Keyboard

Mar-Apr/80 A Publication of Hewlett-Packard Desktop Computer Division

Keyboard

March-April 1980

Cover

Fading light silhouettes workers at the 500-kW geothermal power plant at East Mesa in California's Imperial Valley. The photograph was taken by Doug McWilliams, photographer for Lawrence Berkeley Laboratory.

The laboratory administers geothermal operations in Southern California for the U.S. Department of Energy.

Photo and artwork credits

cover, page 4 — Doug McWilliams, Lawrence Berkeley Laboratory
pages 1, 3 (left), 10, 11, 13 — Paula Dennee
page 2 — Jolly Tiger Restaurants and Donald J. McClenahan, U.S. Atmospheric Sciences Research Center, University of Albany, Albany, NY.
page 3 (right) — Barber-Nichols Engineering Company, Arvada, Colorado
pages 5 through 8 — Chris Eilbeck, Heriot-Watt University, Edinburgh, U.K.
page 12 — HP photo

Breaking the energy habit

by John Monahan
Hewlett-Packard Company
Desktop Computer Division

It's easy to quit smoking. I've done it hundreds of times myself, U.S. humorist Mark Twain once said.

And so it seems with energy: breaking the habit of using it exorbitantly is more easily said than done.

But recently, systems monitored by desktop computers have begun yielding substantial energy savings for the people employing them.

Widespread applications

A gas company in California, a restaurant in New York, a consulting firm in Tennessee and an engineering company in Colorado are examples of how energy can be conserved when there's a will and a way.

These applications show how energy is being saved in a prototype house, at a small business, in large commercial structures and in computerized research on an alternative form of energy.

HOUSE

Twenty percent of the energy consumed in the U.S. is used in houses, according to the Southern California Gas Company, Los Angeles, California. If such consumption could be reduced by 50%, 300 000 barrels of oil could be saved every day, according to their research.

California being the land of mass-produced, or "tract" houses, Southern California Gas in 1976 built two three-bedroom, Spanish-style tract houses to demonstrate that energy can be saved without compromising the conveniences Americans generally consider necessities. Data was collected from dozens of sensor-based instruments using a Hewlett-Packard 9825A Desktop Computer.

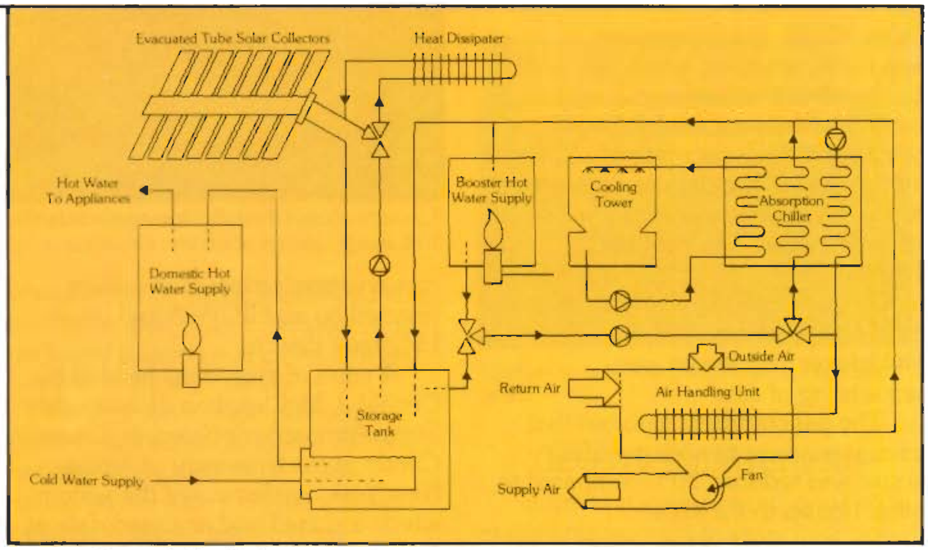

Diagram shows relationship between major components of a minimum energy dwelling, as developed by Southern California Gas, Mission Viejo Construction Company and the U.S. Department of Energy.

The project, called MED for Minimum Energy Dwelling, was undertaken by the gas company along with the Mission Viejo Construction Company, Mission Viego, California; and the U.S. Department of Energy, Washington, D.C.

One of the houses was used for testing and was occupied by a family; the other, essentially identical to the test house, was built for public display purposes.

Materials used in their construction are generally available.

Innovations

Construction innovations in the houses can be divided into those you can see and those you can't.

The latter category includes insulation in the walls and underneath the house, plus a plastic sheet (rather than standard tar paper) between exterior and interior walls. Its purpose is to seal out moisture and heat more effectively; keeping cool, not staying warm, is how Southern Californians spend most of their energy.

Visible features include sun screens, foliage that provides shade, an air-lock-like entryway, a red tile roof and louvered windows.

An "economizer" system monitors differences between inside and outside heat and humidity. When the outside temperature drops, the system automatically shuts off the air conditioner, draws in the cooler air and expels the hot air trapped inside.

Energy-saving appliances and devices that save water also were used.

When necessary, the house is heated or cooled by natural gas-assisted solar collectors, which also preheat the domestic water supply.

Data acquisition

In order to measure how all these innovations worked, especially in relation to an ordinary tract house, Southern California Gas selected a data acquisition system that included a Neff 620 Multiplexer interfaced to the HP 9825 Desktop Computer by an HP 98032 16-bit Parallel Interface. HP's 98034A HP-IB Interface permitted an HP timer and an HP real time clock to be connected to the computer.

Every 30 minutes data was gathered by sensors at 97 locations.

Using HP-IB, this information — about temperatures; water, air, and gas flows; and electricity consumption — was sent to the computer so that calculations could be made on site. The data was integrated by the 9825, then stored on magnetic tape cartridges. The HP 9866A Thermal Printer provided output on request, while the real time clock and timer provided the 9825 and Multiplexer with the proper sequencing of data.

The gas company reported that energy consumption for the MED house was reduced 50% compared to other houses in the area.

PUBLIC BUILDINGS

While some engineers are looking into ways of conserving energy where people live, others are finding ways to reduce energy consumption where people congregate; places like restaurants.

On winter days customers expect a restaurant dining room to be as warm as toast, while on warm days they want to dine in cool surroundings.

A restaurant faces unique problems in energy conservation — from maintaining the comfort of its patrons to ensuring the freshness of its food.

Conservation/reclamation

The Jolly Tiger restaurant of Albany, N.Y., was designed to turn problems into advantages by using a sophisticated energy conservation/reclamation system directed by an HP 9825 and other HP-made devices.

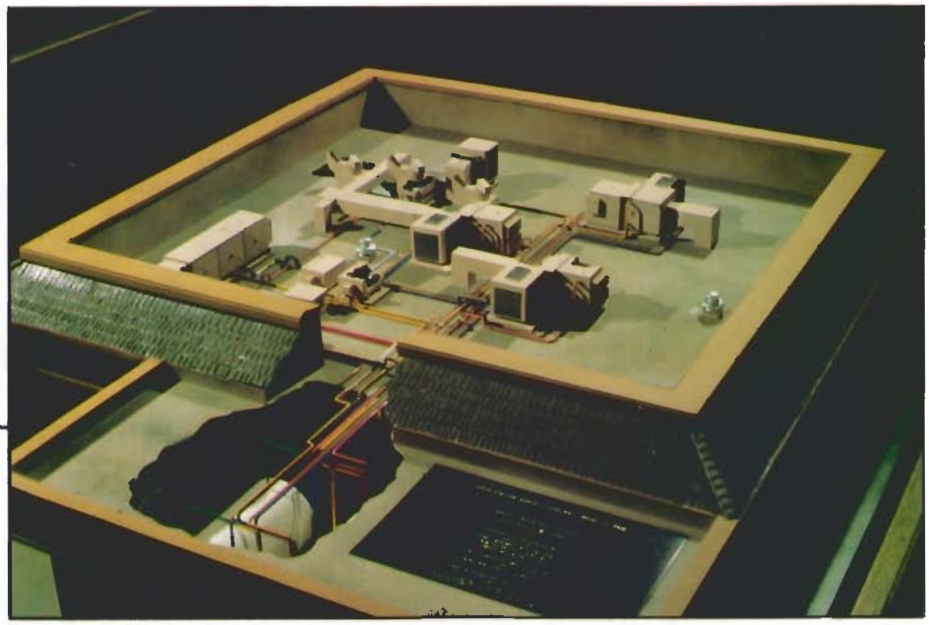

Rooftop model of the Jolly Tiger restaurant shows the equipment which made much of the business's 30% energy savings possible.

