HEWLETT-PACKARD JOURNAL

HEWLETTPACKARDJOURNAL

Technical Information from the Laboratories of Hewlett-Packard Company

MARCH 1981 Volume 32 • Number 3

Contents:

New Display Station Offers Multiple Screen Windows and Dual Data Communications Ports, by Gary C. Staas It's four terminals in one, bringing new flexibility to computer application systems.

Display Station's User Interface Is Designed for Increased Productivity, by Gordon C. **Graham** Easy access to an extensive feature set requires a thorough, thoughtful approach to the user interface.

Hardware and Firmware Support for Four Virtual Terminals in One Display Station, by Srinivas Sukumar and John D. Wiese The goals were 2645A compatibility, improved pricel performance and reliability, and ease of use, manufacturing, and service.

A Silicon-on-Sapphire Integrated Video Controller, by Jean-Claude Roy Integration was considered mandatory to make the display system practical and reliable.

SC-Cut Quartz Oscillator Offers Improved Performance, by J. Robert Burgoon and Robert L. Wilson It's more stable and less noisy, warms up faster, and uses less power.

The SC Cut, a Brief Summary, by Charles A. Adams and John A. Kusters First introduced in 1974, the stress compensated cut has many virtues.

Flexible Circuit Packaging of a Crystal Oscillator, by James H. Steinmetz Selectively stiffened flexible circuitry is a radical approach that meets tough objectives.

New Temperature Probe Locates Circuit Hot Spots, by Marvin F. Estes and Donald Zimmer, Jr. It provides fast, accurate temperature measurements for design, diagnostic, and testing applications.

In this Issue:



The shiny rock on this month's cover is a piece of cultured (laboratory-grown) quartz. The thin transparent disc mounted in its holder in front of the rock is a quartz crystal of the type used for frequency control and timing in many electronic devices, including the quartz wrist-watches that some of us wear (not all crystals are as large as this one). When these thin slices of quartz are subjected to an alternating voltage, they vibrate, and they vibrate much more strongly at one frequency than at any other. It's this property that makes them useful as frequency and time references.

Quartz is a crystalline material, which means that its atoms line up in a regular pattern or lattice. Thin crystal discs for frequency control are taken from large quartz rocks by slicing the quartz at specific angles to the crystal lattice. Different angles produce different sets of desirable properties. The new HP crystal oscillator described in the article on page 20 derives its frequency stability from a crystal that has been cut to make it relatively insensitive to temperature variations. This property and some state-of-the-art circuit design give the new oscillator. Model 10811A/B, better stability, lower power consumption, and faster warmup than earlier HP crystal oscillators. Model 10811A/B is designed to serve as a highly stable frequency or time reference in precision laboratory instruments, especially those that have to operate for long periods without adjustment.

Pages 3 through 19 of this issue describe a new computer terminal, Model 2626A Display Station, that meets a need of many computer users for a terminal that can handle relatively complex operations. Internally the 2626A can be set up to function as up to four separate "virtual" terminals. The operator can see on the screen what's happening in one, two, three, or all four of these virtual terminals at the same time, and can affect what's happening in one virtual terminal at a time, using the keyboard. Any two virtual terminals can communicate with different computers or the same computer at the same time. This kind of flexibility opens up many new possibilities for computer application systems.

Wrapping up the issue is an article about the 10023A Temperature Probe, a simple device that helps circuit designers find and eliminate circuit hot spots that may indicate problems or likely failure sites.

-R. P. Dolan

New Display Station Offers Multiple Screen Windows and Dual Data Communications Ports

This versatile computer terminal can act like four virtual terminals. It's designed for data entry and program development.

by Gary C. Staas

EW DATA ENTRY and program development capabilities are provided by a new HP CRT terminal, Model 2626A Display Station, which lets the user display, compare, and combine data from two different computers and four different memory workspaces simultaneously. The new terminal (Fig. 1) has dual data communications ports for connection to computers or peripherals and the user can divide both the display memory and the display screen into as many as four independent work areas. Other features are line widths up to 160 characters with horizontal scrolling, screen-labeled softkeys, an optional built-in thermal forms-copy printer, programmable tones for audio cues, and keys for interactive forms design.

Design Objectives

The objectives that guided the design of the 2626A Dis-

play Station are reflected in many of the characteristics of the terminal. An important characteristic is compatibility with the HP 2645A Terminal. Terminal drivers and application programs that work with the 2645A will also work with the 2626A, thereby protecting users' software investments.

The new terminal is also reliable and is easy to build, check out, and service. The reliability goal was 8,000 hours MTBF (mean time between failures). An extensive set of selftests is built into the terminal, and some of these tests can isolate failures to the component level. One test, which repeatedly executes most of the other terminal tests, is used after terminal assembly to spot failures in the factory. To make the terminal easy to build, there are very few options and the 2626A has many components in common with other terminals in the 262X line. The factory does very little configuration of the terminal. Since that is an easy process,



Fig. 1. Model 2626A Display Station is a multi-workspace, multiwindow computer terminal that has dual data communications ports. Its capabilities may be configured dynamically as four logically independent virtual terminals. The terminal can handle line lengths up to 160 characters and offers foreign language options and an optional built-in printer. it is almost totally left to the user.

User sophistication varies considerably, from the data entry clerk to the very knowledgeable OEM customer. Each user needs to take advantage of a different set of features. A department manager of a data entry operation, for example, needs to configure the terminal for use by data entry clerks. The terminal provides configuration menus for this purpose. On the other hand, these configuration menus can be locked out and made unavailable to the clerks to avoid confusing them with details they don't need.

Among many new concepts embodied in the terminal is the ability to handle more than one job at a time. Like several sheets of paper on a desk, each relating to a different task, the 2626A Display Station splits its display screen into as many as four windows, each with independent data. The dual datacomm ports allow these windows to communicate with more than one computer program at once.

The optional built-in thermal printer avoids the problems of expensive, distant, large, and noisy impact printers. It was considered essential to be able to print whatever was on the screen, such as forms and special character sets, and the integral printer makes this possible.

To allow a user to draw simple bar charts without graphics capability, one of the terminal's character sets has been expanded to include appropriate characters. The forms-drawing keys make it easy to design data entry forms.

Foreign language support is an important goal in the international marketplace. The 2626A provides six European languages, including mute and overstrike characters.

The 2626A uses HP's silicon-on-sapphire (SOS) largescale integrated circuit process. This and other design features make it possible to offer a terminal with a much improved and expanded feature set for a cost comparable to an HP 2645A.

Softkeys

To allow easy access to terminal functions, the 2626A has eight softkeys that do not have fixed functions. A two-rowhigh label at the bottom of the screen just above each softkey indicates its current function. A function key that indicates a terminal mode, such as REMOTE/LOCAL, has an asterisk on its label when the corresponding function is on. The softkeys can be locked so the terminal stays within a group



Fig. 2. In this example, portions of three 2626A workspaces are displayed in windows on the screen. Workspace 4 is not displayed.

of softkey levels having similar functions.

Workspaces and Windows

When the terminal is powered on, it partitions memory into displayable lines of equal sizes. The user can select a length from 80 to 160 characters per line. Lines 132 characters wide, for instance, are useful for holding data to be sent



Fig. 3. (a) In this example, the terminal has four workspaces, each holding a form. Workspace 1 is displayed in the keyboard window and the user can type data into form 1. Workspaces 2, 3, and 4 are not displayed in a screen window. The host computer is connected to workspace/window 1. (b) After the user presses the **ENTER** key the host displays workspace 2 in the keyboard window and workspace 1 is no longer displayed. The host remains attached to workspace 1 and receives the data just entered while the user fills form 2 in workspace 2.

to a computer line printer. The user can group these lines into workspaces, each with a fixed number of lines. For example, the user could set up the terminal to have two workspaces of 80-character lines, workspace 1 with 40 lines and workspace 2 with 70 lines. Up to four workspaces are allowed. Multiple workspaces are useful for doing several different jobs on the terminal at the same time.

All workspaces can be displayed on the screen simultaneously to give an overall view of what the terminal is doing. The screen can be divided into display windows that show all or part of the workspaces. To continue the above example, workspace 1 might be viewed in a display window from the first through the tenth screen rows, and workspace 2 might be seen in a window from screen rows eleven through twenty-four. Dotted lines separate windows on the screen to avoid confusion. The terminal also has a vertical border on the screen to allow left and right windows. A workspace need not be displayed to be functional. A workspace that is closed (i.e., has no display window) can receive and send data. Fig. 2 shows an example of a workspace/window configuration.

To avoid ambiguity, all keyboard input affects only one workspace at a time, the one in which the cursor resides. Typed data appears only in this keyboard workspace/ window, not in any other window. Local editing keys, such as delete character, take effect only in this window, and so on. The user can change the keyboard window by means of a softkey.

