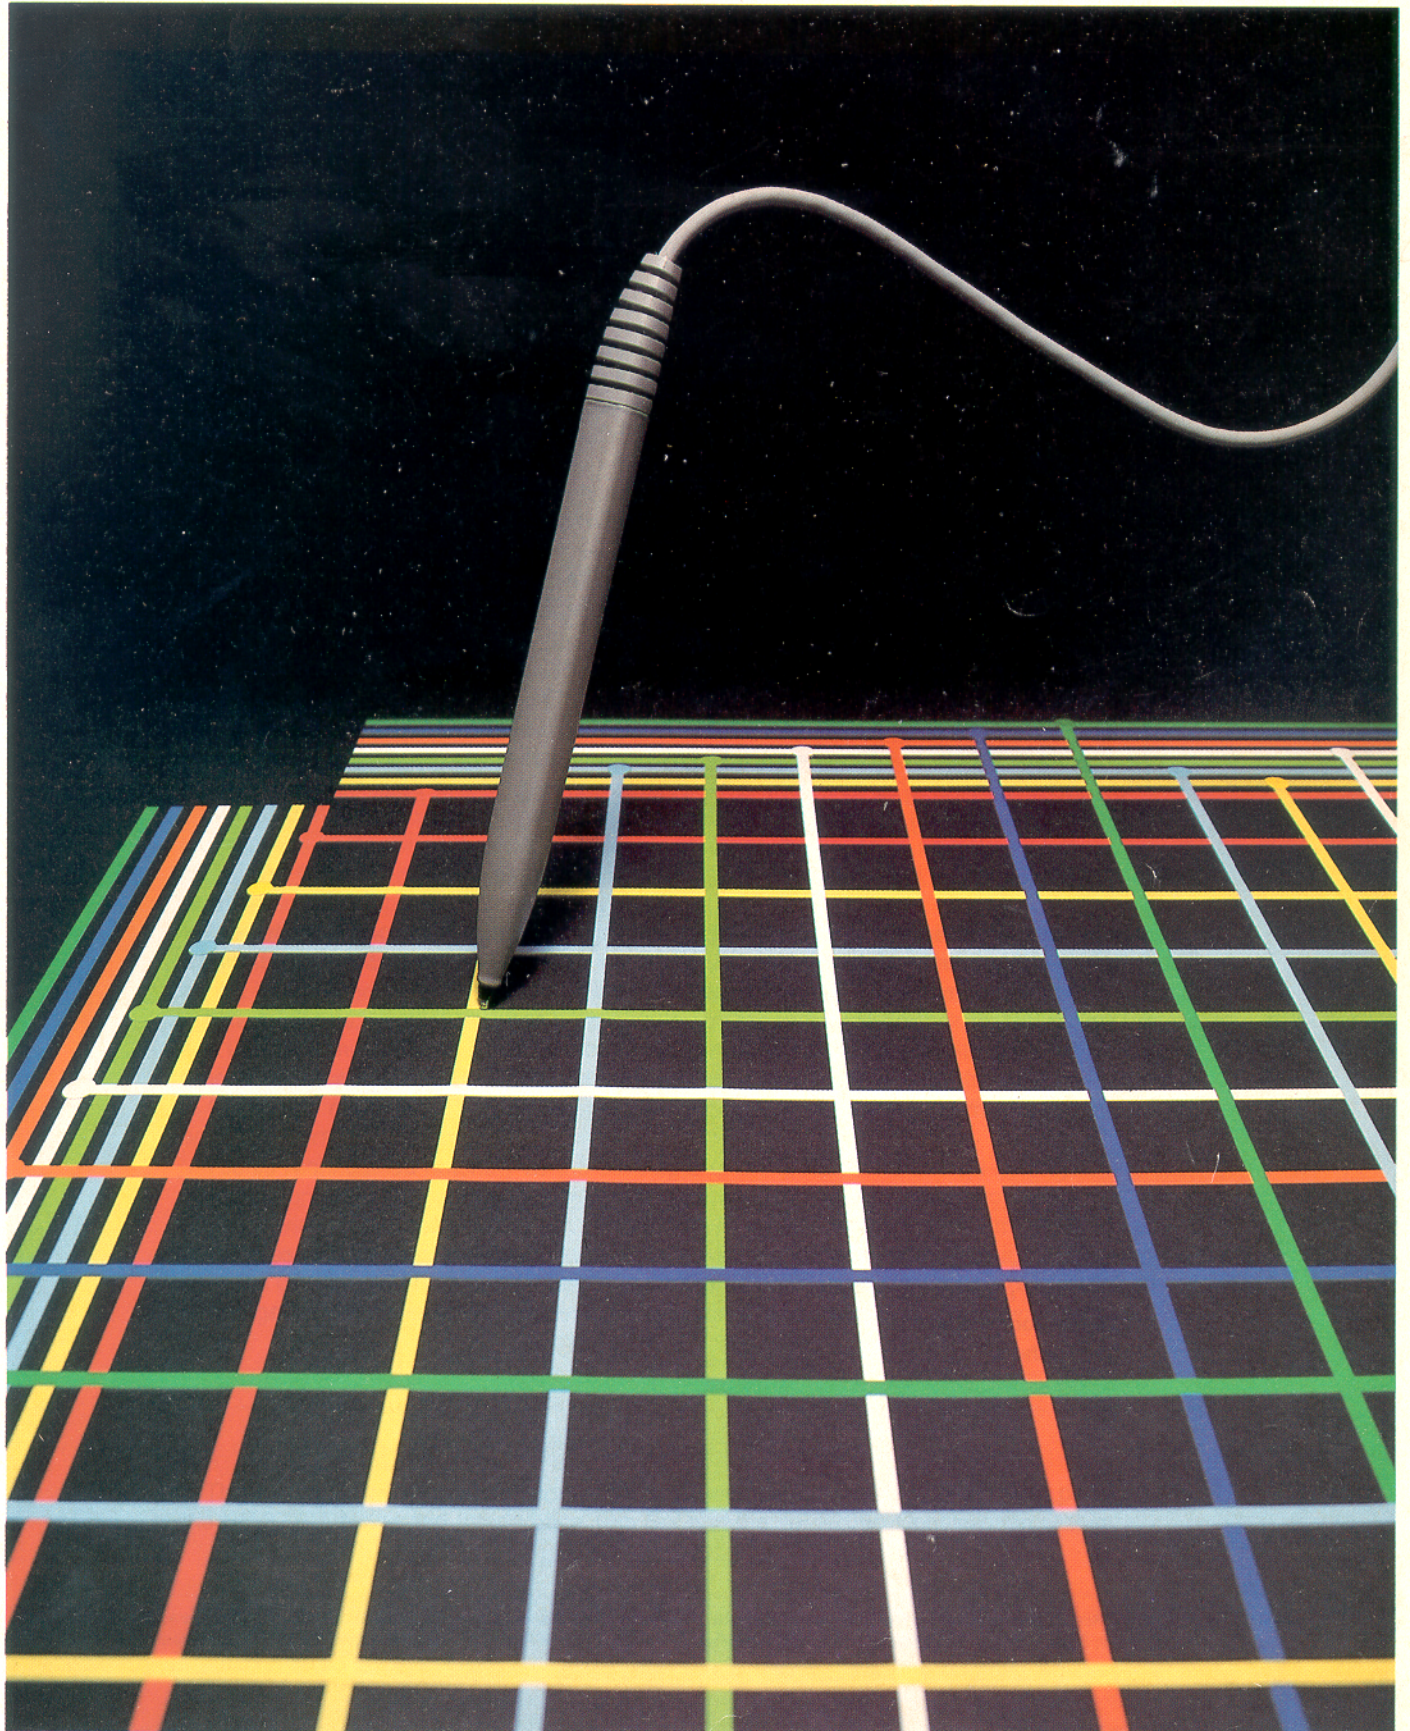# HEWLETT-PACKARD
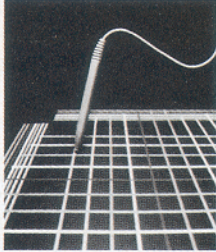# JOURNAL