A team of engineers, headed by Donald J. McClenahan of the United States Atmospheric Sciences Research Center at the University of Albany, New York, implemented the system, which acquired and analyzed data at critical points, while monitoring it for possible trouble spots. The research started in 1976 and the initial phase was completed in 1978.

The Jolly Tiger energy system was built first to conserve energy, then to reclaim energy used as a necessary part of business.

Consumption was reduced by employing three heat pumps, air conditioners, a photocell-controlled lighting system, utilization of outside air for some cooling, and a clock-controlled, variable-speed, exhaust fan over the cooking grill. There were several other, more discrete, subsystems as well, including the refrigeration rack, dishwasher drain and the cooling-grill, exhaust-heat pump.

Reclamation efforts were based around two, 11 350-liter (3000-gallon), water tanks located in a separate room. It was to them, for instance, that heated water from the refrigeration and air conditioning compressors was sent, to be used later in the dishwasher and for public use.

Reclaimed heat from the dishwasher drain was also recovered and sent to the storage tanks for reuse. Ingenious measures such as these were key to making energy do double duty.

Data system

The data acquisition system needed to run such an encompassing system had to be powerful and flexible, and of course its cost was an important factor, says McClenahan. Low maintenance was another consideration, he says. The engineering team chose the HP 9825 Desktop Computer and other devices made by Hewlett-Packard to interface with components of the data acquisition system.

Every two minutes the 9825 scanned readings from the 215 sensors located in the various critical areas, such as the heat pumps, cooling area exhaust, refrigeration units, dishwasher and lights.

The desktop computer then made real-time analyses and printed the results with the HP 9871 Printer. Results were stored in the HP 9877 Tape Drive. An unanticipated advantage, McClenahan says, was that periodic readings on the performance of the equipment could be made without interrupting normal data acquisition. This proved "invaluable" for troubleshooting the restaurant's heating and cooling system, he says.

The second phase of McClenahan's research involves the use of 305 square meters (1000 square feet) of solar collectors on the roof of the restaurant. How much the collectors have affected energy savings has yet to be determined, but McClenahan has estimated that the equipment employed in phase one saved the Jolly Tiger 30% in energy usage.

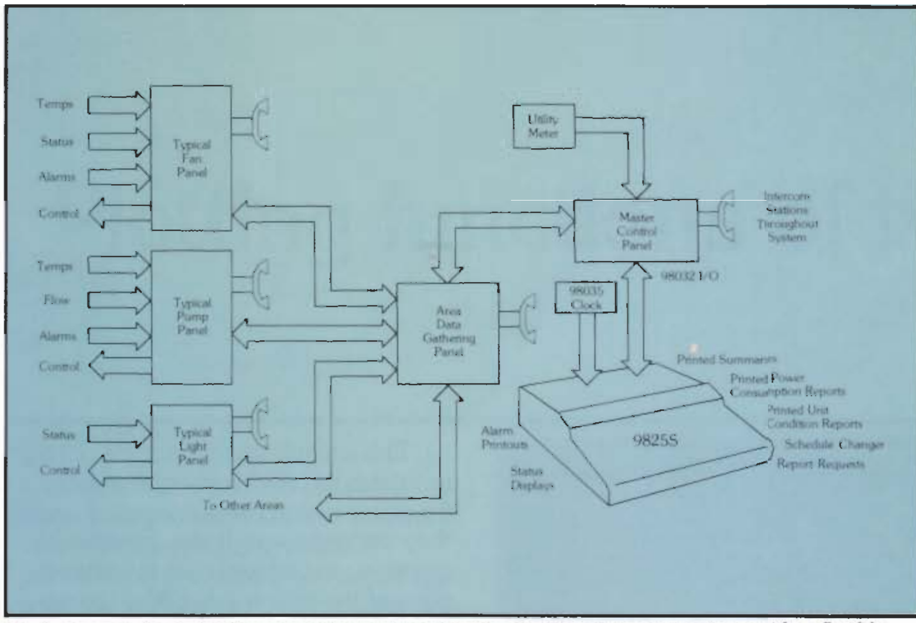

Block diagram illustrates the concepts developed by the PreCon Corporation to provide a flexible system to maintain optimum conditions in large buildings at minimum cost.

LARGE STRUCTURES

So far we've seen energy conservation efforts in buildings that are small compared to the 6 x 10^5 square meters (2 x 10^6 square feet) a manufacturing plant might contain. The PreCon Corporation of Memphis, Tennessee, tailors systems to suit the conservation requirements of large, complex installations such as industrial plants, hospitals, campus buildings and similar facilities.

PreCon initially offered manually operated multiplex panels used for monitoring and controlling heating, cooling and ventilation systems in complex installations. But in 1975, as operator turnover became a problem and it became necessary to keep tighter control over energy costs, the company began looking at the feasibility of using a computer to control the panels.

The idea was inspired, says PreCon's president David Zietz, by another computer he purchased for office use. It initiated him to computers, and he shopped around for one that wouldn't require a human operator, he says.

He decided on the HP 9825 Desktop Computer, which he now also uses in his office. Having it there, he says, has allowed him to see the computer's capabilities, which in turn helps him better determine solutions for his customers.

Complex problems

Their energy problems are myriad, exacerbated by rising energy costs, different heating and cooling needs as workers change shifts, malfunctioning equipment that eludes detection and inefficient systems that can warm air they have already cooled.

To help solve these problems, PreCon devised software for the 9825 that effectively runs the PreCon control panel, based on parameters the system owner and PreCon establish.

Control/data acquisition

The computer performs the essential energy management functions, Zietz says. It permits equipment to operate only when its operation is necessary to maintain minimum acceptable conditions in those areas the equipment serves; it alters the operation of that equipment to help minimize unacceptable utility billing penalties; and it helps assure equipment is running efficiently.

Should a breakdown occur, the computer "beeps," then reports the time and place of the malfunction.

No other computer-related equipment, such as printers or extra memory, is needed, Zietz says, but he adds that the HP 2621 Remote Terminal can be used.

The data fed through the control panel to the computer is collected by thermistors and other sensors.

Zietz says that PreCon's solution is a tool that permits experts such as plant engineers to go about their business without arcane knowledge of computers. It has also, he says, shown managers things about their plants they hadn't detected before, such as employees tampering with thermostats.


Barber-Nichols 500-kW geothermal power plant.

GEOTHERMAL ENERGY

Beyond conserving our present energy resources is the quest for new types of energy that we now use poorly, if at all.

Barber-Nichols Engineering of Arvada, Colorado, is engaged in two projects that can help harness geothermal energy, superheated water found far below the earth's surface. These efforts can ultimately help conserve other dwindling resources.

The process of converting geothermal energy into a usable form such as electricity involves several steps. Hot, miniralized water called brine is first pumped from beneath the earth.

On the surface it gives up its energy through a heat exchange process that heats a working fluid such as isobutane, which vaporizes and runs a turbine to generate electricity. The process is known as the Rankine Cycle.

Design research

This simple description overlooks a major problem: The pump which is used to draw up the brine must be able to withstand fierce temperatures, pressures and huge variations in brine composition.

Brine makeup, for instance, can be corrosive and abrasive, and varies widely even in areas geographically similar.