Each workspace operates independently, so that workspaces can be used to perform different tasks with a computer. Most operating modes of the terminal are or can be on a workspace basis. Some of the modes are always associated with a workspace, such as FORMAT mode, which allows a window to be used as a form, limiting the user to typing data into unprotected fields. Other modes are grouped into a terminal configuration, which may be freely attached to workspaces. Suppose that in terminal configuration 1, REMOTE mode is on, and in terminal configuration 2, REMOTE mode is off. If terminal configuration 1 is used for both workspaces, 1 and 2, then both will have REMOTE on. If workspace 1 uses terminal configuration 1 and workspace 2 uses terminal configuration 2, then workspace 1 has REMOTE on and workspace 2 has REMOTE off. There are four terminal configurations so that each workspace can have its own configuration. This is desirable when workspaces are to be used for different kinds of jobs.

The workspace/window configuration can be set locally on a menu or remotely by escape sequences sent from the host computer.

A data communications port can be attached to any one of the workspaces at a given time. Data flow is to and from that workspace only; other workspaces are unaffected. The data entry example in Fig. 3 illustrates how this feature is used. A data entry application program in the host computer first creates four workspaces in the 2626A terminal and attaches its host's data communications port to each workspace in turn to transmit a form to it. After it sends all the forms, it displays the workspace that has the first form and makes it the keyboard workspace so the user can begin entering data into the form. When the user presses the **ENTER** key, the program displays the second workspace with the next form and attaches the keyboard to it. The host program then attaches itself to the first workspace to receive the data just entered on the form. The user continues entering data into the second form while this transfer is in progress. This concurrency results in greater throughput and operator productivity.

To fine-tune a window configuration, the user can move the horizontal and vertical borders with softkeys. Another softkey moves the cursor to the next screen window, and another displays the next workspace that is not currently displayed on the screen. Whenever the line length is greater than the screen width of a window, the data may be scrolled



Fig. 4. Uses of the dual data communications ports. (a) Port 1 is attached to a computer. Port 2, when not in use as a datacomm port, can be attached to a serial printer or other RS-232-C device. (b) Each port is attached to a different host computer. (c) Both ports are attached to the same computer.

horizontally. If the vertical border were at screen column 30, for example, a right window would have only 50 columns displayed and horizontal scrolling would be necessary to view all the data.

Dual Data Communications Ports

To further enhance its multitasking capability, the 2626A has two ports for data communication. When a port is attached to a workspace, it means that data received from that port is processed and/or displayed in that workspace. If the user enters data in that workspace (by pressing **ENTER** while that workspace is the keyboard workspace/window, for example) the terminal sends data through that port. A port, workspace and terminal configuration together make up a virtual terminal entity.

Fig. 4 shows different possibilities for using both ports. The two ports can be cabled to the same host or to two different computers because each port can use a different data communications protocol. Fig. 5 shows an example using the two ports attached to two different computers. Both workspaces that are attached to ports can be displayed on the screen at the same time in two windows. To make the terminal truly useful for handling multiple jobs, both ports/workspaces can be receiving and sending data and a user can be entering data in a third workspace simultaneously. A useful single-host application is to use one port and workspace for a console and another as a user terminal, or a programmer might examine the output of a compiler for errors in one workspace and edit the source file in another workspace.

Soft Configuration

The 2626A represents an advance over previous terminals in its flexibility and ease of configuration. There are no hardware configuration switches inside the terminal. A user configures the terminal by choosing values on menus built into the terminal. There are configuration forms for workspace/windows, data communications for each port,



Fig. 5. In the top window/workspace, a user can examine a data base listing from a computer attached to port 1, and in the bottom window/workspace can get information needed for a data entry program running on a second computer attached to port 2.



Fig. 6. New line drawing characters make it easy to draw high-quality bar charts.

terminal modes for each workspace, and global terminal items. A host computer program can also set any configuration item on any menu. The terminal saves configuration data in battery-powered RAM when ac line power is off. All configurations can be locked to prevent local alteration in applications involving less sophisticated users. If the optional thermal printer is present or if an external printer is attached to port 2, the user can print the configuration menu for further backup.

The 2626A provides both point-to-point and multipoint data communications as standard features. Data communication configuration is done by means of six general menus that cover most full and half-duplex configurations as well as multipoint (multidrop bisynchronous) communications.

Device Control

The 2626A can copy data from one workspace to one or more other workspaces. If the receiving workspace is in REMOTE mode, it will transmit the data just as if it were typed into that workspace. This allows data transfers from one host to another. An external printer can be connected to the second data communications port, which is a standard RS-232-C port when it is not used as a data communications port.

One of the 2626A options is a thermal printer integral to the terminal. The printer is capable of printing all of the terminal's character sets to give exact screen copies. It also provides normal, expanded (40 characters/line), and compressed (132 characters/line) printing modes. When RE-PORT mode is on, the printer provides three blank lines at the top of each page, 60 lines of text, and three blank bottom lines to format standard $8\frac{1}{2} \times 11$ -inch sheets.

A data logging mode enables the terminal to record all the data it receives on a printer. When LOG TOP mode is on, the terminal prints lines as they disappear from the top of a workspace to the destination printer(s) when all display lines for the workspace are used up and the data would be otherwise lost. LOG BOTTOM causes the terminal to log lines as they are received.

Ordinarily, one workspace at a time is copied to the destination device(s). SCREEN COPY copies the screen to the printer(s) exactly as it appears, so that all windows displayed are copied.

User-Definable Softkeys

When they are not being used for terminal configuration or control, the 2626A's softkeys can be used to perform user-defined functions. The user-defined functions are enabled by pressing the **USER KEYS** key. Each softkey may be given a definition consisting of as many as 80 displayable characters. Subsequently, pressing a softkey has the same effect as pressing all the keys in its definition. Two eightcharacter rows on the screen label each softkey; the labels are defined by the user. Video and character set enhancements may be embedded directly in the softkey labels and the 80-displayable-character key definitions. This allows eye-catching labels and more information in the definitions. The **ENTER** and **RETURN** keys are also softkeys. Another key sets the default values of the user softkeys.

Sketch Forms Facility

The HP 264X family of terminals has a line drawing character set to represent forms. The sketch forms facility of the 2626A makes it much easier to use this set by providing a set of keys to draw horizontal and vertical lines. The drawn lines do not overwrite text or fields. When horizontal and vertical lines cross, the correct intersection character is automatically selected. The user initially selects through softkeys the line type desired (single, double or bold) and the video enhancement desired for the line. Another set of keys automatically draws a box around either the margins (left, right, top and bottom) or the displayed data.

Video Enhancements

Video enhancements and character sets were previously selected only by escape sequence. The 2626A user can also select them with softkeys. The line-drawing, math, and large-character sets have all been expanded to 96 characters. With the new line-drawing characters, for instance, the user can create bar charts as shown in Fig. 6. If a portion of the workspace has the new security enhancement, it displays as blanks, no matter what data is there. This enhancement is useful for password fields and for other secure fields. For users who prefer black-on-white lettering, the black background can be switched to inverse video.

Foreign Languages

The 2626A has ISO character sets for the following European languages: Swedish/Finnish, Danish/Norwegian, French, German, United Kingdom, and Spanish. Keyboards are available for all these languages. The terminal is set to operate for a given language by making a selection on the global configuration menu. All the language capability is built into the terminal; a user needs only the appropriate keyboard with the special character keycaps and key placements and the extra character set ROM.

Fig. 7 illustrates the new language option characters. Some of the languages provide a mute and overstrike capability. For example, when the terminal is configured for

		DECIMAL VALUE									
LANGUAGE	35	64	91	92	93	94	96	123	124	125	126
USASCII			C	1	1	^		{	1	}	~
Swedish/Finnish		É	Ä	ö	A	ū	é	ä	ö	â	ü
Danish/Norwegian			Æ	Ø	A		×	æ	ø	å	~
French	£	à	•	ç	ş	•		é	ù	è	
German	£	ş	Ä	ŏ	ü	•	×	à	ö	ü	ß
United Kingdom	£		۵	1	3	•	•	{	1	}	~
Spanish			i	Ñ	5	•		(ñ	}	~

Fig. 7. Special characters, including overstrike characters, are available.

French with mutes enabled, typing a circumflex displays it without moving the cursor. When the letter a is typed, the character is replaced by the letter a with a circumflex over it. The terminal provides an eight-bit mode for the HP 300A in which characters are shifted into the alternate character set by setting the eighth bit. When the foreign characters or Roman extension set is selected as the alternate, the terminal can operate with an HP 300A doing foreign processing.

Terminal Tests and Error Messages

The 2626A offers a powerful set of internal tests. Whenever the terminal is powered on, it performs a self-test to give the user an immediate indication of terminal malfunction. A more comprehensive test is used by production during the burn-in time of the terminal. This test logs any errors found into the battery-powered RAM. The data communications test facility is menu-driven and allows a variety of tests on each port. The integral printer test prints a test pattern on the thermal printer. Each ROM in the terminal has identification information that can be displayed on the screen; this is useful to customer engineers in determining the exact version of firmware in the terminal. All of these tests can be invoked by softkeys and do not require downloading a diagnostic into terminal memory. A customer can perform a test and pass on the results to a customer engineer over the telephone, making service calls less frequent and more efficient.

