corner of the clip, can be inserted at any lead position. It grounds reference planes within the clip that in turn contact the barrel of the probe(s). Risetimes as short as 1.3 ns are preserved by this arrangement (Fig. 6).




Fig. 6. Fast step as viewed on a 275-MHz oscilloscope using a type 10017A miniature 10:1 probe. Signal, oscilloscope, and probe are identical in all three photos; the differences are the result of the grounding attachment hardware: (a) using the 10024A IC clip (Fig. 5) with integral grounding; (b) using the spring-tip ground (Fig. 3); (c) using a conventional IC clip, hook tip, and 20-cm ground lead (Fig. 4). Sweep time is 10 ns/cm.

DIP Probing

In addition to the above, the probe with its insulating sleeve removed may be used with an accessory clip (10024A) for attaching directly to the pins of 14- and 16-pin DIP packages without danger of shorting adjacent pins.

In this application, the clip is installed on the DIP package, the furnished ground-reference pin is inserted into the appropriate position, and one or more probes are inserted to mate with the desired package leads (Fig. 5). The ground-reference pin connects reference planes in the clip to the package ground to provide a ground reference for any probe inserted in the clip. This grounding arrangement is extremely effective; high-speed pulse fidelity achieves a level previously associated only with probe-to-BNC adapters and high-frequency, point-to-point probing (Fig. 6).

Pins with body coatings that insulate them from the ground plane are also available to enable other types of probes or contact arrangements to be used for coupling signals into or out of an IC while the clip is in place.

Probe Body Design

A circuit diagram typical of the new probes is shown in Fig. 7. While the circuit shows no departures from established probe design, virtually every component and assembly procedure had to be developed "from scratch" because of the very small size.

To minimize input capacitance, a thick-film resistor and "dot-on-the-wire" capacitor are placed right at the tip of the 10:1 divider probes. These are encapsulated in the 0.1-inch diameter barrel. Two major difficulties were encountered in this area: ruggedness and insulation.

The ruggedness problem was attacked by using an assembly sequence that builds the probes in concentric layers around the electronic components, which are held in the proper relationships through the assembly process by a traveling holder arrangement. The necessary insulation thus becomes one of the layers added to the assembly, and the stainless-steel barrel becomes the outside layer and ultimate support of the completed device.

The cable is substantially smaller in diameter than any other presently produced for similar applications. This reduction in size would, traditionally, translate to a reduction in strength of the individual components of the cable; the center conductor, for example, would become so fragile in its unsupported state that it could not be processed unless a new

Probe Cable Design

The cable is substantially smaller in diameter than any other presently produced for similar applications. This reduction in size would, traditionally, translate to a reduction in strength of the individual components of the cable; the center conductor, for example, would become so fragile in its unsupported state that it could not be processed unless a new




Fig. 7. Circuit diagram of the Model 10017A 10:1 divider probe. The compensation assembly at the end of the cable attaches directly to the oscilloscope vertical input connector.

fine-wire extrusion facility were developed for our facility. Two advances over HP's present technology were incorporated into the new miniature cable to alleviate this situation, providing one of the more durable cables presently produced, size notwithstanding.

The woven braid of the cable is its primary strength-providing member. Going from copper wire to a higher-tensile strength, proprietary wire gives tensile strengths in the small cable approaching those of traditional, full-size probe cables. Secondly, the functional lifetime of the fragile center conductor has been extended by lubricating the polyethylene insulator-to-conductor contact so that it does not bind and kink with changes in axial translation.

To summarize, the new miniature probe system makes contributions in the following areas: (1) ability to attach directly to dual-in-line packages for high-speed measurements; (2) increased ability to access densely-packed miniature components; and (3) decreased capacitance loading of the circuit under test.

Acknowledgments

The authors would like to thank: Jim Freeman, Pete Rawson, and Jim Williams for contributions to the basic concept during the product definition phase; Jim Carner, Rick James, and Bill G. Smith for accessories design; Carl Glitzke, Mike Keegan, and John Crowninshield for process development and produc-

tion start up; and Blair Harrison, Walt Fischer, and Stan Lang for continued support. 