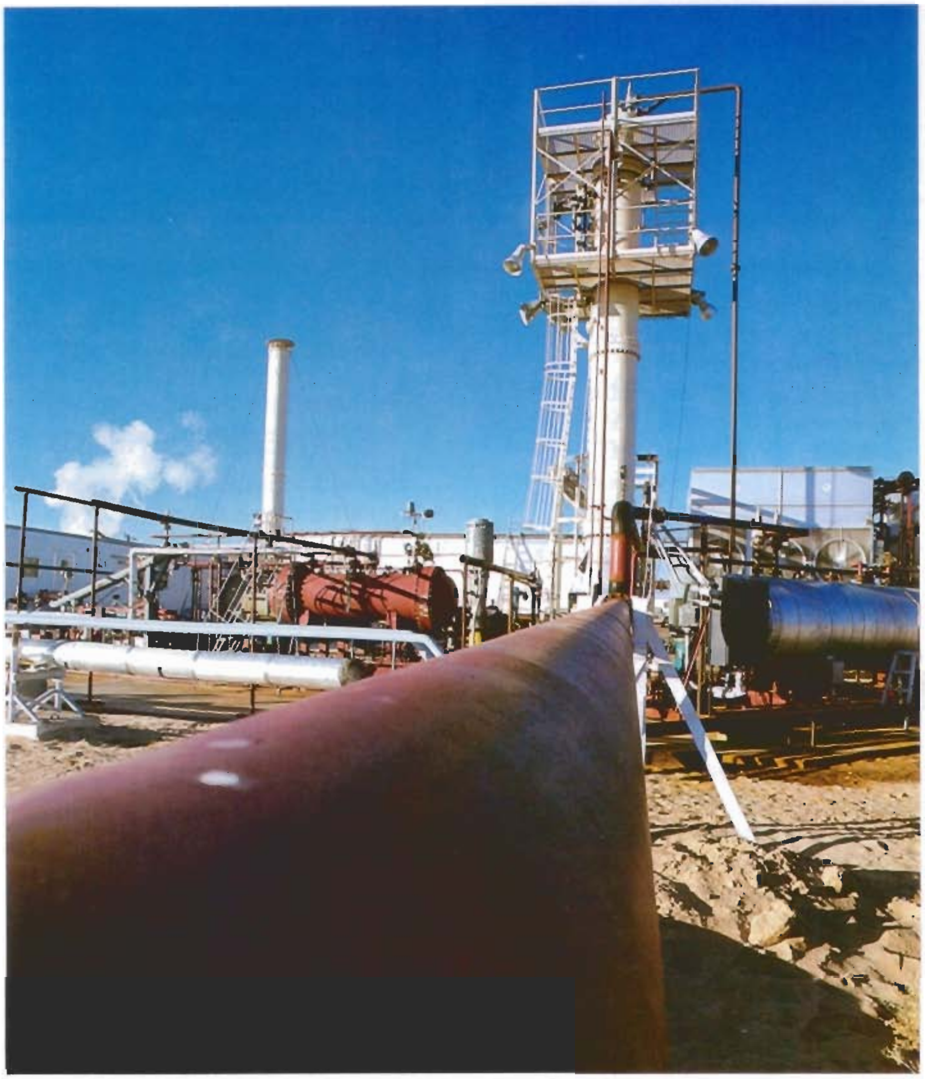

The Barber-Nichols experimental geothermal power plant is currently being tested by the U.S. Department of Energy at the geothermal fields near East Mesa, California.

To determine how pumps hold up under different conditions, without going to the considerable expense of dropping them down a geothermal well, Barber-Nichols developed a method for testing them above ground. A Hewlett-Packard 9825A Desktop Computer monitors equipment and acquires data.

In effect, Barber-Nichols took the drillhole and put it on top of a semi-trailer, in the form of a tower-like rig which houses the pump being tested. With the test rig on site, brine from a geothermal well is admitted to the rig, the pump is activated, and its performance monitored by thermocouples, transducers and other sensors.

A large number of different types and brands of sensors gather data about flow rates, temperatures and pressures. Their findings are reported to the 9825 via the Hewlett-Packard Interface Bus (HP-IB). Data is then output to an HP 9871 Printer

This unique testing process closely simulates the environmental stresses the pumps would be subjected to were they underground. It also permits the company to test well sites in various parts of the country, because the rig can be transported by truck.

Turbine tests

The second Barber-Nichols project demonstrates the potential of geothermal energy.

Once past the pump, the brine is used to vaporize isobutane, which in turn runs a turbine that can produce 500 kilowatts of electricity.

This project also uses the 9825 plus two HP 9885 Flexible Disc Drives for data storage. As the brine is pumped, the 9825 monitors flow rate, temperatures and pressures, while concurrently monitoring the isobutane flow rate, along with temperatures and pressures at various points along the isobutane loop. Results are output to an HP 2631 Printer.

In the United States, as these stories show, and around the world, efforts to conserve energy are underway. Future issues of *Keyboard* will again deal with this subject and how desktop computers are being used to break the habit of squandering energy that, like Mark Twain and his smoking habit, we've broken hundreds of times before. K

Plotting functions of two variables

by Dr. J.C. Eilbeck

Often the output of a calculation or a laboratory experiment is a quantity which depends on one or more parameters; mathematically speaking, a function of one or more variables. In order to visualize the results it is useful to plot the data.

This is a simple process if we have only one variable, but more difficult and time-consuming if the function depends on two variables. For more than two variables we must restrict ourselves to cross-sections to reduce the problem to two variables.

Two plotting methods

Two common methods of plotting functions of two variables are contour plots and 3-D plots with hidden line removal.

The contour plot method draws lines in the x-y plane connecting points of constant function value, whereas the 3-D plot involves drawing graphs of the function at various fixed y values to present the appearance of a three-dimensional solid model whose vertical height at any point is the function value.

These two methods have been chosen with relatively small and slow machines in mind, since they use a minimum of storage and are reasonably fast in execution. The simplicity of these methods means that they can be easily programmed by the user, who can add special features as desired once the basic principles of the algorithms are understood.

The results of the two methods can be judged from Figures 1 and 2 where the function $f(x,y) = r^4 \exp(-r)(1 - 6\cos^2\theta + 9\cos^4\theta)$ is plotted, with $r = \sqrt{x^2 + y^2}$, $\tan\theta = y/x$. This function is the density of the electron wave function $|\Psi|^2$ of the 3d orbital in the hydrogen atom, rotationally symmetric about a horizontal axis.

In both methods the function is

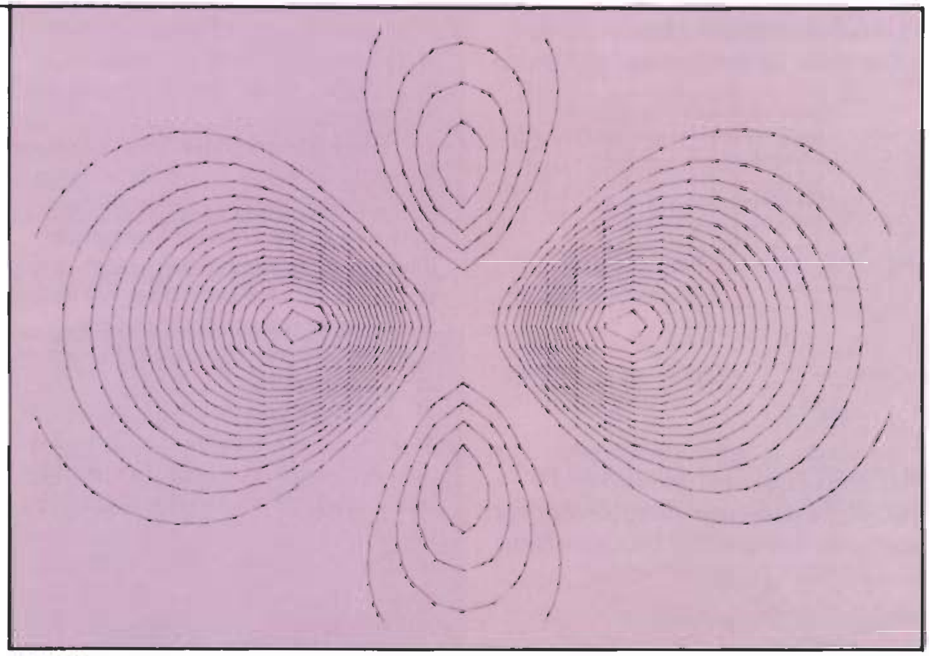

Figure 1
Contour plot of the 3-D hydrogen orbital. The function is axially symmetric about a horizontal axis passing through the large peaks. The grid spacing used in the plot is about 1/30 of the total plot width.