All of the self-tests produce error messages indicating any problems found. Some user operations can also produce errors, such as certain illegal settings on the configuration menus. If a user attempts to set such a configuration, the terminal displays an explanatory message. All error messages appear on the bottom two lines of the screen replacing the softkey labels until the message is cleared. The rest of the screen remains visible so the user can more easily determine the error.

Additional Features

Among the 2626A's features is a programmable bell with 15 tones, 16 durations, and two volume levels. Several new kinds of terminal status are also offered. For example, window status shows the current workspace/window config-



Gary C. Staas

Gary Staas was born in Dayton, Ohio and attended the University of Dayton, graduating in 1969 with a BS degree in mathematics. In 1971 he received his MS degree in statistics from Stanford University. After three years developing software for hospital information systems, he joined HP's Data Terminals Division, where his responsibilities have included QA of various terminals, data entry system design, and firmware development and feature design for the 2626A and other terminals. Gary is a member of the ACM and a resident of Scotts Valley, California. He lists his

interests as skiing, eastern philosophy, running, backpacking, public speaking, and humor. Recently he left HP to pursue his career with another computer firm.

uration, and terminal ID identifies the terminal as a 2626A.

Modify modes allow a user to modify a line on the screen and retransmit it to the host. When the TAB = SPACES feature is on, pressing the **TAB** key transmits the appropriate number of spaces to take the cursor to the next tab stop rather than sending a tab character. **BACK TAB** sends the appropriate number of backspaces to the previous tab stop. This makes it easy to create text files for applications that do not understand tab characters.

Acknowledgments

I would like to acknowledge the HP 2626A project manager, Prem Kapoor, for his many helpful suggestions on features and continual striving for excellence in the product. Stan Telson contributed much to the self-test features and keyboard hardware. Products of this complexity cannot be realized without the active support of lab and marketing management. I would like to thank Lance Mills, lab manager, Tom Anderson, marketing manager, and Terry Eastham and Ken Blackford, product managers, for their encouragement and support.

Display Station's User Interface Is Designed For Increased Productivity

by Gordon C. Graham

PRIMARY CONTRIBUTION of the HP 2626A Display Station is its advanced user interface, which allows quick access to the many features of the terminal. An important design goal in this project was to make the terminal easy to use. Because this terminal has more features than previous alphanumeric terminals developed by HP, easy access required a thorough, thoughtful approach to the user interface.

Elements of the User Interface

All user interfaces have four elements in common: User Model. This is the mental model that the user of a product forms as to how the equipment functions. If the product has been carefully designed, the user automatically develops a good understanding, or model, of how it works, and will use the product's features naturally instead of resorting to the operator's manual every time a new situation arises. This area is a very important part of the total design. Each new feature and every change to an existing feature must be carefully evaluated to ensure that it is completely consistent with the operation of all other features, and that its use is natural to the operator.

Command Language. The command language is closely related to the user model and is in fact a concrete representation of it. It is by means of the command language that the user accesses the features of a product. The commands should all relate to each other in a consistent, systematic way, and should include provisions for aborting commands and handling errors.

Feedback from the Instrument. Feedback from the instrument in response to commands tells the user about the completion status of any command. Proper design of feedback mechanisms inspires confidence that the instrument is being used properly and is performing the desired functions. Feedback is also used to point out errors to the user as soon as they are detected. **Information Presented to User.** This is closely related to feedback. The central issue is how to display information in a manner that promotes the most effective interaction between the user and the instrument. Both visual and audible means can be used. Generally, the form that visual information takes must be tailored to the characteristics of the display.

Application of these guidelines to the 2626A is mainly reflected in two areas: screen-labeled softkeys and menudriven configuration.

Screen-Labeled Softkey

Earlier terminals provided access to their features by means of function keys on the keyboard. These keys had a fixed task associated with them, such as home the cursor or clear the display. This approach was satisfactory as long as the function set was small. Later terminals, with many more functions, required the user to type in escape sequences (the ESC key followed by one or more other keystrokes) to exercise many of the functions. These escape sequences were designed for use by a host computer system driving the terminal and so were designed for compactness rather than ease of use. The result was that many of the advanced features of these terminals were used only by sophisticated users willing to memorize or look up the escape sequences. To solve these problems, more recent terminals have introduced the concept of tree-structured, screen-labeled softkeys. The 2626A has extended this approach.

Softkeys are keys that do not have a dedicated function, but instead perform many different functions depending upon the state of the terminal. To indicate the current function of a 2626A softkey, a visual representation of the key is displayed on the screen with a label describing the key's current function (see Fig. 1). A softkey may either perform a terminal function or cause a new set of softkey labels to be displayed on the screen, thus assigning a new set of features



Fig. 1. The relationship between the eight softkeys and their screen labels.

to all of the softkeys. Each set of labels can be thought of as a node of a softkey tree, hence the name tree-structured, screen-labeled softkeys.

Within the terminal family of which the 2626A is a member, a softkey labeling convention has been established. A key that performs a function is labeled in capital letters, while a key that performs a branch to a new set of softkeys is labeled in lower-case letters. Fig. 2 indicates the labels displayed at the top level of the softkey tree. These are all lower-case, so each branches to a separate functional area of the terminal. This is the functional choice level of the softkey tree, often referred to as the AIDS level. There is, in fact, a dedicated key labeled **AIDS** that always displays these labels when pressed, thus providing easy access to the top of the softkey tree.

An example will illustrate several points. Suppose you wish to enhance a certain section of text on the CRT screen by displaying it in blinking inverse video. Pressing the **AIDS** key returns the terminal to the top of the softkey tree. Pressing the fifth softkey, labeled enhance video, then causes a branch to the video enhancement level and displays the labels shown in Fig. 3. The keys labeled INVERSE VIDEO and BLINK VIDEO (**15** and **16**) can then be pressed to indicate the desired combination of enhancements. In response to this, asterisks appear in the lower right-hand corners of these key labels to indicate they have been selected. Next, the cursor is positioned on the screen at the starting position for the selected enhancements. When this is accomplished, the



Fig. 2. The top level of the softkey tree. Each key causes a branch to a separate functional area of the terminal.

SET ENHNCMNT key is pressed to propagate the selected combination. To end the enhancement the cursor is placed at the ending position and SET ENHNCMNT is again pressed. That's all there is to it. The softkey labels direct the user through the entire operation. If a mistake is made when selecting enhancements, pressing the erroneous key a second time toggles the displayed asterisk off, thus reversing its selection.

Contrast this to the older method. First there was no indication to the occasional user that video enhancements were available in the terminal (the enhance video key serves this function in the 2626A). Once aware, the user would have to position the cursor to the desired starting position and type in

escape & d C

to propagate the enhancement, then reposition the cursor and type

escape & d @

to terminate it. These escape sequences had to be memorized or found by looking through the terminal reference manual, either of which requires a much higher level of sophistication than letting the softkey labels direct the user's actions to the desired result. It should be noted that all 2626A features are still accessible to the host system through escape sequences. The softkey tree is simply a user-oriented means of accomplishing the same thing.

The 2626A softkey tree is shown in Fig. 4. Each row of softkey labels represents one level (or node) of the softkey tree and the labels within that level identify the functions available there. The arrows between rows indicate the transitions that are possible between levels using the branching (lower-case) keys.

Relating 2626A softkey operation to the four areas of user interface design leads to the following conclusions. First, extensive use of tree-structured, screen-labeled softkeys significantly improves the user model by breaking the terminal feature set into smaller functionally related areas, each accessible through branching keys. Second, the use of softkeys makes for a particularly simple command language. All features are available with a few keystrokes (often one) and there is virtually no syntax to learn or



Fig. 3. The video enhancement level of the softkey tree.





memorize. Consistent labeling conventions further simplify the command language by visually differentiating between function softkeys and branching softkeys. Third, the structure of the softkey tree presents several opportunities for enhancing the display of information to the user. Since the softkeys are tree-structured, only relevant information is presented. The use of sixteen characters to label key functions exceeds what can be placed on a normal keytop and thus makes it easier to understand the function of the key.

Menu-Driven Configuration

Early terminals were designed for specific purposes and their feature sets did not have a great deal of breadth. Later designs offered features that covered larger portions of the total terminal market. Some of these features were added as optional extras, while some were available as strap options within the terminal. Straps usually took one of two forms, either a physical switch on a printed circuit board or some type of wire jumper. The terminal could be made to exhibit a different feature set simply by modifying the arrangement of internal straps. The user became aware of the straps and their related meanings by reading the reference manual for the terminal. To alter the straps the user had to unplug the terminal, open the case, locate the board containing the strap, remove it, locate the strap, make the adjustment, reassemble the terminal, turn it on, and verify the operation of the terminal with the modified strap.