In this Issue

Our cover story this month is one of those classic engineering tales. Once upon a time, two HP engineers were discussing how nice it would be to have graphics tablets for their home computers. Wasn't it too bad that the tablets—even HP's—were so expensive? In their free time, they began to design graphics tablets, and they eventually hit on a simple, elegant scheme that reduced the number of parts needed, and therefore the cost, without reducing resolution. The idea is to use just a few conductive traces under the tablet surface instead of hundreds, but to use them over and over in different orders at different locations. They call it permuted trace ordering. Our cover photo illustrates it using a different color for each trace. The graphics tablet, the HP 45911A, costs less than a quarter of what previous tablets cost for the same resolution. The full story is told by Tom Malzbender in the article on page 4.

A graphics tablet is a device that a human uses to communicate graphical data to a computer by pointing with a stylus. You can use it for sketching, drawing, computer-aided design, or menu picking. On the other hand, you might choose some other device, such as a mouse, a touchscreen, the keyboard, a digitizer, or a knob. All of these computer input devices, and others too, operate at human speed, which by computer standards is pretty slow. Hewlett-Packard has a low-cost standard interface for connecting devices of this kind to personal computers and workstations. Called the HP Human Interface Link, or HP-HIL (not to be confused with the HP Interface Bus, HP-IB, or the HP Interface Loop, HP-IL), it allows you to connect up to seven devices to a single port on the computer. To find out how it works, read the article on page 8.

If you're interested in AT&T's UNIX® operating system or in HP's version of it, HP-UX, you've probably already read a lot of the extensive literature on the subject. Even so, you may find some new insights in the paper on page 26, which compares the use and performance of the various interprocess communication facilities available in this multiprocessing operating system. Signals, pipes, shared memory, semaphores, and message queues are ranked for various uses and data is presented to support the ranking.

Branches are decision points in computer programs. Branch analysis is a method of assessing the thoroughness of software testing by keeping track of how many branches have been executed by the test procedure and how many have not. Although it sounds simple enough, the first HP software laboratories that imposed branch coverage requirements on their testing projects found that there are many pitfalls, such as attempting to meet the coverage goal by testing all the easy branches instead of the critical ones. In the paper on page 13, three HP software engineers warn of the pitfalls and lay out a comprehensive methodology for avoiding them and reaping all the benefits of branch analysis.

Yoshio Nishi gained notoriety as the developer of the first commercial 1M-byte dynamic read/write memory chip. That work was done when he headed Toshiba Semiconductor Group's semiconductor device engineering laboratory. Brought to HP by an exchange program between the two companies, Dr. Nishi now directs HP Laboratories' silicon VLSI research laboratory. On page 24, he gives us his view of the current status of CMOS technology and lists some of the engineering challenges facing this technology as we approach the era of ultra-large-scale integration (ULSI).

-R. P. Dolan

What's Ahead

The July issue tells the design story of two instruments for evaluating digital radio performance: the HP 3708A Noise and Interference Test Set and the HP 3709A Constellation Display.

Permuted Trace Ordering Allows Low-Cost, High-Resolution Graphics Input

A scheme that substantially reduces the number of trace drivers required provides an inexpensive, but high-performance graphics tablet for HP's HP-HIL family.

by Thomas Malzbender

THE TASK OF ANY GRAPHICS TABLET is to provide the host computer with information corresponding to the position of a pen-like stylus relative to the top surface of the tablet, commonly referred to as the platen. This capability allows the user to input graphical data in a more natural manner for applications such as menu picking, CAD (computer-aided design), sketching, and drawing.