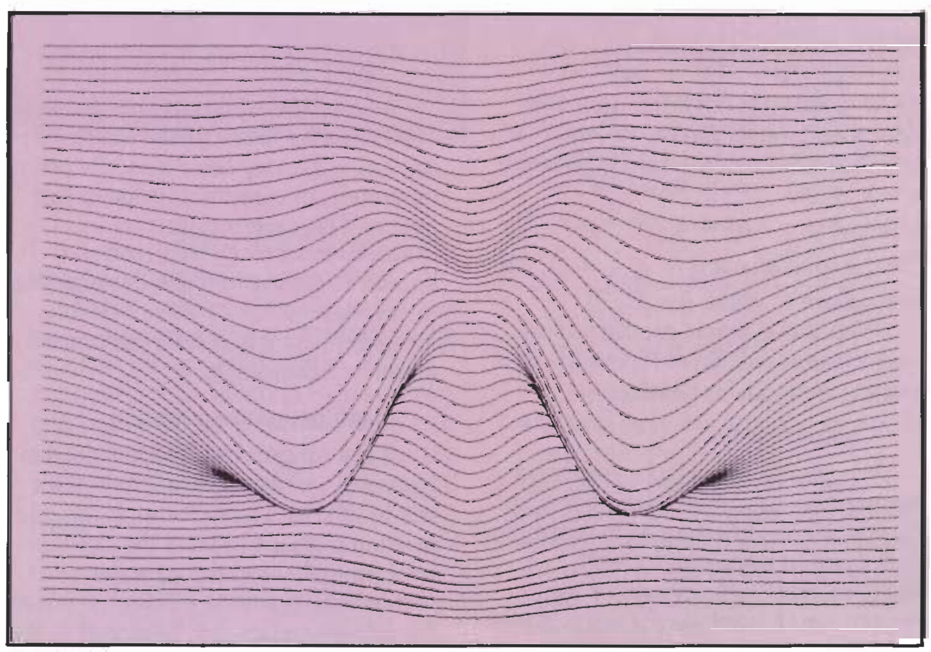

Figure 2
Three-D plot, with hidden line removal, of the same function as in Figure 1.

Each graph represents a 'snapshot' at a fixed time, and the full plot represents the time development of the process.

calculated or recorded on a rectangular grid in the (x,y) plane.

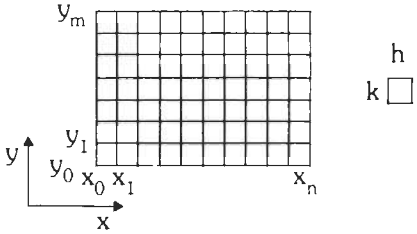

$x_i = x_o + ih, i = O,1,...,n$

$y_i = y_o + jk, j = O,1,...,m$

$f_{i,j} = f(x_i, y_j)$

In order to approximate any function g(x,y) between two grid points, we use linear interpolation. For example, between the two grid points (x_i, y_j) and (x_{i+1}, y_j) we have

$$g(x,y_j) = h^{-1}(x_{i+1} - x)g_{i,j} + h^{-1}(x - x_i)g_{i+1,j} \quad (1)$$

where $h = x_{i+1} - x_i$

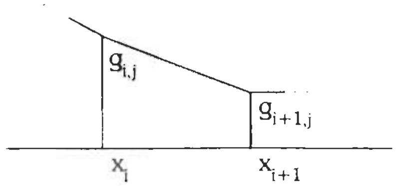

If g_i and g_{i+1} have the opposite sign, then g has a zero in this interval, and the position of the zero is given by $x = x_i + \theta h$

where $\theta = g_i / (g_i - g_{i+1})$. (2)

For the special case $g_i = g_{i+1} = O$, it is convenient to make $\theta = O$.

With these simple mathematical preliminaries we can examine the two plotting methods in detail.

Contour plot

We give here a simple version of the method described in the book by Roach (1972). We scan each rectangle in the grid in turn, and for each function value whose contour we require, say f_c, we look for zeros of the function $g = (f - f_c)$ on each side of the rectangle. This is done by testing the sign of $(f - f_c)$ at each corner of the rectangle, and if a zero is present, its position is calculated using equation 2 and stored.

There are only three possible cases for each rectangle: (i) no zeros, (ii) two zeros, and (iii) four zeros. In the special case where the contour line passes through a corner this will count as two zeros, i.e. one for each side. We ignore the exceptional possibility of the contour line passing through more than two corners.

In case (i) no plotting is required and the grid search moves on to the next rectangle. For case (ii) a straight line is drawn between the two zeros, and if case (iii) occurs, all zero-pair combinations (six in all) are joined.

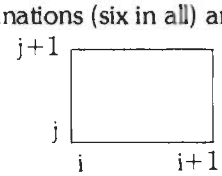

case (i)

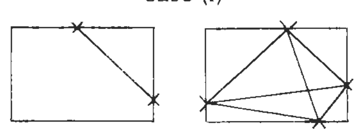

case (ii) case (iii)

The combinations in case (iii) take in all possible forms of the function f, since $(f - f_c)$ may have a saddle point, minimum, maximum or ridge in this region. This case is quite rare in practice, but if more information is known about the function it is possible to calculate which of these possible forms is present and case (iii) can be modified accordingly.

Machine implementation

If the function is already stored in an array, we can draw each contour line in turn by a scan through the grid, pausing perhaps to label the contour appropriately, and this process is reasonably rapid and efficient. If not all the tabular values can be stored in the machine at one time, the plotting area can be subdivided into two or more and each subregion plotted in turn.

However, if the grid point function values are calculated according to a mathematical formula, the evaluation of the function will probably be the most time consuming part of the process. In this case it is perhaps better to have an inner loop to draw each of the different contour lines on each grid rectangle in turn. If the grid search is progressing along the y direction, it is useful to store $f_{i+1,j}$ so that it will not need to be recalculated when the adjacent column is scanned.

The smoothness of the contour lines will depend on the grid spacing — a suitable compromise between the running time (proportional to n x m) and the best possible smoothness will have to be chosen. A further example of a contour plot, showing the concentrations of two chemicals in a model embryo (Catalano et. al. 1980), is shown in Figure 3.

3-D Plot with hidden line removal

In this type of plot we draw a succession of one-dimensional graphs, $f(x,y_j)$ as a function of x, for different values of a fixed y value, y_j. Each graph is raised slightly with respect to the previous one, to present the impression of a 3-dimensional surface viewed from above. To maintain this illusion, all lines which would be invisible to an observer are removed.

This method is especially appropriate when y is the time variable: then each graph represents a "snapshot" at a fixed time, and the full plot represents the time development of the process. An example of a nonlinear wave oscillation, taken from a recent paper by Calogero and Degasperis, (1978) is shown in Figure 4.

In the 3-D plot the width of the plotting area is the x-axis, and again this is subdivided by a grid of spacing h, $x_i = x_o + ih, i = O,1,...,n$. For each

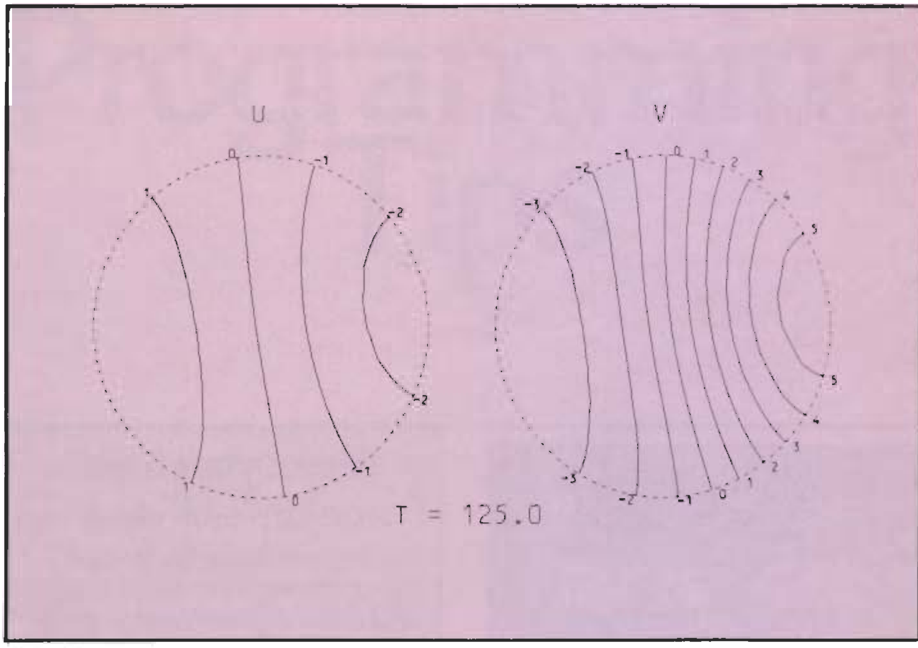

Figure 3
Contour plot of the concentration of two chemicals, u and v, generated by a mathematical model used in embryology.