Marvin F. Estes



Marv Estes joined Hewlett-Packard in 1973 after spending two years at NASA, Huntsville, Alabama, working on magnetic levitation, electrets, and a magnetic hammer. He has a BSEE degree (1965) from Purdue University (Lafayette, Indiana), an MSEE degree (1967) from Clarkson College of Technology (Potsdam, New York) and has done post-MS work at North Carolina State and Colorado State Universities. Off

hours, Marv likes to go camping with his family (wife, 3 sons, and 1 daughter) and he's working towards completion of his 30-foot-long greenhouse from which he hopes to get fresh vegetables year around.

Lawrence A. Gammill



Before joining HP in 1967, Larry Gammill spent 10 years in the aerospace industry, designing models for wind-tunnel testing and contributing to the Apollo program and the Titan missiles. Since being at HP, he has done product design for a number of 1400- and 1800-series oscilloscope plug-ins and 1700-series oscilloscopes, obtaining two patents along the way. He obtained his basic engineering education at Mesa College (Colorado) and at the University of Colorado at Colorado Springs.

Married, and with four children, Larry works with church youth groups and he also likes to hunt with a muzzle-loading rifle he built from scratch himself.

Carolyn M. Finch



A 1968 graduate (BSME) of the University of Houston, Texas, Carolyn Finch worked three years in boiler and pressure-vessel design and in pneumatic conveyor design before switching to homemaking for two years. She then joined HP in 1975 to work on the design of oscilloscope accessories. Outside of working hours, Carolyn is helping her husband design a stained-glass production line for his antique shop, besides devoting time to their family (three girls, 7, 4, and 2)

SPECIFICATIONS Miniature Oscilloscope Probes

Probe Model No.	Approx. Overall Length In Metres (ft)	Division Ratio	Input R	Shunt C	Compensates Scope Input C
10017A	1 m (3.3)	10:1	1 M Ω	8 pF	9 to 14 pF
10018A	2 m (6.6)	10:1	1 M Ω	10 pF	9 to 14 pF
10040A	1 m (3.3)	10:1	1 M Ω	9 pF	20 to 30 pF
10041A	2 m (6.6)	10:1	1 M Ω	12 pF	20 to 30 pF
10042A	3 m (9.8)	10:1	1 M Ω	15 pF	20 to 30 pF
10021A	1 m (3.3)	1:1		36 pF	
10022A	2 m (6.6)	1:1		62 pF	
10026A	1 m (3.3)	1:1	50 Ω		
10027A	2 m (6.6)	1:1	50 Ω		

ACCESSORIES SUPPLIED: One each retractable hook tip, pincer tip, alligator clip, 20-cm (8 in) ground lead, and grounding spring; four indicator sleeves (A,B,C, and D).

ACCESSORIES AVAILABLE: 10024A IC Test Clip; BNC-to-Probe adapter.
PRICES IN U.S.A.: 10017A/18A/40A/41A/42A, \$90. 10021A/22A/26A/27A, \$45. 10024 IC Test Clip, \$15. BNC-to-Probe Adapter (HP P/N 1250-1454), \$8.25.

MANUFACTURING DIVISION: COLORADO SPRINGS DIVISION
1900 Garden of the Gods Road
Colorado Springs, Colorado 80901 U.S.A.

A Small, Solid-State Alphanumeric Display

Serial loading and on-board storage of data on this dot-matrix display reduces the cost and complexity of supporting circuitry substantially. The 5 × 7 array of LEDs gives full alphanumeric capability (upper and lower case letters, numerals, special symbols).

by John T. Uebbing, Peter B. Ashkin, and Jack L. Hines

UNTIL VERY RECENTLY, full alphanumeric LED displays have not been used nearly as widely as the popular numeric-only displays. This is not only because of the cost of available alphanumeric devices, but also because of the complexity and expense of the associated addressing and driving circuitry.

A new, small, rugged, alphanumeric LED display, the Hewlett-Packard type HDSP-2000, now reduces the cost of using alphanumerics by minimizing needed supporting circuitry. Coupled with the recent availability of low-cost microprocessors and semiconductor memories, this should contribute to a rapid expansion in the use of displays with full alphanumeric capability (upper and lower case alphabet, numerals, and special symbols) where previously they have not been practical. Some of the potential applications for low-cost alphanumeric display systems are in calculators, portable computer terminals, data-entry stations, and other readout devices, all of which are becoming more common in business, medical, and telecommunications systems.