Based on a new input technology, the HP 45911A Graphics Tablet (Fig. 1) represents a significant contribution in price/performance for this class of graphics input devices. Less than a quarter of the cost of earlier HP graphics tablets, the HP 45911A offers a resolution of 1200 lines per inch (0.02 mm) with essentially no jitter at this high resolution. Its active area was chosen to be 11 inches per side to accommodate standard overlays produced by third-party software vendors. Ergonomically, the HP 45911A features a minimal footprint, low-profile package designed to be used in front of large workstations like HP's Vectra Computer without restricting easy access to the system's disc

drives. In addition, its standard width of 325 mm allows it to be stacked on top of the system when not in use.

The development history of the HP 45911A is reminiscent of HP's early development style in which projects were initiated by lab engineers with a need for a new product and who believed they had a good idea on how to construct it. Early in 1984, Mike Berke and I were griping about the high price of the HP 9111A (HP's only tablet back then) and how useful a good inexpensive tablet would be for our home computer systems. So when time would allow it, we started experimenting with various tablet designs. After prototyping several technologies (electrostatic, magnetic, optical, and ultrasonic), it became clear that an electrostatic approach was our only choice. Magnetic technology also promised high resolution, but required higher current consumption and had a significant problem with sensitivity being highly dependent on the angle of the pen to the tablet surface.

In an electrostatic design, traces underneath the active
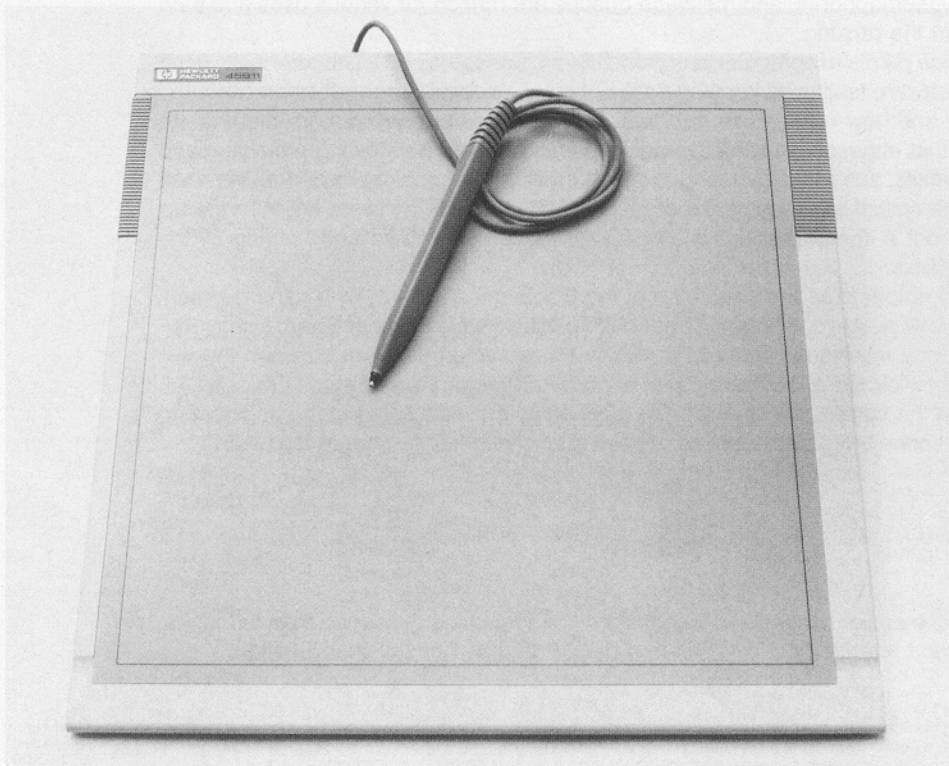
Fig. 1. *The HP 45911A Graphics Tablet features a high resolution of 1200 lines per inch over an 11 × 11-inch active area for a low price. A member of HP's Human Interface Link (HP-HIL) family of input devices, it can be used with a variety of HP computers and workstations.*

area are sequentially pulsed and these pulses are capacitively coupled to the tip of a stylus. The amount of coupling is a function of the local dielectric coefficients (which are normally constant) and the spatial separation between the stylus and any specific trace. Hence, the stylus position can be accurately determined by the relative strength of the signals coupled back from the traces as they are pulsed.