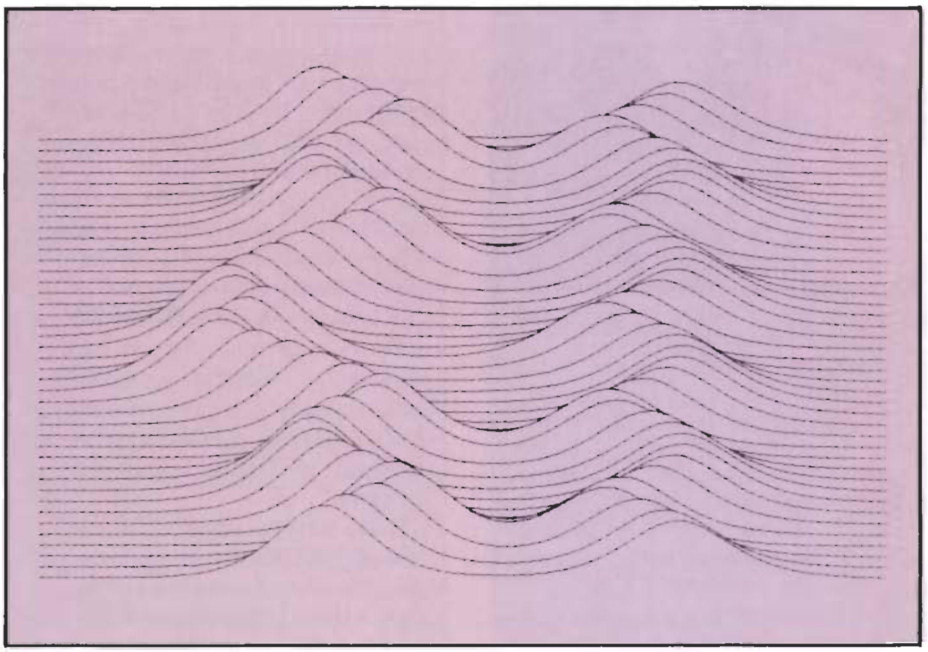

Figure 4
3-D plot of nonlinear wave oscillations

value of y_j we plot $F(x,y_j) = f(x,y_j) * s + kj$ where s is a suitably chosen scaling constant and k is the spacing between the different graphs.

Hidden line removal

The hidden line removal effect is achieved by starting at the front of the plot and proceeding backwards: at each stage a "mask function" m_i is stored, which is the maximum value of $F(x_i,y)$ for all the graphs plotted up to this point. If the current plot is higher than this mask it is "visible" and drawn, the mask being updated to this new value. If the plot is lower than the mask, it is invisible and the pen is moved across in the up position.

In the detailed implementation, we must also consider the possibility of the function being partly visible: in considering the function between two grid points on the x-axis we have four possibilities, depending on the relative signs of $(F_{i,j} - m_i)$ and $(F_{i+1,j} - m_{i+1})$.

These are:

(1)

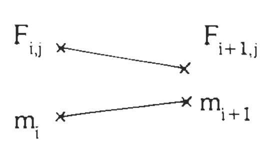

Full line segment visible: connect $F_{i,j}$ and $F_{i+1,j}$ and update m_i, m_{i+1} to $F_{i,j}$, $F_{i+1,j}$.

(2)

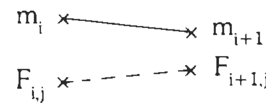

Full line segment invisible: move pen in up position to $F_{i+1,j}$.

(3)

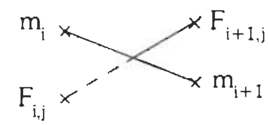

Part line segment visible: calculate and move to intersection point, then draw to $F_{i+1,j}$. Update m_{i+1} to $F_{i+1,j}$.

(4)

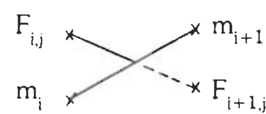

Part visibility, draw to intersection point then move to $F_{i+1,j}$. Update m_i to $F_{i,j}$.

Initially, we set the m_i to be large and negative so that the first graph $F(x,y_o)$ will always be visible. When intersection occurs (cases 3 and 4 above), the coordinates (X,Y) of the intersection point will be given by equations (1) and (2) for $g = F - m$. The calculation gives

$X = x_i + \theta h$, where

$$\theta = \frac{F_{i,j} - m_i}{(F_{i,j} - m_i) - (F_{i+1,j} - m_{i+1})}$$

$Y = (1 - \theta) F_{i,j} + \theta F_{i+1,j}$.

The plotting calculations are simple and fast: only the function m_i need be stored and the function $f_{i,j}$ can be calculated at each point if necessary. Care should be taken in scaling the plotting area and chosing the scale factors so that the top and bottom graphs do not extend outside the plotting area.

Comparisons

In general, the contour plot requires more plotting time since the plotting calculations have to be repeated for each contour value. However, it is much easier to get quantitative information from a contour plot than the 3-D plot.

The effectiveness of each graph will also depend on the shape of the functions in question: the 3-D plot is better perhaps for conveying a visual impression of the form of the function, but in some cases the hidden line feature may mask some interesting behaviour behind a large peak.

A case where the 3-D plot is perhaps the most appropriate is where the variable y represents time. As in all graphics problems, it is the user who must make the final choice, depending on his experience and personal preference.

Both these methods were originally coded in BASIC for the HP 9830A, and I will be glad to make copies of these programs available to any interested user. Only minor modifications to the plotting commands are required to run the programs on other Hewlett-Packard desktop computers with BASIC. Fortran versions are also available if required. K

References

Calogero and Degasperis, (1978), to be published in Springer monograph edited by R.K. Bullough & P.J. Caudrey.

Catalano, G.; Eilbeck, J.C.; Monroy, A.; and Parisi, E. (1980). A mathematical model for pattern formation in biological systems. To be published in Physica D: Nonlinear Phenomena.

Roach (1972), "Computational Fluid Dynamics," (Hermosa Publishers, Albuquerque, New Mexico), page 332.

Chris Eilbeck currently is Lecturer in Numerical Analysis in the Department of Mathematics, Heriot-Watt University, Edinburgh. His main research interests are in the numerical solution of partial differential equations arising in applied mathematics, physics and biology.

Dr. J.C. Eilbeck
Department of Mathematics
Heriot-Watt University
Riccarton, Currie,
Edinburgh EH14 4AS, U.K.

Send your Tips on tape or card

If you are planning to send a Programming Tip to *Keyboard*, please remember to send your Tip already recorded on a tape cartridge or magnetic card. Also, clearly indicate the model number of the computer for which your tip is intended.

By doing this, you reduce the chances that errors may occur in typesetting the Tip, while also saving us the step of keying your program into a computer.

In return, we will send you an additional tape or card for every tape or card you send containing a Programming Tip which we publish. In other words, we will return your original tape or card, plus a new blank tape or card for your desktop computer. Even if your Tip is not published, your original tape or card will be returned.