The new display is shown in Fig. 1. Each unit is a four-character LED cluster with each character consisting of a 5×7 dot matrix. Characters are 3.8 mm (0.15 inch) tall spaced on 4.5 mm (0.175 inch) centers. Each cluster is in a 12-pin, dual in-line package 17.7 mm (0.7 inch) long that can be end-stacked to provide displays with as many characters as desired while maintaining constant digit-to-digit spacing.

The new display requires much less complex drive circuitry than earlier units because data is entered serially and stored on board. Traditionally, LED dot-matrix displays have been organized in an X-Y addressable array requiring as many as 12 interconnect pins per digit (with strobed operation). The HP type 5082-7101 four-digit cluster,¹ for example, has 20 pins for driving the columns (five per character) and seven pins for driving the seven rows (similar rows of all characters are tied to a common pin).

During strobed operation, the row pins are grounded one at a time while current is supplied to the appropriate column pins, illuminating the diodes needed for that row in all characters. By strobing all seven rows 100 times per second, the characters appear fully illuminated with no flicker.

Organization of the New Display

In the new HDSP-2000, the 5×7 diode array for each character is organized with the anodes of the seven diodes in each column tied in common and the cathodes of the five diodes in each row tied in common (see Fig. 2). Like columns in the four characters of a cluster are tied together and brought to a single pin (e.g., column 1 of all four characters are connected to pin 1). The five cathodes in each row of an individual character are tied to an on-board current sink




Fig. 1. New, small, dot-matrix display has four LED alphanumeric characters in a 12-pin package having standard 0.1-inch lead spacing. With a significant portion of the operating electronics on-board the display unit itself, the display is easier to integrate into a system than earlier dot-matrix displays.




Fig. 2. The organization of the HDSP-2000 Display is shown by the block diagram at top. Circuit details are shown in the diagram below.

driven by one stage of an on-board shift register. Thus, any diode in the cluster is addressed by shifting

data to the appropriate shift register location and applying a voltage to the appropriate column pin.

An abbreviated diagram of the circuitry internal to the HDSP-2000 is also shown in Fig. 2. The serial organization of the shift registers allows long strings of row data to be entered through a single interconnect pin. Hence, only six pins are needed to drive the four characters (the other pins are for the shift register supply voltage and clock, the brightness control, a data output, and ground). The TTL-compatible data output can feed the data shifted out of the last shift-register stage directly to the data input of the succeeding cluster. Hence, all the data for an extended display can be supplied through a single TTL-compatible input. Examples of control circuitry needed for driving the display are described in the box on the next page.

Categorization of the LED brightness enables the user to stack any quantity of clusters of the same brightness category and achieve uniform brightness in a multicluster display. The LEDs are pre-tested for brightness in wafer form and classified in eight overlapping categories so each cluster can be made up of LEDs having the same brightness.

Display Development

Several design approaches were considered for the LED array. The first approach involved a monolithic (single chip) 35-LED matrix fabricated in GaAsP. Prototype chips were made but after some experience had been gained, an analysis of the yield indicated that a hybrid array would be more cost effective.

Because of the small size of the character and the close dot spacing, the use of individual LED chips was not feasible. Instead, all five diodes of each row are on a single, narrow chip or "sliver" with diode center-to-center spacing of 0.46 mm (some development was required to insure that the wafers could be

scribed and broken into the relatively long, narrow slivers with good yields). The seven slivers for each character are then mounted on the ceramic substrate on 0.56-mm centers, as shown in Fig. 3.

Package Development

The initial package design goals were: (1) DIP geometry; (2) end stackability; and (3) low cost. Because of the large number of wire bonds, and the need for good reliability during temperature cycling, an air-gap construction was the logical choice.

An early design used a ceramic substrate with brazed-on copper leads and a polycarbonate cover that slipped over the substrate. The substrate itself is an alumina ceramic which, with the copper leads, provides a low thermal-resistance path for enhanced power dissipation capability. The metallization on the substrate was carefully configured to provide minimum electrical resistance for the relatively high currents that flow in the ground and column lines.