The HP 45911A offers greatly reduced hardware complexity over comparable tablets by using a technique (patent applied for) that reduces the number of trace drive lines from over 110 down to 16. This scheme, called permuted trace ordering (PTO), allows us to drive the traces directly from the on-board microprocessor, eliminating the need for any separate driver ICs, which usually represent a large fraction of the cost of a graphics tablet. To accomplish this reduction, the same drive lines are used over and over again on the 112 vertical and horizontal traces on the tablet by varying the sequential ordering of the traces along the tablet surface. In this way, a unique signature is coupled into the stylus at all points on the platen. Fig. 2 demonstrates this for a section of the tablet platen board. It shows, for both axes, the drivers associated with each of the traces shown. The tablet operates by activating the trace drivers singly, in sequence, and reading the stylus response for each trace driver.

For example, if the stylus is located as shown, the outcome might resemble the list of values shown in Table I.

The units in the response column are merely relative values and could be viewed as results from an 8-bit analog-to-digital (A-to-D) conversion.

Table I
**X-axis trace driver responses**

Driver	Response
0	160
1	60
2	30
3	225
4	100
5	25
6	35
7	50

Observe how the stylus response is a function of stylus-to-trace distance. Driver 3 is the closest and gives the highest response, followed by driver 0, then driver 4. These top three responses can be formed into a code, say 304, which, by design, is unique to that coarse position on the tablet surface. The magnified section of Fig. 2 shows coarse position codes for both X and Y in a small region of the tablet. Note that two different code values occur within a trace spacing. Each of the drivers is pulsed one at a time, the
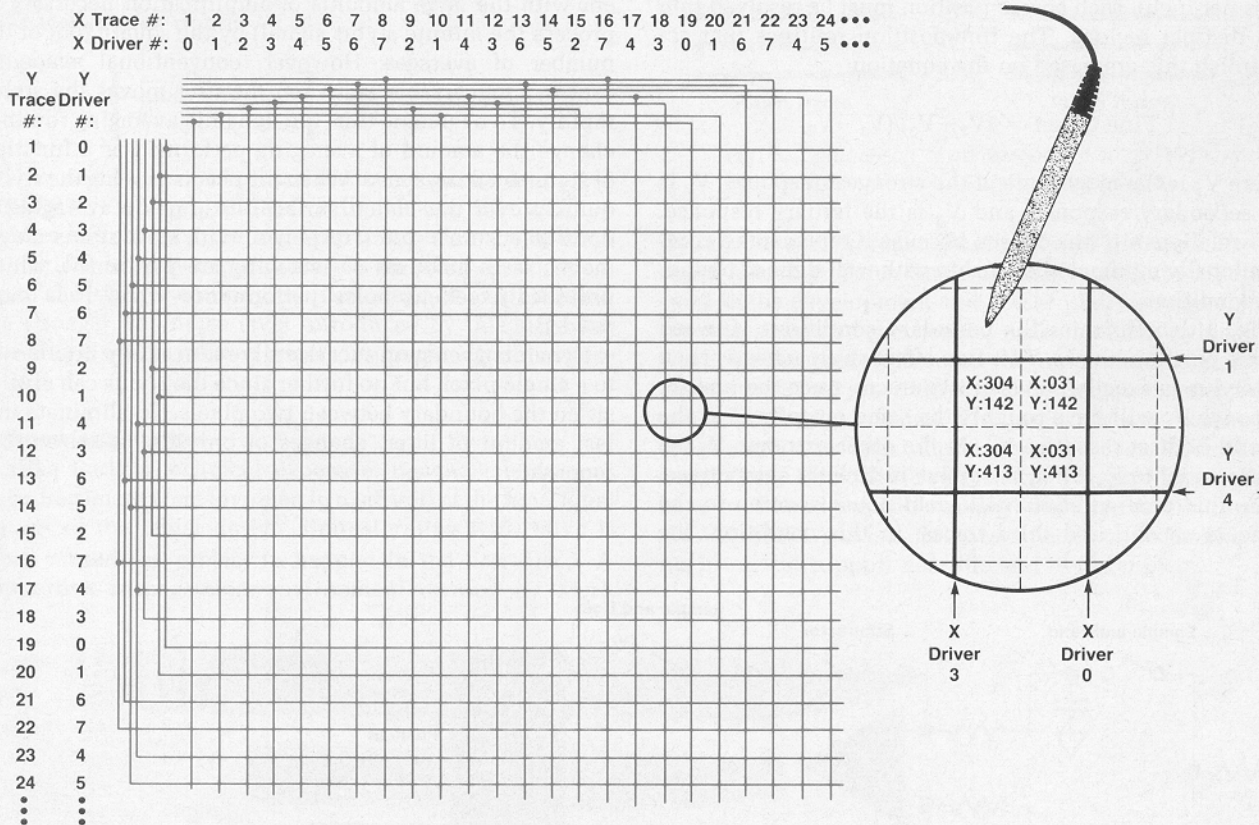