By sending us your Tip already recorded on a tape cartridge or magnetic card, you will help us publish your Tip sooner. K

Programming Tips

On error statement for 9825A/S

An on error statement can be used to suppress the usual error message when an error occurs, and instead branch to a suitable error recovery routine. The statement often is used in programs which are free from bugs, but for which an error condition can nevertheless arise, as for example, when a peripheral device fails to perform as expected.

Sometimes an error may occur which was not anticipated, and therefore not provided for in the recovery routine. How can one ensure that if this happens, the usual error message is displayed? One simple way is by putting a line of the following type at the end of the error recovery routine:

```
26:  erl-26→X;
  jmp X
```

If an error occurs which does not correspond to any of the anticipated errors, this line is executed and results in a jump to the line in which the error first occured. The on error statement is now no longer in force, and when the same error occurs the second time, the computer stops with a beep and the usual error message is displayed.

This program requires the operator to enter the track and file number from which data is to be loaded. If an unsuitable file is specified, the program branches to "loaderror" and an appropriate message is printed. However, if an error occurs at any line other than 12, or if the error does not correspond to any of the anticipated errors, then the usual error message will be displayed.

When a line such as 26 is used in this way, there are two points which should be noted. The first is that if the program is changed, the line will need to be revised so that the constant value (26 is this case), continues to match the line number. The other point is that if the jmp X statement causes a jump to a line which begins with an on error statement, a continuous loop could form. To avoid this, it is advisable to exclude any other statements from a line containing an on error statement.

For some additional comments on this program, write to *Keyboard* at the address listed on the back page.

by S. Brumby
Department of Chemistry
University of Malaya
Pantai Valley
Kuala Lumpur
Malaysia

Data file editing on the System 45

Editing a data file can become quite a headache, especially when it is a task delegated to data entry technicians who are not overly familiar with the operation of a computer, and are somewhat intimidated by the machine. Here at Electrodyne, we have eliminated the problem through the use of a conversational editing program, whch instructs the operator through all phases of the editing process.

The following program was written for use with the company's System 45, but it is applicable to any machine which supports BASIC software. It can be easily tailored to fit any individual application (e.g., numeric, alphanumeric, differing data formats, etc.) by changing the lines: 230, 350, 590, 900, 920, 940, 1100, and 1220 to suit the particular application.

This program possesses an advantage over programs that modify the file directly, in that the corrected file (or portions thereof) may be examined by the operator before the original data file is replaced.

As the program is presented here, the two arrays D and H require around 20K of memory, limiting its use to those machines with the 201 or 203 Memory Options. However, the amount of memory required may be reduced (or a larger data file edited) by modifying the program to deal with only one portion of the data file at a time, and writing the modified portions out to a scratch file until the entire data file has been edited. The modifications necessary to do this are as follows:

```
171   CREATE "SCRTFL",XX,XX
201   ASSIGN #2 TO "SCRTFL"
202   FOR K=1 TO 2500
203   D(K)=0
204   NEXT K
1220  PRINT #2; D(J)
1231  T=T+1
1232  IF T<3 THEN 202
```

The terminal value of 'T' depends on the number of times the data file which is being edited must be broken up, and may be any constant of a variable.

For a listing of this program, please write to *Keyboard* at the address on the back page and request the data file editing program from this issue.

Warren Harrison
Data Processing Manager
Electrodyne Surveys
P.O. Box 12595
Reno, Nevada 89510 U.S.A.

DMA: the I/O superhighway

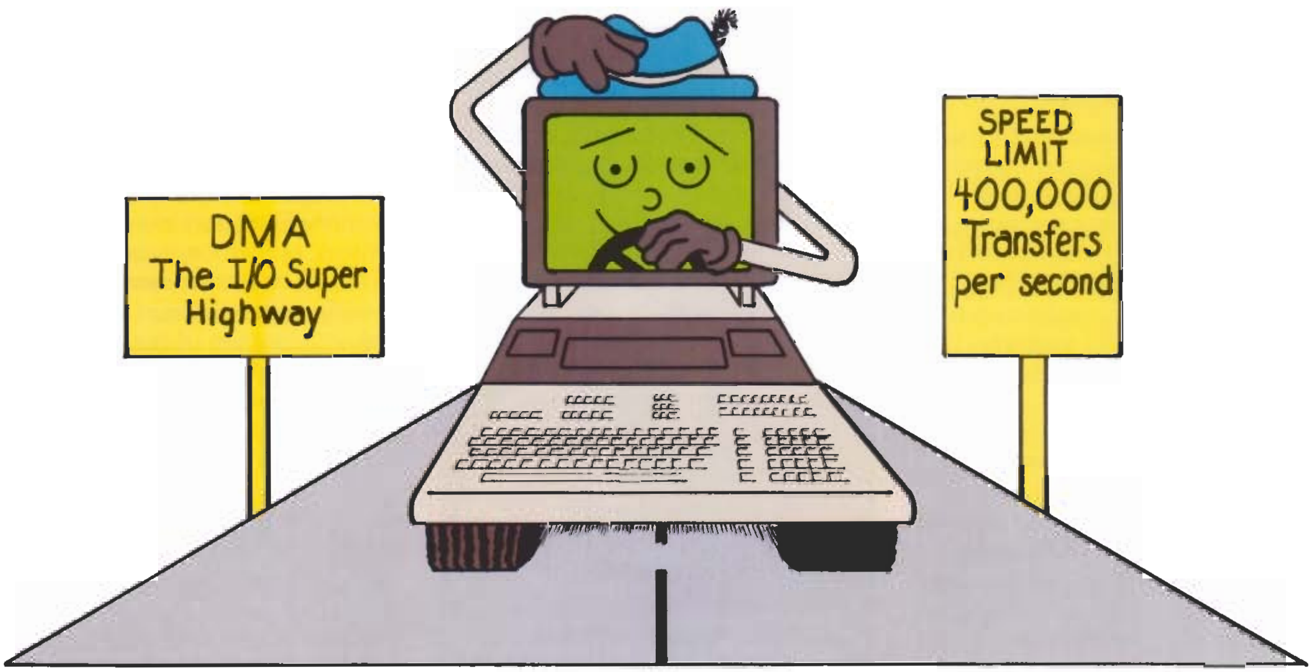

by Steve Leibson
Hewlett-Packard Company
Desktop Computer Division

The articles in this series have described the hardware and circuitry necessary to interface peripheral devices with computers.

All the discussions thus far have assumed that the computer processor is in control of the data transfer process. This is true for many of the devices interfaced. The processor is usually fast enough that the peripheral device determines the data transfer rate.

Some devices, however, are too fast for processor-controlled I/O. These devices are capable of data rates approaching the speed of the computer memory and require a different I/O technique. The technique for interfacing such fast peripherals is called direct memory access (DMA).

In the previous article, we discussed interrupt I/O, which is used for interfacing wth devices so slow that it is very inefficient to have the processor wait for the completion of each I/O transfer.

Instead, the processor initiates a transfer and then continues with other processing. When the peripheral device is ready for the next transaction, it interrupts the processor and reminds it of the previous I/O committment.

The interrupt I/O technique is used as a software transformer to match the slow peripheral with the fast processor. If the peripheral device is faster, the computer processor may only be able to execute the few machine instructions necessary to perform the I/O transfer before the peripheral is ready for another. Here there is a good match between the I/O software and peripheral speeds, and programmed I/O is sufficient for the task.

Speedy peripherals

Ultimately, there is a class of peripherals too fast for even the few instructions needed to perform programmed I/O. As long as these peripherals are not faster than the computer's basic memory cycle, there should be a method for performing the required I/O. There is and it is called direct memory access (DMA).

In order to discuss DMA and how it works, we must return to the model of the processor-memory-I/O system discussed in the first article in this series. Recall that the processor is linked to the memory via a set of lines called a memory bus and to the I/O interfaces via an I/O bus.

Both busses require the processor to generate address signals and control signals to synchronize the flow of data over these busses. Generally, I/O consists of taking information from the interfaces through the I/O bus and transmitting this information to the memory using the memory bus or vice versa.