The polycarbonate cover would have allowed magnifying lenses to be molded in economically. It soon became apparent, however, that magnification was not needed with the sliver array, so the design was changed to use a glass window to gain the advantages of quasi-hermetic sealing. The ceramic substrate was redesigned to accommodate a metal sealing ring with an epoxy sealant.

The glass window is a red, optical quality glass that has 93% transmission at a wavelength of 655 nm, about the same as clear glass but with the contrast enhancement that comes from blocking other wavelengths. The primary problem with the glass was the difference in coefficient of expansion between the glass and the ceramic substrate. The glass-filled epoxy sealant that was chosen is able to accommodate this difference—units have survived 100 temperature cycles between -55 and $+100^{\circ}\text{C}$ and have passed a 5×10^{-7} atm-cc/s helium leak test and the MIL-STD-883 Method 106 moisture resistance tests.

Integrated-Circuit Development

Placing the shift registers and current sinks on the display substrate was a major factor in achieving simpler support circuitry. It was quite obvious that to provide a compact assembly, the shift registers and current sinks would have to be on IC chips fitted between the LED arrays. This led to the choice of a column strobing technique using a 7-bit shift register for the rows in each character. Shift registers and associated current sinks for two adjacent characters are on each IC chip.

For the IC chip itself, design goals were: (1) good assembly yield in a complex hybrid; (2) chip width less than 1.3 mm to permit fitting between the LED arrays; (3) ability to sink a constant LED current irres-




Fig. 3. Close-up view of the display substrate with the five-diode "slivers" mounted in place.

Generating Characters on a Dot-Matrix Display

Dot-matrix alphanumeric displays usually require any character to be generated as a combination of several subsets of data. The character is displayed by presenting the data subsets sequentially to the appropriate locations on the display matrix, a technique known as strobing. If this is repeated at a fast enough rate, the entire character appears as a continuous image.

For the 5x7 dot matrix of the HDSP-2000, each character is made up of five 7-bit subsets. For a four-character display, the 28 bits representing the 7-bit subsets for the first columns in all four characters are loaded serially into the on-board shift registers and the first columns of all the characters are energized for a short period, usually less than two milliseconds. Then the 28 bits of the next subsets are loaded and the number two columns are energized. The process continues for all five columns and then repeats.

If the repetition rate is 100 Hz or greater, the display appears to be lit continuously with no flicker. For most applications, a duty factor [(diode "on" time)/(refresh period)] of 10% provides more than satisfactory display intensity. The duty factor (DF) can be determined by the formula

$$DF = \frac{T}{N_s(T + t)}$$

where T is the diode "on" time, t is the time required to load each 28-bit subset into the shift registers, and N_s is the number of data subsets, in this case 5. N_s(T + t) is the refresh period and is 10 ms for a 100-Hz refresh rate. A 100-character display, for example, using a shift-register clock rate of 1 MHz can be operated with a duty factor of 13%.

Microprocessor Control

When the display is to be used with a microprocessor-based system, the microprocessor can decode the data and load it into the display's shift registers. There is a trade-off, however, between the amount of time the microprocessor spends servicing the display and the complexity of the interface hardware. Because of the on-board storage of data on the display, the microprocessor can do its regular tasks during the "on" time of the display. Loading column data can be on a priority basis so if a more urgent task calls the microprocessor, the display may flicker, but not very much. It has been estimated that a microprocessor like the MC6800 requires less than 15% of its time to service a 20-character display using a minimum-hardware approach. Circuit details for a typical microprocessor-driven display system are shown in Fig. 2.




Fig. 1.




Fig. 2.

Hardware Control

When a microprocessor is not used, a character generator and a refresh memory are needed to operate the display. A block diagram of such a display-driver system is shown in Fig. 1. Data representing the characters to be displayed, usually supplied in one of the standard data codes such as ASCII or EBCDIC, are loaded into the refresh memory. The characters are read into the character generator under control of the 1/N counter, which addresses the stored characters one by one.

The column-select counter determines which column subset of a character is to be displayed. The column data is loaded in parallel into the shift register of the parallel-to-serial converter and then transferred serially into the HDSP-2000 display.