Strapping the terminal in this way, and setting a few external switches, such as the baud rate and duplex controls for data communications, is called configuring the terminal.

To enhance the user interface in this area and make the full terminal feature set accessible to all users rather than just the skilled few, the 2626A uses a menu approach to accomplish the same result. There are no physical straps inside the 2626A, so a user will never have to open the case. All configuration items are selected simply by choosing menu items displayed on the CRT screen during configuration mode. This approach not only improves the user interface but also eliminates keyboard and internal switches. Thanks to the relative ease of this approach it is now feasible to reconfigure the terminal quickly for different applications.

The terminal's feature set is broken into four smaller areas, each with its own configuration menu. The four areas are: global configuration, window configuration, terminal configuration, and datacomm configuration. Since the 2626A can act like four virtual terminals, there are four terminal configuration menus. Similarly, there are two datacomm configuration menus to accommodate the two datacomm ports.

Global Configuration. Contains configuration items global to all four virtual terminals. Examples are the terminal language (USASCII or one of the six international languages) and frame rate (50 Hz or 60 Hz).

Window Configuration. Contains all items associated with the window/workspace relationships within the terminal. This configuration menu is used to partition the terminal's memory into workspaces and to partition the CRT screen into windows.

Terminal Configuration. Contains all items specific to a given virtual terminal, i.e., virtual terminal modes (remote, block, format, etc.), handshaking requirements, and choice of alternate character sets.

Datacomm Configuration. Contains all datacomm items.



Fig. 5. The configuration level of the softkey tree.

One first chooses from a set of menus describing the various datacomm protocols supported by the terminal and then fills out the configuration items for that particular protocol.

An example will best illustrate the use of configuration menus. Suppose that the terminal language is to be changed from USASCII to Danish. Like all other features the configuration menus are accessed through the softkeys. Pressing the **AIDS** key returns the terminal to the top softkey level. Then, pressing config keys (**f8**) causes a branch that displays a set of softkey labels allowing access to all of the configuration menus in the terminal (Fig. 5). Since the terminal language is a global item, pressing **f1** (global config) brings up the appropriate menu (see Fig. 6). Note that the first configuration item on the second line of this menu is the language specification.

To change this the cursor is positioned in the field and the softkey labeled NEXT CHOICE (**13**) is used to scroll through a list of all available languages. Each depression of this key displays the mnemonic associated with one of the languages supported by the terminal. For our example, two depressions are required to display the desired mnemonic, DANSK/NORSK (Danish/Norwegian). At this point any other changes required in the global configuration menu can be made using a similar approach. When all changes have been completed and the SAVE CONFIG key (**11**) pressed, all changed items are activated. Typically, this is all that is required to change any configuration item in the terminal. Once again, the softkey labels combined with Englishlanguage field names in the configuration menu lead the user through the operation.

Within a configuration menu the cursor can be quickly advanced from field to field with a single depression of the **TAB** key. the **BACK TAB** key returns it to the previous field, allowing quick access to all menu items.

Values are selected for each configuration item in one of two ways. If the value field of the configuration item is underlined, it implies that the NEXT CHOICE/PREV CHOICE softkeys are to be used to scroll forward/backward through a predefined list of values for that field (as in our example). If not underlined, then the user types in the choice directly from the keyboard. Depressing **f4** (DEFAULT VALUES) or **f5** (POWER ON VALUES) causes the entire menu to be refreshed



Fig. 6. Global terminal configuration menu.

with either the default values stored in ROM or the poweron values stored in nonvolatile RAM.

Let's review the above features with respect to the four guidelines for designing user interfaces. As with the softkeys, the configuration menus enhance the user model by dividing the entire terminal function set into smaller related groups. All items associated with each functional area are presented on the same configuration menu, while unrelated items are always contained on other menus. This helps the user develop an understanding of the terminal by understanding its subfunctions.

The command language used to configure the 2626A is very simple. Once the appropriate menu is displayed, items to be modified are selected with the cursor and the user either types in the new value or selects from a predefined list of values. Again, there is no syntax to learn or remember.

Feedback to the user takes several forms during configuration. When using the NEXT CHOICE/PREV CHOICE keys the value of each choice is fed back via an English-language mnemonic in the selected field. When values are entered directly from the keyboard, the range of each value is checked and the consistency of all menu items is verified. When erroneous entries are detected the user is informed by error messages in the softkey label area, and the cursor is placed in the offending field to aid in correcting the error.

It is particularly in the area of presenting information to the user that the configuration menus make a large contribution. In the older schemes, each strap was given a letter identifier. This letter was then associated with its function in the reference manual for the terminal. With configuration menus, each "strap" is given an English-language mnemonic describing its function. The equivalent letter identifier is parenthetically included to maintain consistency with older products and with existing applications that may refer to straps by letter only. The fact that all choices available to the user are displayed at once greatly enhances the user's ability to make quick, accurate changes in the configuration without resorting to a manual. Since the NEXT CHOICE/PREV CHOICE keys present only valid data to the user, the possibility of error is greatly reduced.

Acknowledgments

The design of the user interface for the 2626A represents the contributions of a large number of people. In particular, Maxine Brown developed the framework for both the softkey tree and the soft configuration menus. Prem Kapoor, project manager for the 2626A, was instrumental in developing the softkey tree to its present state of refinement.



Gordon C. Graham

Gordon Graham did firmware programming for the 2626A Display Station. He joined HP in 1979 with several years of experience in programming and circuit design. A native of Covina, California, he received his BSEE and MSEE degrees from the University of California at Los Angeles in 1971 and 1975. He's married, has a daughter, lives in San Mateo, California, and enjoys photography and bicycling.

SPECIFICATIONS

HP Model 2626A Display Station

SCREEN SIZE: 150 mm × 215 mm (6 in × 8.5 in) diagonal measurement 262 mm (10.4 in).
SCREEN CAPACITY: 24 lines × 80 columns (1920 characters). 25th and 26th for labeling of function/softkeys or as message/status lines.

CHARACTER GENERATION: 7 × 11 enhanced dot matrix with interstitial dots; 9 × 15 dot character cell; non-interlaced raster scan.

CHARACTER SIZE: 2.4 mm × 3.5 mm (.094 in × .138 in).

CHARACTER SET: Upper/lower case, displayable control codes, extended line drawing. Optional math symbols, large characters, Finnish/Swedish, Danish/Norwegian, French, German, Spanish and U.K. characters.

CURSOR: Blinking-underline, blinking-square for insert mode.

DISPLAY ENHANCEMENTS: Inverse video, underline, blink, security

REFRESH RATE: 60 Hz (50 Hz optional)

TUBE PHOSPHOR: P4.

IMPLOSION PROTECTION: Tension Band.

- MEMORY: 80 characters by 119 lines including a default datacomm buffer. (Buffer expandable to 2K bytes—reduces memory to 80 characters by 107 lines.) 128 bytes nonvolatile configuration memory (battery powered).
- KEYBOARD: Full ASCII code keyboard; eight screen-labeled keys; cursor controls; 14-key numeric pad; auto-repeat; N-key rollover; detached with 1.2-m (4 ft) cable. Optional Finnish/Swedish, Danish/Norwegian, French, German, Spanish, and U.K. keycaps.

Data Communications

- DATA RATE: 110, 134.5, 150, 200, 300, 600, 1200, 1800, 2000, 2400, 4800, 9600 baud and external. Operation with control codes, escape sequences, integral printer, or baud rates above 4800 may require CPU supplied delays or handshakes. Typical printer throughput is 60 cps. Full 24 × 80 line character screen copies in 18 seconds.
- PORT 1 ASYNCHRONOUS/SYNCHRONOUS INTERFACE (50 PIN): EIA standard RS-232-C; fully compatible with Bell 103A, 202C/D/S/T modems. Choice of main channel and/or reverse channel in half-duplex. CCITT V.24 hardware handshaking available. Accessory pods provide current loop (13266A), asynchronous or synchronous multipoint (13267A first terminal, and 13268A daisy-chain terminal. 300 Baud Modem 13268A U.S. only).

TRANSMISSION MODES: Full or half-duplex, asynchronous, synchronous,

- PORT 2 ASYNCHRONOUS/SYNCHRONOUS INTERFACE (25 PIN): EIA standard RS-232-C; fully compatible with Bell 103A moderns, 202C/D/S/T moderns. Choice of main channel and/or reverse channel in half-duplex. CCITT V.24, hardware handshaking for control of external printer.
- TRANSMISSION MODES: Full or half-duplex, asynchronous, synchronous (does not support HP Multipoint polled protocol).
- OPERATING MODES (BOTH PORTS): On-line; off-line; character, line modify, line and block.

General

PARITY: Selectable; even, odd, none, 0, 1.