Inefficient throughput

During this transfer the processor is also using the memory and memory bus to supply machine instructions so that it knows how to effect the data

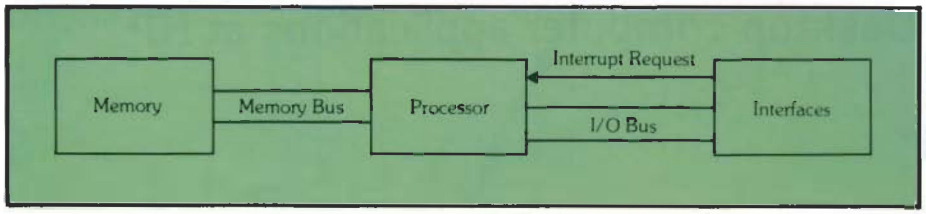

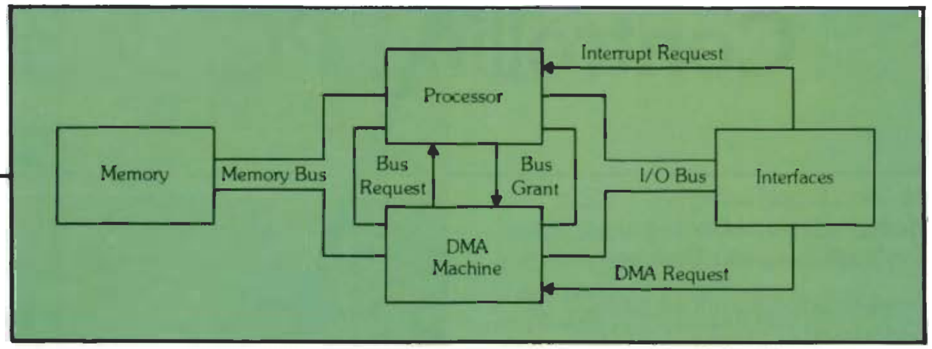

Diagrams above illustrate the differences between a system that does not include a direct memory access machine, top, and one that does, bottom.

transactions. If we assume that it takes only nine machine instructions to perform one data transaction, we can see that the effective I/O throughput is only 10% of the rate that the memoy could support.

That is, for every 10 memory cycles, nine are used to instruct the processor and only one is used to place data for I/O. Only very simple data transactions can be performed with nine machine instructions. If formatting or code conversions are necessary, many more instructions are needed.

Bypassing the processor

The only way to speed up the I/O process is to eliminate the slowest link in the data path. For high-speed peripherals, the slowest link is clearly the processor itself! How can we eliminate the processor when that is the component that links the I/O and memory busses and is required for the generation of the signals that actually make these busses work?

The answer is to build a specialized circuit that is designed to transfer data at the full memory speed. Because the only function this circuitry must perform is this transfer, the capability may be wired into the circuit and instructions from memory are not needed and do not reduce the effective memory bandwidth.

If we place this specialized circuitry so that it, too, bridges the I/O and memory busses and if we also give it the capability of generating the address and control signals required by these busses, then we have a machine that is capable of performing I/O at the full memory speed. This specialized circuitry is called a direct memory access or DMA machine. All that remains is to select which device will have control of the busses, the processor or the DMA machine.

Controlling DMA

Normally, the processor will have control of the busses because the DMA I/O must be infrequent enough to allow at least some processing to be done. It is therefore necessary for the DMA machine to acquire bus control from the processor whenever necessary.

The processor can enable the DMA machine to request bus control, but it is the interface that must actually request service through the DMA machine. Only the interface knows when the attached peripheral requires DMA service. Thus we must add some connecting signals between the interface and the DMA machine, and between the DMA machine and the processor.

DMA handshake

The interface must have some means of requesting service from the DMA machine. A signal called DMA Request (DMAR), added to the collection of signal lines on our I/O bus, will be sufficient. Upon receipt of this request, the DMA machine must request bus control from the processor.

The processor may decide that the time of the request is inopportune and wish to hold off the transfer of control temporarily — This is a job for the everpresent handshake!

We will create two handshake lines called Bus Request and Bus Grant. The DMA machine will ask for bus control with Bus Request and wait to actually take control until it receives a signal on Bus Grant. Thus the processor can maintain control of the memory and address busses as long as required.

Burst and cycle-steal

The DMA that we have been discussing is called burst DMA because data transfer is done in a burst where the DMA machine totally controls the I/O with the full speed of the memory bus at the expense of completely halting any processor activity.

If half the memory bus bandwidth is sufficient to solve the high speed I/O problem, another type of DMA can be employed. Called cycle-steal DMA, the DMA machine alternates control of the busses with the processor, each unit using every other memoy cycle. Cycle-steal DMA allows the processor to operate at 50% efficiency while still providing relatively high speed I/O.

At this point in the I/O series, we have discussed the basic hardware needed for interfacing computers to peripherals. We have covered the four basic interfaces: Parallel, BCD, HP-IB, and Serial and we have discussed specialized I/O; interrupt and DMA. Now that we have our devices talking, we will discuss how to overcome the language barrier. Next issue: character codes. K

Controlling PC electroplating

by John Monahan
Hewlett-Packard Company
Desktop Computer Division

A desktop computer capable of performing closed-loop process control has increased production of printed circuit (PC) board panels 80% for Hewlett-Packard's Loveland Instrument Division.

The HP 9825 Desktop Computer is controlling the process of electroplating printed circuit boards manufactured on panels. Each panel contains from one to several hundred PC boards.

The project is directed by Eugene Dick, associate engineer for the Loveland division's Engineering Group, who says that the system has increased production to 1800 panels per day from 1000.

The PC boards are used in equipment produced by nine of HP's manufacturing divisions. Consequently, Dick says, there is an overriding need for quality as well as quantity.

Copper, nickel and gold

Each PC board when finished is plated with five micro-inches of gold that covers an initial layer of copper and a second layer of nickel.

The gold, Dick says, provides superior durability and reduced contact resistance, while nickel guarantees hardness and brightness.

Each PC board receives its chemical treatment on one of eight panels attached to a mobile rack. The computer monitors some of the electroplating process, and directs a controller which guides the movement of racks through a series of 53 stations or tanks.

Here they are washed, rinsed, treated wih acids, and electroplated in solutions of copper, nickel and gold. At no time, except when placed on the rack, is human assistance required, because automated cranes move the racks through the U-shaped chain of stations.

Worker loads PC boards into rack for automated plating.

Brains and brawn

This mechanical process is controlled by an Allen-Bradley PLC processor, which receives its instructions from the 9825 via an HP 98032A 16-bit Parallel Interface.

The PLC processor requests instructions through an interrupt and sends a 16-bit word regarding the crane's location and identificatin. The 9825 then issues a four-digit BCD command that includes the crane's identification number and a three-digit command code.

The PLC processor initiates and monitors each command from the computer. When the order is completed, the next crane is activated.

Dick says this separate-controller process frees the computer for computational tasks, and permits each machine to be used for its special purposes. Thus the PLC processor is the brawn, the 9825 the brain.

Traffic control

The computer keeps the racks moving properly, using a subroutine that prevents them from piling up on one another. It also computes how long they have been in the solutions and warns of amperage aberrations in the gold solution.

Each panel is submerged in the copper solution for 60 minutes, in the nickel 29 minutes, and in the gold 35 seconds. These times are controlled by the computer using an HP 98035A Real Time Clock and counted down on the screen of the HP 1350 Display System, interfaced to the 9825 using HP-IB.

The display also depicts the relative positions of the cranes, panels and immersion tanks. An error information block appears at the bottom of the screen. Its purpose is to alert the operator to stoppages or other problems by flashing a message denoting the error.

Also depicted are the time, date, process in use, and file being used. Disabled tanks are shown with an X, and should a tank become inoperable, the computer prohibits new racks from being placed into it.

Guarding the gold

Gold plating, the most expensive of the electroplating processes, is monitored by the 9825 using an HP 9878 I/O Expander and Allen-Bradley multiplex system. Should electrical current levels in the gold solution or alkaline cleaner become too weak or quit, a message appears on the 9825-controlled graphics display system, along with a flashing indicator ("LOW," or "NC" for not current). A diamond-shaped graphic appears around the malfunctioning station, and an audible alarm and red light are turned on.