The operation of this system for a 4-character display illustrates the process. For this system, the 1/N counter is a 1/4 counter. With the four characters to be displayed loaded into the refresh memory, the 1/N counter selects the rightmost character first and loads it into the character generator. The column-select counter indicates column 1 and the column 1 data is supplied to the parallel-to-serial counter. It is then loaded into the display's shift registers while the ÷7 counter counts down the seven clock pulses needed for loading the data. The ÷7 counter then increments the 1/N counter, which loads the next character into the character generator, and column 1 data for that character is loaded into the display.

This sequence repeats for the remaining two characters. On the next clock pulse, the 1/N counter overflows, starting the display-time counter which counts down clock pulses for the selected display time. The overflow also inhibits clock pulses from stepping the shift registers in the display and enables the

output of the column-select decoder, which turns on the drive transistor for column 1.

When the display-time counter overflows, it resets the ÷7 counter, disables the column decoder, turning off the diodes, and increments the column-select counter. Data for the second columns in all four characters is then loaded and displayed in the same manner. This process then repeats for the remaining columns following which the column-select counter resets and the entire sequence repeats.

Easy Expansion

To extend this system to strings of more than four characters, the size of the refresh memory is increased accordingly and the modulus of the 1/N counter is made equal to the number of digits. Regardless of the number of digits, only five column-drive transistors are needed because like columns of all the characters in a display string are enabled simultaneously.

The column drive voltage can be anywhere between three and six volts but because of the on-board constant-current sinks, light intensity does not change noticeably if the column drive voltage varies. A simple, unregulated, full-wave-rectified power supply therefore may be used.

Power dissipation, typically 900 mW per cluster, is relatively high for such a small package so thermal design must be considered carefully. For short display strings (4 to 8 characters) maximizing the amount of metal surface on the supporting printed circuit board is usually sufficient. Longer displays usually require a heat sink (full details are given in HP Application Note number 966, "Applications of the HP HDSP-2000 Alphanumeric Display").



Fig. 5. Model 9825A Calculator was the first product to use the new HDSP-2000 display. The calculator's display is controlled by a dedicated custom LSI chip shared with the calculator's printer. The calculator's processor feeds the characters to a ROM in the custom LSI chip, and the chip's circuits then process the character data for display. The 32-character display has a refresh rate of 128/s and a duty factor of 17%.

pective of changes in total chip power dissipation and chip temperature; and (4) capable of being fabricated on a conventional bipolar IC line already making analog circuits with thick epitaxial silicon and a long minority-carrier lifetime.

Each shift register stage is an edge-triggered flip-flop that is extremely tolerant of harsh operating environments and of shifts in the processing parameters. Noise margins are sufficient to prevent problems with high junction temperatures and high ground noise. Circuit details are described in the box on the next page.

A logical "1" loaded into a shift register stage turns




Fig. 4. Current sink uses low-power transistor pair (T1-T2) to control 10× greater current in T3.

on its associated current sink so the diode at the intersection of that row line and the active column will turn on. Besides permitting the use of a poorly regulated supply to drive the columns, the constant-current sinks prevent brightness from being visibly affected by variations in series resistance of the LEDs and substrate traces.

A constant-current circuit often used in analog designs is used here. It works on the principle that two identical transistors have identical collector currents when the base-emitter voltages are the same. A corollary is that if one transistor has an emitter K times as large as the other, then the ratio of collector currents is $K:1$ when the base-emitter voltages are equal.

The circuit used in the HDSP-2000 display is shown in Fig. 4. The emitter of transistor T3 is 10 times as large as the emitter of transistor T2 so the current in T3 can be controlled by a current one-tenth as large in T2. A logical "1" in the corresponding shift register stage turns the current source "on" by saturating transistor T1, allowing current to flow. Resistor R is selected such that the T1 emitter current $I_e = (V_B - V_{be} - V_{sat})/R = 1.3$ mA with the brightness control voltage V_B at ≥ 1 volts (neglecting base currents since β is large). Hence the collector current I_c in T2 is also 1.3 mA and since T2 and T3 have the same base-emitter voltage, the collector current in T3 is 13 mA.

A logical "0" in the shift register cuts off T1 and only leakage currents then flow.