 POWER REQUIREMENTS

 INPUT VOLTAGE:
 100/120V (+5%, -10%) at 60 Hz (±5%) 100/220/240V (+5%, -10%) at 50 Hz (±5%)

 with Option 050
 115V (+10% -25%)/230V (+10%, -15%) at 50/60 Hz (±5%)

 POWER CONSUMPTION:75W
 with Option 050

 with Option 050
 120W

Srinivas Sukumar performed the detailed design of the configuration menus. Special appreciation is given to Tom Anderson for his help in defining the user interface for the windows and workspaces in the 2626A. The untiring efforts of Gary Lum and Frank Santos in verifying proper operation of the 2626A deserve special mention.

ENVIRONMENTAL CONDITIONS

TEMPERATURE, FREE SPACE AMBIENT: Non-Operating: -40 to +60°C (-40 to +140°F).

Operating: 0 to +55°C (+32 to +131°F).

WITH OPTION 050: Operating: +5 to +40°C (+41 to +104°F).

HUMIDITY: 5 to 95% (non-condensing).

WITH OPTION 050: 5 to 80% (non-condensing).

ALTITUDE: Non-Operating: Sea level to 15240 metres (50,000 ft). Operating: Sea level to 4572 metres (15,000 ft).

VIBRATION AND SHOCK: Vibration: 0.38 mm (0.015 in) pp, 5 to 55 Hz, 3 axis

Shock: 20 g, 11 ms, ½ sine,

Type tested to qualify for normal shipping and handling in original shipping container.

Prices in

PRODUCT SAFETY

Product meets the requirements of the following safety agencies for EDP equipment and/or office equipment in the following countries: Canada—CSA, Finland—FEI, Germany—VDE, Switzerland—SEV, U.K.—BSI, United States—U.L.

PHYSICAL SPECIFICATIONS

DISPLAY MONITOR WEIGHT: 16.8 kg (37 lb).

With Option 050: 19.0 kg (42 lb)

KEYBOARD WEIGHT: 2.0 kg (4.4 lb).

DISPLAY MONITOR DIMENSIONS: 380 mm W \times 475 mm D \times 440 mm H (15.0 \times 18.7 \times 17.3 in); 665 mm D (26.2 in) including keyboard.

KEYBOARD DIMENSIONS: 430 mm W \times 190 mm D \times 75 mm H (17 \times 7.5 \times 3.0 in). ORDERING INFORMATION:

		U.S.A.
2626A	Display Station. ASCII keyboard with numeric pad, 80 character by 119 lines display memory, serial I/O port, 60 Hz, 110 volt operation.	\$4150
Option		
001	Finnish/Swedish character set and keyboard†	265
002	Danish/Norwegian character set and keyboard†	265
003	French character set and keyboard†	265
004	German character set and keyboard†	265
005	United Kingdom character set and keyboard†	265
006	Spanish Language character set and keyboard+	265
013	240V, 50 Hz operation	
014	100V, 60 Hz operation	***
015	220V, 50 Hz operation	777
016	100V, 50 Hz operation	***
050	Integral forms copy thermal printer, 120 characters-per-second using 8½-inch-wide paper	1210
201	Math and Large Character sets. (Standard with any language option)	265

†Deletes U.S. keyboard and includes math and large character sets.

MANUFACTURING DIVISION: DATA TERMINALS DIVISION 974 East Arques Avenue Sunnyvale, California 94086 U.S.A.

Reference

1. W.M. Newman and R.F. Sproull, "User Interface Design," Principles of Interactive Computer Graphics, Second Edition, Chapter 28. McGraw-Hill Book company, New York, 1979.

Hardware and Firmware Support for Four Virtual Terminals in One Display Station

by Srinivas Sukumar and John D. Wiese

HE 2626A DISPLAY STATION is a sophisticated CRT terminal with features that form a superset of the 2645A Terminal's feature set. With its windowing capabilities and dual communications ports, the 2626A can function as up to four virtual terminals. It has a custom SOS video controller chip, 80K bytes of ROM, 32K bytes of display RAM and 2K bytes of program RAM for variables. A fast character ROM containing ASCII and extended line drawing characters is a standard feature.

The 2626A represents a significant contribution to CRT display terminal capabilities. The following objectives guided its design and development:

- 2645A-compatible feature set. Because of HP's large established customer base, it was considered mandatory that all host applications and device drivers developed for the 2645A function as well as with the 2626A.
- Improvement in price/performance ratio. To meet the needs of increasingly sophisticated users and take advantage of new technology, a significant enhancement of the feature set was required.
- Ease of use. There are many useful features in the 2645A that are tedious to use and require a detailed knowledge of the terminal's control sequences. A good example is the design of forms. The HP 2626A provides easier access to all terminal features. There are no hardware straps to be set by the user. Instead, menus provide a highly visible means of configuring the terminal and displaying its current state.
- Ease of manufacturing, reliability, and easy serviceability.

Firmware Design

The terminal's feature set is provided by a combination of hardware and firmware (microprograms stored in read-only memory). One of the main goals of the firmware design was modularity. This is achieved by a clear specification of the firmware interfaces. The firmware modules called intrinsics in the 2626A provide the interface between the hardware and the main code so that the main code is not overly burdened with the hardware aspects of the design.

The firmware in the 2626A is divided into six major parts: the main code, the operating system, the display intrinsics, the keyboard intrinsics, the datacomm intrinsics, and the printer intrinsics. The main code controls all the terminal's features, thereby providing the 2626A with its personality. All input processing beyond the interrupt service routines is done by the main code.

Traditionally, terminals do not have an operating system. The 2626A has a simple operating system that is adequate for its needs. In studying the various functions that the terminal has to perform, the following tasks can be identified:

- Input from datacomm port(s)
- Datacomm output(s)
- Keyboard input
- Block mode process. A block of data is transmitted when the ENTER key is hit.
- Device transfers. Data is transferred from the datacomm port to devices like workspaces and printers.
- User softkey processing. A user-defined string of characters is processed when a softkey is pressed.

The operating system on the HP 2626A is designed to control the execution of these tasks as the system (terminal) is responding to interrupts from its I/O ports.

A task in the 2626A may be in one of three states: waiting, scheduled or executing. If a task is currently executing, in the absence of a hardware interrupt it remains in control until it gives up control to the operating system by WAITing or YIELDing. That is, the operating system never specifically takes control away from an executing task. After the task WAITs or YIELDs, the operating system looks to see if another task has been scheduled and executes the first one that it encounters. It examines all tasks in a round-robin fashion such that the last task executed is the last task checked when the search is made to find another task to which to give control.

How long a task remains in control depends on the device or function associated with the task. For example, all tasks associated with I/O devices WAIT after all the characters in a burst from that device have been processed. But these tasks YIELD after processing a certain fixed number of characters in a burst. The datacomm tasks process 256 characters before YIELDing, but the keyboard YIELDs after each character processed. Tasks controlling BLOCK mode and device transfers WAIT after the complete data transfer is done. But they YIELD after processing one line of display memory.

Interwoven into this multitasking system is the concept of windows and workspaces. A workspace is a block of display memory that can be associated with a display screen window. Traditionally a workspace is all of display memory in a terminal and a window is the 24-line screen that is currently visible. The 2626A has four possible workspaces, with the number of columns in each configurable from 80 to 160 characters. The total in all workspaces is 119 lines of 80 characters. As the line width increases the number of lines available decreases proportionally.

Each task in the terminal is associated with a workspace. This association is established through the terminal workspace/window configuration, which can attach a datacomm port to a workspace and the keyboard to another workspace. Thus the 2626A can process inputs from both datacomm ports and the keyboard simultaneously.



Terminal Intrinsics

The display intrinsics are designed to interface with the video controller (see article, page 16) and the display memory. These also have the workspace/window concepts built into them. At any point in time all the functional intrinsics operate on the currently active workspace. Whenever a device becomes active, the workspace associated with that device gets activated and the function to be performed by the character is executed in that workspace, whether it involves displaying the character or executing the control function associated with it. The display intrinsics can be classified into six major groups:

- Configuration of workspaces
- Configuration of windows
- Screen related functions like scrolling
- Display of characters
- Control of cursor movement
- Display memory partitioning, allocation and deallocation.

The keyboard intrinsics are completely table-driven, so that the various keyboard options like the Danish, German, Swedish, French, Spanish and Katakana keyboards can be supported in addition to the standard typewriter-style U.S. keyboard. The keyboard intrinsics isolate the hardwaredependent aspects of the keyboard from the logical operations and data needed by the main code.

The datacomm intrinsics provide the interface between the main code and the datacomm hardware. This is a character interface independent of the current protocol. That is, the main code does not have to know at any point whether the current datacomm configuration is point-topoint or multipoint. The datacomm code operates using a context area that is set up during terminal configuration. The point-to-point driver is completely reentrant so that the same code can be used for both ports even when one of the ports is configured as a printer port. Fig. 1. The 2626A display section is controlled by the video control chip (VCC). This large-scale integrated circuit makes the display section very simple.