Fig. 2. A section (left) of the HP 45911A platen with an expanded view (right) showing coarse position codes for X and Y coordinates. The first digit of the code corresponds to the closest trace driver, the second digit corresponds to the second closest, and the third digit to the third closest. The driver sequence is chosen to assure unique codes for all coarse positions.
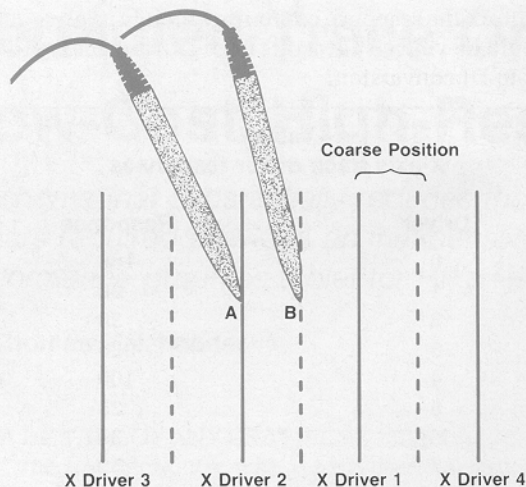
Fig. 3. *At stylus position A, the fine offset value is 0 and at position B the value is 1. Continuous values between 0 and 1 exist between positions A and B.*

responses are measured and sorted, and then the driver numbers for the three strongest responses are combined into a code word. This code word then becomes an address for accessing a coarse-position lookup table.

These algorithms determine coarse position with a resolution of 2.7 mm. To achieve a resolution of 0.02 mm (1200 lines per inch), each coarse position must be resolved into 128 distinct regions. The fine-position routines that accomplish this are based on the equation:

$$\text{Fine Offset} = (V_2 - V_3)/(V_1 - V_3)$$

where V_1 is the magnitude of the strongest response, V_2 is the secondary response, and V_3 is the tertiary response. This relationship was chosen because it represents a computationally minimal relationship with well-defined boundary conditions.

Fig. 3 demonstrates the boundary conditions between coarse position blocks. This fine offset approaches 1 when the stylus is exactly between two traces since the highest response V_1 will have roughly the same magnitude as the second highest response V_2. At the other extreme, V_2 becomes equal to V_3 when the stylus is directly over a trace, since this configuration will yield equal spacing to the adjacent second and third traces. In this condition, the

numerator and the fine offset itself approach zero. Between these two extremes, the values are continuous but not necessarily linear. Linearization is achieved through the use of a lookup table within the HP 45911A's microprocessor, and we are left with a flat position response at high resolution.