The IC also has buffers that allow the CLOCK, DATA IN, and V_B inputs to be driven by low-power TTL signals so the display won't load down system

A Highly Tolerant Shift-Register Flip-Flop

The on-board IC shift registers designed into the HDSP-2000 dot-matrix display are insensitive to variations in supply voltage, noise, and other interference. Each stage in the shift register is based on a widely used binary-counter circuit* that was modified to function as an edge-triggered, D-flip-flop, shift-register stage by the addition of two decoupling diodes, shown as D1 and D2 on the diagram below.


Operation is as follows. T1 and T2 form the master section of the flip-flop and T3 and T4 the slave. When the clock input is at 0 volts, T1 and T2 are off and draw no power while T3 and T4 hold the state of the stage. If the state were a logical "1", T4 would be saturated with its output, \bar{Q}_N , at $V_{sat} = 0.2V$ and T3 cut off with its output, Q_N , at $[(V_{CC} - V_{be})R7/(R5+R7)] + V_{be} = 2V$.

Shifting occurs when the clock line rises towards V_{CC} . Suppose the previous stage (N-1) contains a logical "0" (i.e., $\bar{Q}_{N-1} = 2V$ and $Q_{N-1} = 0.2V$). The T1-T2 master of stage N senses this condition through resistors R3 and R4 as the clock rises.

Since Q_{N-1} is at 0.2V, the emitter of T2 is at 0.2V and diode D2 is reverse biased, effectively decoupling this side of the master from the slave. \bar{Q}_{N-1} being at 2V forward biases D1 through R3. With T1 and T2 off at the time the clock starts to rise, the base and collector voltages of these transistors also rise. When the clock reaches $V_{be} + V_{sat}$, T2 starts to turn on and regenerative action keeps T1 turned off while T2 saturates. The relatively large emitter current in T2 flows into the base of T3, causing T3 to saturate with the Q_N output at logical "0" and the T3-T4 bistable action switches \bar{Q}_N to a logical "1".

Because the T2 emitter current is much greater than any current through R3 or R4, once the T1-T2 master goes bistable, subsequent changes in the output of the previous stage are locked out until the clock makes its next positive-going transition. Hence, data is transferred reliably from stage to stage on the rising edge of each clock pulse.

*A. Richardson and A. Foss, "New Binary Counter Circuit," *Electronic Letters*, December 1965.



buses. The buffers are fairly standard designs modified to minimize the input bias current, I_{IL} , so several clusters can be driven in parallel.

Modifications were made to the original IC chip design as an understanding of the overall display and packaging system evolved. The first version had "industry standard" bonding pads, an unbuffered analog brightness control circuit for supplying V_B , and buffered reset circuits to reset all 14 shift-register stages to logical "0". After some experience it was determined that no one used the reset feature, the unbuffered brightness control presented an unrealistically heavy load on the user's brightness control circuit, and the small size of the bonding pads hindered good yields in hybrid assembly.

Version two had an op-amp that buffered the brightness control circuit, no reset circuits, and larger pads. The third and final version replaced the analog brightness control circuit with an on-off digital control circuit so brightness could be controlled easily by pulse-width modulation techniques.

Final Assembly

Stringent specifications for the substrate assure close tolerances between the package lead positions and the topside stripes where the slivers are attached. Close tolerances are required not only for character-to-character alignment but also for package-to-package alignment. A new jig that enables several dice to be attached simultaneously was developed to insure proper alignment of the 28 slivers in each package.

The matching of the light output of all 140 LEDs in a cluster is an important part of the assembly process. Because of the large number of LEDs in a single cluster plus the possibility of damage during the bonding process, some dim or extra bright diodes may show up in pre-package test when all 140 LEDs are lit for a uniform brightness test. Those that fail the test are replaced by other LEDs from the same brightness category established for that array.

After sealing, leak testing, and temperature cycling, each cluster is given a final test for the IC parameters and re-visualized for LED uniformity, digit-to-digit uniformity, and brightness categorization.