The print mechanism used in the 2626A has a dot interface, which requires that all character dot information be sent to the printer from the main code. The dots sent to the printer are the same character dots that are used by the video controller to display the data. Thus everything on the screen can be printed. The printer intrinsics control the printing of the characters as well as the position of the printhead. Thus some optimization is done to skip over blank characters and to print in both directions.

Hardware Design

To meet the goals of reliability, serviceability and ease of manufacture, the terminal electronics are on only three printed circuit boards: the power supply, the sweep, and terminal logic boards. This structure minimizes the number of interconnections, increasing reliability and making the terminal easy to build and service. Any problem in a terminal can be easily isolated and the faulty module quickly replaced.

The terminal logic board can be divided into four major sections: the microprocessor and its memory, the two datacomm interfaces, the keyboard scanner, and the display section. The microprocessor is a 16-bit HP CMOS SOS device chosen for its speed, low power dissipation, and input/output oriented instruction set. It is a modified MCC¹ called the MC5, which has TTL compatible inputs and outputs.

The program memory consists of ten 64K-bit ROMs (40K words) and 1K words of static RAM, which is used for the stack and for frequently used variables. To preserve the terminal configuration, a 256×4 -bit CMOS RAM is included. It is powered by a battery when the terminal power is off.

The datacomm section consists of two nearly identical ports. Port 1 supports synchronous or asynchronous, fullor half-duplex, point-to-point or multipoint communications protocols. It also provides power to support a family of external datacomm pods (current loop, multipoint, and a 300 baud modem). Port 2 supports only asynchronous, full-duplex, point-to-point datacomm. Although Port 1 requires more control lines to support its additional protocols, both ports appear identical to the microprocessor so that a single firmware driver can be used to service both ports. Both ports use a universal synchronous asynchronous receiver/transmitter (USART) and appropriate interface chips to provide an RS-232-C interface.

To relieve the MC5 of the time-consuming task of scanning the keyboard, the keyboard is scanned by a single-chip microcomputer. It has 1K bytes of ROM, 64 bytes of RAM, 16 input/output lines and an eight-bit interface to the MC5. It is programmed to scan the keyboard, detecting and debouncing changes in the state of all keys. When it finds that a key has been pressed, it interrupts the MC5 and reports both the keycode and the state of the qualifier keys (shift and control). Each key can be programmed to be repeating, either slow or fast, or nonrepeating. Individual keys can also be locked out so they will not be reported to the main processor when pressed. The scanner provides N-key rollover, so overlapping keystrokes of a fast typist will not be lost. Besides scanning the keyboard, the scanner also drives a small speaker in the keyboard. The frequency and duration of the tone are programmable from the main processor.

Display Section

The display section is responsible for storing the data to be displayed, sending the necessary video, horizontal and vertical drive signals to the sweep to refresh the display, and supplying character dot information to the microprocessor for use by the optional thermal printer. A block diagram of the display section is shown in Fig. 1.

The display section is based on the CMOS SOS video control chip (VCC) described in the article on page 16. The VCC controls and refreshes the 16K×16-bit dynamic RAM, which is used to store data and enhancements. The character ROMs are 32K-bit CMOS SOS devices with an access time of 150 ns. Each pair of ROMs contains two character sets. The 2626A has the full 128-character ASCII set (upper- and lower-case Roman, numerals and control characters) and an extended line drawing set as its standard set. Optional sets include math and large characters.

The VCC reads its configuration information from the display memory and also reads a pointer list stored in the memory by the MC5. The pointer list points to rows of characters to be displayed. The VCC picks up a pointer and begins fetching characters and their enhancements at that address. As each word is read from display memory, the character and the character set select bits are used as an address for the character ROMs. The enhancements are read by the VCC, then the dots for the character are read from the character ROMs. The VCC modifies the dots with the enhancements and shifts them out to the sweep.

When the MC5 needs to access the display memory to update the pointer lists or to read or write data, it simply addresses the desired location as though it were normal RAM. The VCC holds off the MC5 until the next time the VCC does not need to use the display memory (usually during horizontal or vertical retrace). The VCC then disables its address bus and enables the MC5 to address the memory. The VCC then cycles the memory and returns a signal to the MC5 that the operation is complete.

The terminal processor board is designed to be easily tested and repaired. An extensive self-test is run whenever the terminal is turned on. Any detected errors are displayed on the CRT. The error messages generally point the repair person to a particular socketed component that has failed. If replacing this component does not fix the problem, then the board is replaced with a new one.

Acknowledgments

Prem Kapoor was the project manager for the 2626A and provided a lot of guidance to the lab team. The datacomm intrinsics were written by Chris Vandever, Sara Hilbert and John Hill. Grant Head was responsible for the operating system. Brodie Keast and Bill Rytand helped him in designing it and wrote the printer intrinsics. We would also like to thank Janelle Bedke and Ed Tang for their contributions.

Reference

1. B.E. Forbes, "Silicon-on-Sapphire Technology Produces High-Speed Single-Chip Processor," Hewlett-Packard Journal, April 1977.



Srinivas Sukumar

Srinivas Sukumar received his BE degree in electrical engineering from Victoria Jubilee Technical Institute in Bombay in 1973 and his MS degree in electrical engineering from Washington State University in 1975. After two years of software design for medical information systems, he joined HP in 1977 and contributed to the firmware design of the 2626A Display Station. He's now a project manager with HP's Data Terminals Division. Born in Bombay, Srinivas is married and now lives in Sunnyvale, California. He enjoys racquetball, gardening, and table tennis.



John D. Wiese

Born in Norman, Oklahoma, John Wiese received his BSEE degree from Stanford University in 1969 and joined HP the same year. He's done hardware design for the 2570A Coupler/Controller and the 2313A Analog/Digital Subsystem, hardware and firmware design for the 2240A Measurement and Control Processor, and design of the processor board and self-test code for the 2626A Display Station. He's now a project manager with HP's Data Terminals Division. John enjoys volleyball, backpacking, and the study of wine. He's married, has two children, and lives in Palo Alto, California.

A Silicon-on-Sapphire Integrated Video Controller

by Jean-Claude Roy

HE 2626A DISPLAY STATION'S multifaceted personality is best seen through its display. The important display design objectives are very high character legibility and quality, hardware-supported windowing and blank filling, and horizontal and vertical scrolling. Other requirements are ease of manufacture, reliability, low cost, and low power consumption.

In support of these objectives a raster-scan display¹ operating at a 24.90-kHz line rate and a 25.7715-MHz dot rate was chosen. The character size is 7 dots by 11 scan lines inscribed within a 9-dot-by-15-line cell. Interstitial dots are fully encoded so that individual dots can be shifted one-half dot space for increased character resolution.

The screen format consists of 26 rows of 80 characters. The display may be configured so that any number of contiguous screen rows starting from the top constitute the upper screen, and the remainder the lower screen. This is easily done by having the 2626A firmware set a value for the demarcation row that indicates where the split is to occur. This allows flexible and efficient vertical screen partitioning into logical blocks. In normal operation the 2626A firmware reserves the upper 24 rows for the user and the bottom two rows for the softkey labels.

Horizontal screen splitting within either the upper screen or the lower screen can be done by configuring the display's upper and lower window screen count parameters. Using these and the demarcation row it is possible for the firmware to divide the screen into four independent windows.

To ease the task of discriminating between windows for the user, a programmable demarcation line can be turned on by firmware. This line coincides with the demarcation row, and its position, thickness, and intensity are controlled by means of appropriate configuration constants. Similarly, a right border may be turned on at the interface between the left and right windows. Finally, a bottom border may be activated on a row basis for each window to allow for further screen splitting in the vertical direction.

The two types of cursor supported by the display system are a blinking underscore and a full character cell rectangle. Cursor type, blinking, and underscore position within the cell are all selected by the firmware through cursor control words. Moving the cursor inhibits its blinking to make it visually easier to track for the terminal user.² This is done by internally comparing the old and new cursor positions for each frame and firing a blink-inhibiting digital one-shot multivibrator if they differ.

The display system supports ten combinations of the blink, underline, inverse video, half-bright, and security enhancements. In addition, a full-screen background setting capability allows the user to change the entire display from white-on-black to black-on-white or vice versa. Four character sets of 128 characters each can be accommodated. Finally, an internal dot and crosshatch pattern generator can be invoked to simplify terminal setup and alignment during manufacture. All of these features are available through configuration parameters and entries in the display memory.

To make this display system practical and reliable, it was felt that integration was mandatory. Since no commercially available CRT controller exists with this flexibility and large repertoire of features, a new video control chip (VCC) was developed. All the circuits required to support the 2626A's display system are included in the chip with the exception of a crystal oscillator, character ROMs, display



Fig. 1. Block diagram of the video control chip (VCC). Several interlocked state machines and counters are driven by parameters read from the display memory (black lines are control paths). memory, sweep circuits, a delay line, and a bus switch.