There is a fundamental relationship between computation speed and noise/jitter performance. Fast position determinations make it possible to use averaging to reduce any noise in the system. For this reason, the fine offset equation is computed in hardware rather than firmware. Referring to the configuration shown in Fig. 4, the subtractor stage is used to generate both the terms $V_1 - V_3$ and $V_2 - V_3$. The first term is applied to the reference input of the analog-to-digital converter (ADC) and the other term is applied to the ADC's signal input. The effect of this is a division of the two terms. Since the speed of this process is limited only by the signal propagation and A-to-D conversion times (dominant here), data can be collected quickly and averaged often. In addition, multiple samples can be taken on the input sample-and-hold circuits, which causes very quick analog averaging to take place there. The result is excellent noise performance.

Noise performance is further improved by two firmware routines, dynamic averaging and antijitter. Dynamic averaging is a technique that offers all the benefits of large amounts of position determination averaging without the drawbacks. Averaging reduces the amount of noise (inherent with the large amounts of amplification necessary to process the minute stylus signal) by the square root of the number of averages. However, conventional averaging causes a perceivable lag when the user moves the stylus rapidly. To overcome this, the dynamic averaging routines change the amount of averaging performed as a function of stylus tracking speed. When the user is moving the stylus quickly over the platen surface, little or no averaging is done to ensure a quick response. With slow stylus movements, large amounts of averaging are performed, which provides excellent noise performance when it is most needed.

Dynamic averaging successfully reduces any jitter down to a single pixel, but no further since the stylus can always sit on the boundary between two pixels. To eliminate this last amount of jitter, changes of only one pixel are not reported.
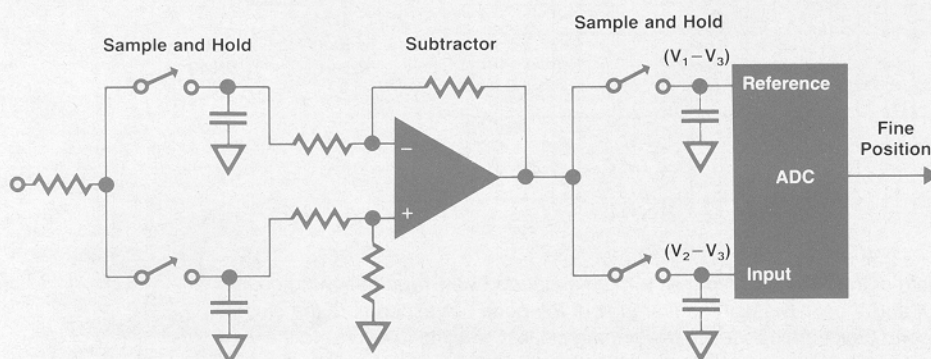
Fig. 4. *Hardware system for computing fine offset position.*

Stylus Design

The conventional electrostatic graphics tablet stylus can be thought of as merely a shielded wire that brings the capacitively coupled trace signals back to the main system electronics for amplification and processing. The stylus for the HP 45911A, on the other hand, is active and the trace signals are amplified at the stylus tip before they are sent back to the main electronic system.

The use of surface mount components (see Fig. 5) let us put the first stages of amplification in the stylus. Although this approach requires power and ground wires to be connected to the stylus, it improves noise performance by roughly an order of magnitude.

Signals seen by the tip are greatly reduced by a parasitic voltage divider formed by any existing tip-to-ground capacitance. In a conventional stylus, the tip and attached wiring running through both the body of the stylus and the cable shield form a considerable parasitic divider.