Acknowledgments

The authors would like to acknowledge the work of David Giuliani and Stan Gage on system architecture, Chen Wu, Jake Haskell and Bob Weissman on the LED display chip development, Les Brodie on debugging the IC, Ted Howell on sliver development and IC characterization, Bob Teichner on the prototype packages, John Scarangelo on product testing, and Ron Pitman on reliability testing. 📧

Reference

1. H.C. Borden and R.L. Steward, "Solid-State Alphanumeric Displays," Hewlett-Packard Journal, July 1970.



John J. Uebbing


John Uebbing worked seven years on such varied projects as high-energy molecular beams, III-V photocathodes, and electron spectrometers before joining Hewlett-Packard in 1973 as project leader on the small alphanumeric display. A member of the International Society of Microelectronics, John obtained his education at the University of Notre Dame (BSEE, 1960), Massachusetts Institute of Technology (MSEE, 1962), and Stanford University (PhD, 1967). Now living

in Palo Alto, California, with his wife and two children (6 and 8), John does some hiking and enjoys sailing on San Francisco bay and the ocean.



Peter B. Ashkin

Born in New Rochelle, New York, and raised in Terre Haute, Indiana, Peter Ashkin served as a cadet at the U.S. Coast Guard Academy before switching to the Massachusetts Institute of Technology, where he earned BSEE and MSEE degrees. He joined HP's Optoelectronic Division in 1973, initially working on IC development for displays. Obtaining several patents along the way, he is now engaged in the development of CMOS/SOS microprocessors at HP's Data Systems Division. Outside of working hours, he works at restoring a vintage Lotus sports car and he also does some mountaineering. He and his wife live in Campbell, California.



Jack L. Hines

Jack Hines had already been involved 8 years in transistor packaging equipment engineering at the time he joined HP's first integrated-circuit operation in 1966, where he's been doing package design ever since. A native of Akron, Ohio, Jack spent seven years in the U.S. Navy before going to college under the GI Bill to earn a BSME (University of Minnesota, 1954). He worked as a plastic molding engineer before getting involved in semiconductor packaging. Living now in Sunnyvale, California, Jack, who has two grown sons, enjoys sailing the ocean and San Francisco bay in his 30-foot craft. He's also been heavily involved in archery, both hunting and target, gaining professional status some years ago.

SPECIFICATIONS

HDSP-2000 Solid State Alphanumeric Display

POWER SUPPLY:

V_{CC} : 4.75—5.25 volts.

I_{CC} : 70 mA, typical.

OPERATING COLUMN VOLTAGE: 3.0—6.0 volts.

COLUMN CURRENT (all drivers "on"): 350 mA, typical.

POWER DISSIPATION: 15 LEDs on per character, $V_{col} = 3.0$ volts: 600 mW typical.

DATA INPUT AND OUTPUT:

LOGIC TYPE: Low-power TTL.

MAXIMUM CLOCK FREQUENCY: 1.5 MHz.

BRIGHTNESS INPUT:

BRIGHTNESS CONTROL VOLTAGE: <0.4 volts, "off"; >2.0 volts, "on".

BRIGHTNESS CONTROL CURRENT: -0.3 mA "off"; 10 μ A "on".

PEAK WAVELENGTH: 655nm (red).

TEMPERATURE RANGE:

-20°C to +70°C, operating

-55°C to +100°C, non-operating

DIMENSIONS: 17.7 mm long \times 7.25 mm wide \times 5 mm above seating plane (0.7 \times 0.3 \times 0.2 inches).

PRICE IN U.S.A.: \$47 each in small (125) quantities.

MANUFACTURING DIVISION: OPTOELECTRONICS DIVISION

640 Page Mill Road

Palo Alto, California 94304 U.S.A.

Hewlett-Packard Company, 1501 Page Mill Road, Palo Alto, California 94304

HEWLETT-PACKARD JOURNAL

APRIL 1977 Volume 28 • Number 8

Technical information from the Laboratories of
Hewlett-Packard Company

Hewlett-Packard Central Mailing Department
Van Heuven Goedhartlaan 121
Amstelveen-1134 The Netherlands
Yokogawa-Hewlett-Packard Ltd., Shibuya-Ku
Tokyo 151 Japan

Editorial Director • Howard L. Roberts
Managing Editor • Richard P. Dolan
Art Director, Photographer • Arvid A. Danielson
Illustrator • Susan E. Wright
Administrative Services, Typography • Anne S. LoPresti
European Production Manager • Dick Leekma

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

CHANGE OF ADDRESS. To change your address or delete your name from our mailing list please send us your old address label (it peels off). Send changes to Hewlett-Packard Journal, 1501 Page Mill Road, Palo Alto, California 94304 U.S.A. Allow 60 days.