The reasons for integration are compelling. Given the complexity of this chip, the equivalent in TTL logic would have required approximately 220 integrated circuit packages and used 775 cm² of six-layer printed circuit board. This logic would have dissipated approximately 15 watts. Custom integration resulted in one component that requires about 29 cm² of board area and dissipates only 900 mW.

The decision to implement the VCC in HP's silicon-onsapphire (SOS) process³ was based on SOS's high speed, low power, static operation, and reasonable density.

Video Control Chip

The VCC has a flexible structure configurable by firmware to allow a variety of screen formats. Thus it may be used not only in the 2626A but also in other terminal products without any redesign. It supports character enhancement and all dot processing, and has its own built-in DMA capability.

The VCC is implemented as a hierarchy of interlocked state machines and counters driven by configuration and control parameters read from the display memory. A simplified block diagram of the VCC is shown in Fig. 1.

The master sequencer, a complex state machine, controls the overall operation of the VCC. The algorithm of the master sequencer is driven both by qualifiers from the raster generator representing the current state of the screen and by configuration information from the firmware. The master sequencer controls the address counter, the vertical window counter, and the memory controller. At chip turn-on it reads the hard configuration area of the display memory and loads the values into the chip's registers. These represent raster-defining parameters such as the number of screen rows and columns (see Table 1). During on-screen activity the master sequencer coordinates the fetching of pointers and characters from the display memory and during horizontal retrace it controls memory refreshing. Once per frame it goes through a soft configuration housekeeping cycle and reads such dynamic information as the current cursor position and the vertical window width.

The address counter is 14 bits wide, giving it an addressing range of 16K words. Under master sequencer control it performs the operations of row pointer and character fetching, display memory refreshing, and the addressing of the hard and soft configuration memory areas.

The memory controller arbitrates processor access of the shared display memory. In normal operation the VCC locks out the processor during on-screen columns and allows access only during horizontal and vertical retrace. Under processor command the VCC may allow itself to be preempted by the processor, which causes screen blanking. One use of this mechanism is the interleaving of VCC and processor accesses. Another use, helpful for whole-screen updates, involves blanking the entire frame from the time the processor desires the memory until the start of the frame following the release of the memory.

The video logic consists of separate dot and enhancement processors. These accept the character enhancement information read from the display memory and the character dot information read from the character ROMs. Together these processors generate the two video signals representing

			Table I
Hexad	ecim	al	
Add	ress		Hard Configuration Constants
	0	VHRC	Variable Horizontal Retrace Columns
	1	AC	Active Columns
	2	CHDR	Clear Horizontal Drive Signal
	3	SHDR	Set Horizontal Drive Signal
	4	(0)	(not used)
	5	SR	Number of Visible Screen Rows
	6	(0)	(not used)
	7	(0)	(not used)
	8	(0)	(not used)
	9	SSL	Starter Scan Lines
	Α	(0)	(reversed for testing)
	В	(0)	(reserved for testing)
	С	RL	Scan Lines per Row
			Soft Configuration Constants
	10	CROW	Cursor Row
	11	CCOL	Cursor Column
	12	CC	Cursor Control
	13	DC	Display Control
	14	DLC	Demarkation Line Control
	15	SLD	Demarkation Line Starting Scan Line
	16	SLDR	Demarkation Row Scan Lines
	17	DR	Demarkation Row
	18	ESL	End Scan Lines
	19	VR	Vertical Retrace Lines
	1A	LSEA	Lower Screen Enhancement Assignment
	1B	USEA	Upper Screen Enhancement Assignment
	1C	LSHM	Lower Screen Hardware Mask
	1D	USHM	Upper Screen Hardware Mask
	1E	LSVW	Lower Screen Vertical Window
	1F	USVW	Upper Screen Vertical Window
	20-2	2 (0)	(not used)
	23	(0)	(reserved for testing)
	24	(0)	(reserved for testing)

full-bright and half-bright bit streams, which go directly to the sweep assembly.

The dot timing generator is driven directly by the 25.7715-MHz dot rate clock. It provides all of the VCC's internal timing including the high-speed clocking required by the video logic's parallel-to-serial converters. It also



Fig. 2. The display memory contains two configuration areas, a pointer list, and the rows of characters to be displayed.

supplies the external timing needed for the display memory and the character ROMs. The 2626A uses an external delay line driven by a synchronization signal to generate the memory RAS and CAS signals (row and column address strobe).

The raster generator consists of the column, line, row, and vertical window counters. These configurable counters represent much of the VCC's flexiblity, allowing the chip to be firmware-configured for a variety of applications. The specific configuration required is loaded into the raster generator during the hard and soft configuration housekeeping cycles. By way of example, two screen configurations are used in the 2626A to select between the 60-Hz United States frame rate and the 50-Hz European rate. By judiciously selecting the appropriate line counts for vertical retrace and by padding the frame with dummy scan lines it is possible to select either rate and still maintain both constant vertical centering and a constant horizontal scan rate. Frame rate selection is seen by the user as a choice between 50 Hz and 60 Hz. When a change is made to the terminal configuration, the firmware modifies two constants in the soft configuration table. Beginning on the next frame, the raster generator responds to the new values by operating at the new frame rate. This type of flexibility totally eliminates the need for the usual plethora of switches and jumpers.

Display Memory Organization

The display memory may be either 4K or 16K words by 16 bits. It is implemented in the 2626A by means of industry standard dynamic RAMs. It could also be implemented in static RAM, PROM, or ROM, provided that the system tim-



Fig. 3. Photomicrograph of the video control chip.

ing requirements are met. Dynamic memories represent the lowest system cost for full-screen displays. However, these memories must be periodically refreshed to preserve their contents. For this reason the VCC performs explicit refresh cycles as part of its horizontal retrace activities. Once per scan line four successive columns are cycled. An internal seven-bit refresh register is kept as part of the address counter for this purpose.

Access to the display memory for refreshing the CRT is done on a real-time basis with no line buffering. The processor is allocated access during both horizontal and vertical retrace and may elect to lock out the video system for large block transfers. This represents a reasonable compromise between memory cost, bandwidth, and memory availability.

The display memory contains the two configuration areas, the pointer list, and the rows of characters to be displayed. Fig. 2 shows the data structure used. At the bottom starting at address 0 is the hard configuration area. This is read once at chip turn-on and is never accessed again. The soft configuration area is located from address 10_{16} to address 24_{16} and is read once per frame to update the dynamic screen parameters.

The pointer list starts at location 25_{16} and contains three words per configured screen row. These are a pointer to the left window characters, a pointer to the right window characters, and two concatenated end-of-data bytes. These bytes represent the number of characters to display in the row's left and right window areas, respectively, before starting hardware blank filling.

The pointers may be manipulated by the firmware to implement scrolling, row insert, and other functions. Each pointer can address the entire 16K range of the display memory. It is the responsibility of the firmware to build a valid and consistent data structure since the VCC cannot write to the display memory nor does it know how big it really is.

The character entries in display memory each consist of seven bits of ASCII representing the character itself, four bits of character enhancement, and two bits for character set selection. The remaining three bits are reserved by the firmware for 2626A features. Mask registers in the VCC may be set to disable the enhancement and character set select bits on a half-screen basis. These are used if a chip's application requires the bits to be reassigned to a different purpose.

The logical length of a row of characters may be different from its physical length on the screen. Portions of longer rows will be off-screen until they are scrolled into view. This is accomplished in firmware by altering the row's pointer to redefine what is to be displayed in the window. Shorter rows may be blank-filled by the firmware at the expense of display memory. Alternatively, the firmware could use the row's end-of-data byte, which indicates how much of the row to display, with the remainder to be hardware blanked.

VCC Implementation

Fig. 3 shows a micrograph of the VCC with the principal blocks labeled. The chip's logic is designed for maximum speed in the critical circuits while minimizing area (the die is 4.9 mm by 5.8 mm). The configuration constants are stored in latches connected to an internal 16-bit data bus and clocked by the master sequencer. The same data bus is used as a stimulus bus when the chip is tested at wafer sort and package sort.

The internal state machines are implemented using D-type flip-flops and PLAs (programmable logic arrays). The chip is entirely static, a feature of CMOS/SOS that simultaneously allows low power dissipation and high speed. Static operation makes testing of this chip relatively easy. A circuit may be set to a state and left there for an arbitrary amount of time without loss of information.

Physical Description

The VCC is packaged on a 64-pin leadless square ceramic substrate measuring 4.4 cm by 4.7 cm (Fig. 4). This package fits into a special zero-insertion-force socket and is held in place with spring clips. Heat dissipation is through the substrate and a temperature rise of 10° C is typical at the 900 mW power level at which it operates.



Fig. 4. The VCC is packaged on a 64-pin leadless ceramic substrate measuring 44 mm by 47 mm.