That is (Fig. 6, left):

$$V_{Stylus} = \left[\frac{C_{Trace}}{C_{Trace} + C_{Body} + C_{Shield}} \right] V_{Platen}$$

Given typical values of 1 pF for C_{Trace} and 100 pF for $C_{Body} + C_{Shield}$, the stylus voltage is approximately $(1/101)V_{Platen}$.

The HP 45911A stylus tip sees only the parasitic tip-to-body capacitance, yielding a signal about ten times stronger at the input to the first stage of amplification.

That is (Fig. 6, right):

$$V_{Stylus} = \left[\frac{C_{Trace}}{C_{Body} + C_{Trace}} \right] V_{Platen} \times Amplifier\ Gain$$

In this case, $C_{Body} = 10$ pF and the stylus voltage is approximately $(1/11)$ V_{Platen} multiplied by the amplifier gain.

After buffering by the low-output-impedance amplifier, any shield or body capacitance has no effect. In addition, since the signals entering the first stage of amplification are stronger, the noise level introduced by this stage has less effect, which yields a greatly improved signal-to-noise ratio.

Acknowledgments
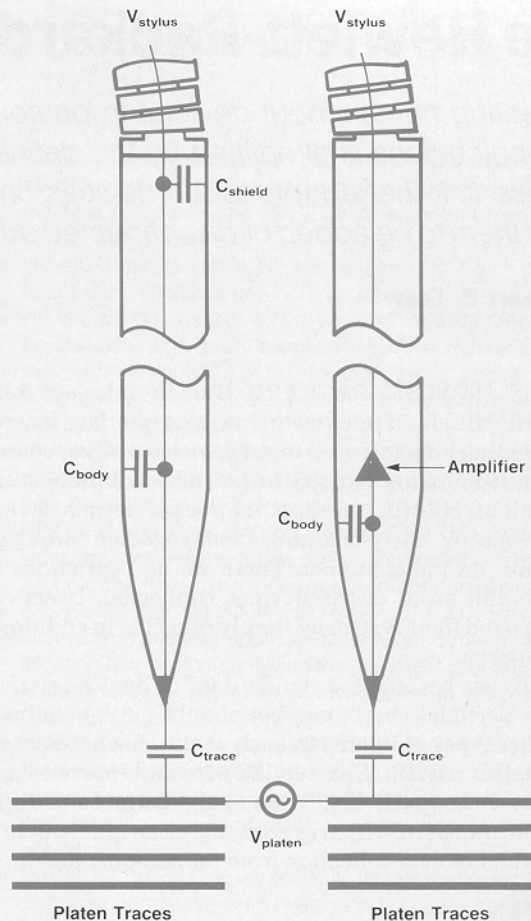
Fig. 6. *Comparison of conventional electrostatic stylus design (left) with HP 45911A stylus design (right).*
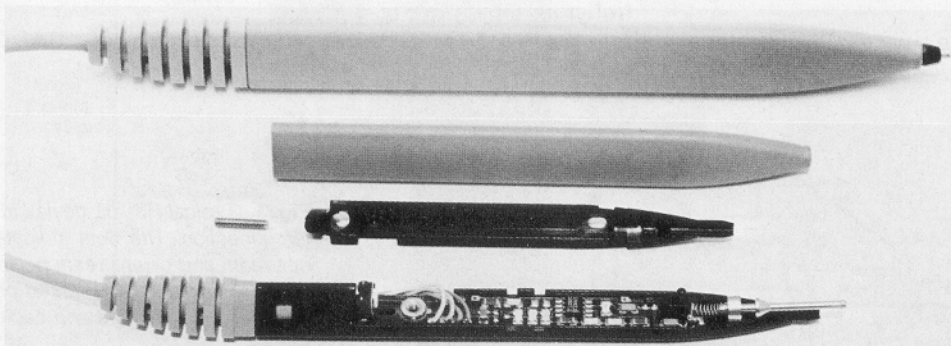

Fig. 5. *Assembled (top) and disassembled (bottom) stylus assembly showing amplifier board using surface mount components.*