Redundancy

The overall 9825-controlled system is backed up by a duplicate system near Dick's desk, plus other measures written into the computer's software.

These include automatic dumping, every 230 seconds, of essential data

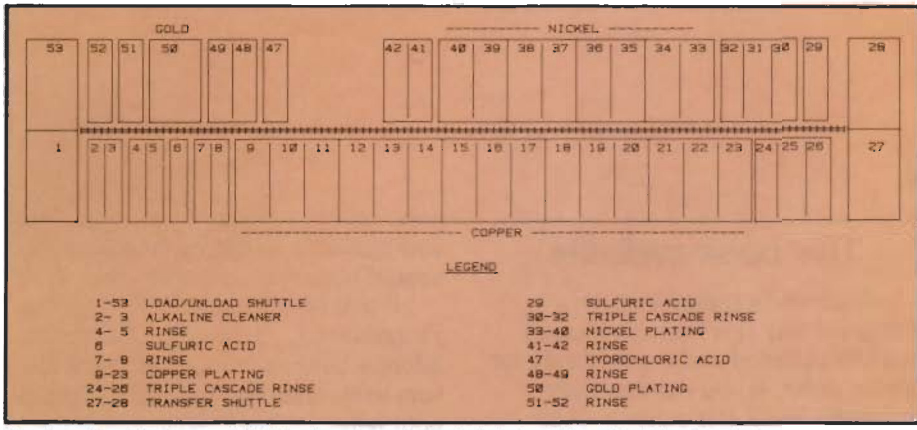

The Loveland Instrument Division's automated printed circuit board plating line includes the stations, or baths, shown above.

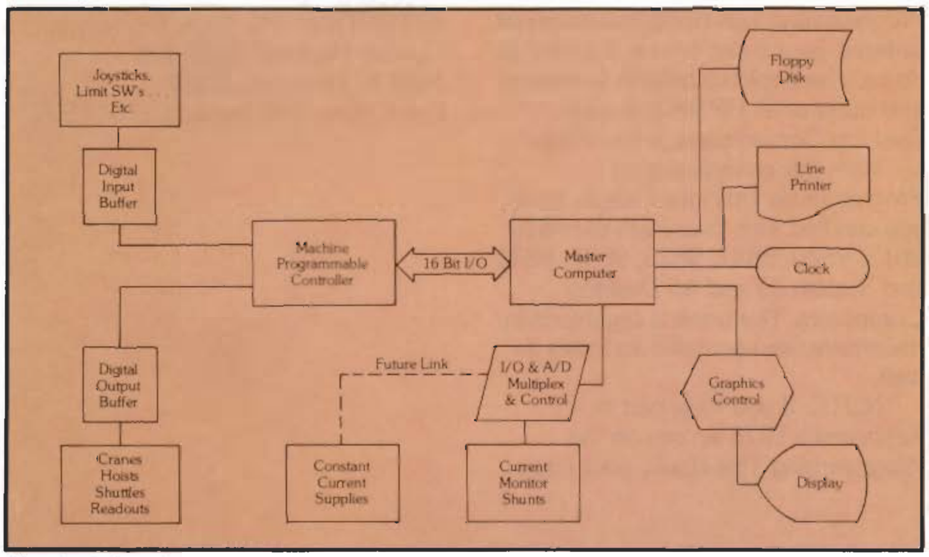

Flowchart shows interrelationships in the system which controls the plating line.

into the HP 9885M Flexible Disc Drive. If a power failure should occur, this procedure allows the system to be brought back on line where it was before power was lost.

Should the program itself be stopped, computing of the panels' elapsed immersion time ceases automatically. A stoppage could occur, for instance, if an engineer were troubleshooting board failures or an operator were retrieving a panel that had slipped from the rack. There would be, says Dick, sufficient time to stop the line, reduce the electric current levels in the plating solution, then resume that current and timing without ruining the panels.

Moreover, a pause function, initiated by the operator, dumps all data into the flexible disc and stops the program. The line is then backed up to its condition at the moment the pause was initiated.

Versus paper tape

The desktop computer, Dick says, can handle line stoppages and the resultant changes in electrical current levels more precisely and more quickly than manual methods. For this reason the Loveland plant replaced its former control system — punched tape — with the computer.

With the tape system, Dick says, it was necessary to create new tapes when a breakdown occured at a station. The operator was forced to select the right tape, then resynchronize the tape with the cranes. The line would be down an undesirable amount of time, and often an empty rack found its way through the stations.

Now the computer decides which station should receive the next rack. If a station fails, it can be neutralized dynamically by the operator and bypassed automatically by the computer.

Reliability

During the last year, Dick's 9825 has failed only twice. One failure was a simple keyboard problem and the other involved a memory board. Dick notes that the computer resides in a hostile environment — sulfuric and hydrochloric acids are used in the electroplating process — and that the system is operated six days a week, 24 hours a day.

Simplified operation

A schedule like that dictates that Dick cannot be present every moment the system is working. As a result, he devised a menu — consisting of seven items — which appears on the single-line display screen to guide other operators through the system.

Printed by an HP 9866 Thermal Printer, the instructions available from the menu include: general directions, disabling/enabling data for the stations, system stop control, system restoration after power failure, use of pause-restore, and electrical current information. The seventh instruction lists his and other engineers' names and home phone numbers.

HP 9825s are being used for two other related tasks. One is controlling the hoist on the electroless plating line, again reducing downtime and errors compared to the tape reader that was previously used.

The chemistry laboratory at the Loveland installation uses a third 9825 to log temperature, pH, and metallic ion concentration in the solution tanks, saving hours of manual labor. K

Update

New 9825 keyboard

New 9825A/S Desktop Computers are being built with an enhanced, typewriter-like keyboard.

The raised keypads result in an easier-to-use and more responsive keyboard that facilitates faster and more comfortable alphanumeric input.

Upgrade kits are available to convert existing 9825s to use the new keyboard. Special Function Key overlays are available to fit the new keyboard.

This keyboard uses the same quality keypad design as keyboards supplied on System 35 Desktop Computers. Part number 09825-67960 is the current color, ASCII standard keyboard; 09825-67961 is the original color. For optional Katakana keyboards, the part numbers are 09825-67962 in the current color; 09825-67963 in the original color.

Tips book available

Keyboard's compilation of Programming Tips from past issues has been printed and is now ready for you to order. If you have already ordered a copy, it is either in your possession now or on its way to you.

If you want to receive the Programming Tips Book, but have not ordered your copy, now is the time to do so. The book is available to owners and users of all HP 9800 Series Desktop Commputers at no charge.

With this compilation of Programming Tips into a single book, you can find Tips to use wih the 9810, 9815, 9820, 9821, 9825, 9830, 9831 and System 35 and 45 Desktop Computers. The book is organized by mainframe and contains an index as well.

NOTE: If you took part in *Keyboard*'s 1978 survey on the Programming Tips Book, your copy was mailed in February, and you should have received it by now.

If you wish to order the Programming Tips Book, write to the address below for your free copy. Be sure to include your mailing address in your letter, or, preferably, attach the mailing label from your latest copy of *Keyboard*.

ATTENTION: Ms. Chris Stumbough
Hewlett-Packard/*Keyboard*
3404 E. Harmony Road
Fort Collins, Colorado 80525 U.S.A.

Keyboard

Hewlett-Packard, 3404 E. Harmony Road
Fort Collins, Colorado U.S.A. 80525
Phone: 303/226-3800

Publications Manager - Ed Bride
Editor - Bill Sharp
Feature Writer - John Monahan
Editorial Assistant - Ms. Chris Stumbough
Art Director - Hal Andersen
Typography - Ms. Paula Dennee

For further information on HP products or applications, please contact your local Hewlett-Packard Sales and Service Office or write to Keyboard

CHANGE OF ADDRESS: To change your address or delete your name from our mailing list please send us your old address. Send changes to Hewlett-Packard Keyboard, 3404 E. Harmony Road, Fort Collins, Colorado 80525, U.S.A. For Europe, send changes to Hewlett-Packard N.V., Post Box 529, 1180 AM Amstelveen, Holland.

Mar 1, 1980