Acknowledgments

No LSI chip design can be done without the help of large numbers of people and organizational entities. My thanks to all of them. Special thanks go to David Parks (now at HP's Vancouver Division) for his design contributions and to Stan Moriya for developing the test and characterization programs. Mention must also be made of Al Desroches and Frank Wilson at the Cupertino Integrated Circuits Operation for all their help in getting the VCC into production.

References

1. J.C. Roy, "A High-Resolution Raster Scan Display," Hewlett-Packard Journal, June 1975.

J.C. Roy, "Display Apparatus Having Improved Cursor Enhancement," United States Patent No. 3,967,266, June 29, 1976.
 B. Forbes, "Silicon-on-Sapphire Technology Produces High-Speed Single-Chip Processor," Hewlett-Packard Journal, April 1977.



Jean-Claude Roy

A native of Montreal, Canada, Jean Roy did his undergraduate work at the University of California at Davis, graduating in 1970 with a BSEE degree. He received his MSEE degree from Stanford University in 1974 and his MSCS degree from the University of Santa Clara in 1979, With HP since 1970, he's been responsible for software design for the 5407A Scintigraphic Data Analyzer, for the organization and logic design of the 2640A Terminal's display, and for the definition, design, and project leadership of the video control chip for the 2626A Display Station. Now a project

manager with HP's Data Terminals Division, he's a member of the IEEE and a registered professional engineer in California. Jean is married and lives in San Jose, California. He's an avid reader and is interested in art, gourmet cooking, real estate, and raising dogs.

SC-Cut Quartz Oscillator Offers Improved Performance

This compact oscillator is designed to serve as a built-in precision frequency source. New technology and packaging provide lower power consumption, faster warmup, better stability, and lower phase noise.

by J. Robert Burgoon and Robert L. Wilson

IGH-STABILITY OSCILLATORS are key components in many instruments and systems. Such oscillators serve as time bases for frequency counters, time bases for navigational systems, frequency sources for synthesizers and spectrum analyzers, and local oscillators for radar and communication equipment.

The new HP Model 10811A/B Crystal Oscillator (Fig. 1) is a component quartz oscillator that puts out a highly stable 10-MHz signal. It is designed to provide an instrument or system designer with a frequency reference or source for applications requiring a low aging rate, a spectrally pure signal, or both.

Intended as a next generation plug-in replacement for the HP 10544A/B/C Oscillator, the 10811A/B offers significant improvements in almost all areas of performance.

Phase noise. The phase noise floor of the 10811A/B is 10 dB to 15 dB lower than the older 10544A/B/C Oscillator. Phase noise is becoming more and more important, especially in applications where the reference oscillator output must be multiplied up to some higher frequency. Frequency multi-



Fig. 1. The new 10811A/B 10-MHz Crystal Oscillator is a component oscillator for instruments and systems. A new quartz crystal cut and a new circuit design give it many improved features including lower phase noise and lower power consumption than previous designs. Two models, A and B, provide a choice of electrical connection methods.



Fig. 2. Typical 10811A/B phase noise plot, showing the noise processes responsible for the overall phase noise.

plication causes phase noise and sidebands to increase by 20 dB for every decade of multiplication. For example, multiplying the 10-MHz output of the 10811A/B to 500 MHz will cause a 34-dB increase in phase noise. Fig. 2 shows a typical 10811A/B phase noise plot. The slopes of the straight lines drawn through the data indicate the different noise processes responsible for the overall phase noise of the oscillator.

Time domain stability. Time domain stability is often divided into long-term and short-term regions. Long-term stability is measured over intervals of days, and short-term stability is measured over intervals of seconds. Long-term stability, also called aging or frequency drift, is the key specification in many applications. It determines how often an oscillator must be adjusted to keep its absolute frequency within useful limits, which vary depending on the specific application. At the factory, all oscillators are monitored in an automated system until their aging is less than 5×10^{-10} /day over a three-day period.

Short-term stability is closely related to frequency jitter and is a type of noise measurement. The 10811A/B's shortterm specifications are factors of 2 to 33 better than the 10544A/B/C specifications. Fig. 3 shows a plot of time domain stability measurements of two 10544B Oscillators and



Fig. 3. Time domain stability (short-term) measurements of two 10811A Oscillators and two 10544B's (an older design).

two 10811A Oscillators.

Warmup. The 10811A/B Oscillator cuts the 10544A/B/C warmup time specification in half: 10 minutes for the 10811A/B compared to 20 minutes for the 10544A/B/C Oscillator. Warmup time is defined as the time between oscillator oven turn-on and the time when the output frequency is within 0.05 Hz of the operating frequency. In portable instruments, where battery weight must be minimized, low power and quick warmup are crucial. Since the instrument warmup time is likely to be much shorter than the oscillator warmup time, the latter usually dominates. A side benefit of the 10811A/B is that its frequency just after oven turn-on is only 100 Hz low compared to 1 kHz low for the 10544A/B/C. This allows faster lock-in for instruments using phase-locked loops.

Oven power. Fig. 4 compares the steady-state oven power of the 10811A with that of a 10544B, both in moving air. At 25°C the 10811A requires less than one-half the power of the 10544B. Power consumptions of the 10544A and the 10811A in still air are also shown in Fig. 4 for 25° C only. Low oven power makes the instrument or system designer's job much easier, especially for portable applications.

Temperature coefficient. The dominant environmental factor affecting the oscillator frequency is temperature. The 10811A/B is specified to be three times less sensitive to temperature than the older 10544A/B/C. The actual performance comparison shown in Fig. 5 is even more striking.

Oscillator Design

Fig. 6 shows a block diagram of the 10811A/B Oscillator. The oscillator loop includes a 10-MHz SC-cut crystal (see page 22) which is mode-suppressed to allow only C-mode operation. The buffer stage transfers the 10-MHz signal to the output stage while isolating the oscillator loop from output effects. The buffer stage also drives the AGC (au-



Fig. 4. Steady-state oven power of the 10811A compared with that of the 10544A/B.

tomatic gain control). The AGC holds the output level constant, but more important, holds crystal current constant to prevent frequency changes caused by variations in the crystal drive level. The oven controller monitors the temperature of a thermistor embedded in the oven mass and accurately controls the temperature by applying power to the



Fig. 5. The new 10811 A/B Oscillator's frequency is much less sensitive to ambient temperature than the older 10544B's.

The SC Cut, a Brief Summary

by Charles A. Adams and John A. Kusters

Quartz crystals for use as frequency determining elements in oscillators are cut from large pieces of quartz, the cuts making various angles with the crystal lattice. Two common cuts, called AT and BT, are single-rotation cuts, which means that the master crystal is rotated about one of its axes to the desired angle and then cut. A doubly-rotated crystal is rotated about two axes before cutting.

The SC (stress compensated) cut used in the 10811A/B Oscillator was introduced in November 1974. It has also been known as the TTC (thermal transient compensated) and TS (thermal stress) cut. Many papers extolling its virtues and faults have been written, and the most-often-aired complaint is that it is difficult to make. It is a doubly-rotated crystal, which adds to its difficulty of manufacturing. However, HP has been making doubly-rotated crystals since 1965, so the SC cut did not present insurmountable difficulties.

A brief comparison of the properties of the SC crystal with those of AT and BT crystals is as follows.

Superior SC Qualities:

Fast warmup. No frequency overshoot

Much smaller second-order temperature coefficient (Fig. 1)

(ing. i)

Almost zero amplitude-frequency effect Essentially no activity dips (coupled modes) Better short-term stability Lower acceleration sensitivity

Equal:

Long-term aging

Shunt capacitance

Inferior SC Qualities:

Tighter angular tolerances required

Q lower than BT Higher series resistance

Five factors affect crystal frequency. They are dimensions, density, elastic moduli, third-order elastic effects, and coupling of unwanted modes. The AT and BT cuts, traditionally used for precision resonator applications, compensate for statically induced temperature changes in the first three factors. The SC is a doubly-rotated cut that compensates for changes in all five factors. Although thermal gradients induced in guartz plates result in momentary frequency excursions that are not predictable by basic elastic theory, recent developments in the theory permit determination of a crystal cut that is compensated both statically, as in the AT and BT cuts, and dynamically for momentary thermal gradients and surface stresses. Dynamic compensation results in a cut that exhibits very small frequency changes at turnover even though the temperature is changing fast enough to cause thermal gradients in the crystal. As an illustration, consider the case where BT and SC resonators are exposed to a thermal shock large enough to induce a gradient through the thickness of the crystals so that at equilibrium the temperature has risen 1°C. The instantaneous frequency change of the BT resonator would be some four orders of magnitude larger than that of the SC. For an actual case of a fast warmup oven, the BT would require about 20 minutes to stabilize within 1×10^{-9} of its final frequency, whereas the SC in the same oven would require less than six minutes and there would be no frequency overshoot. It should be pointed out that the SC is not necessarily operated at a turnover point. The 10811A/B Oscillator is designed to work with the crystal between 80° and 84°C. At the operating temperature, the crystal's frequency-temperature slope can be no larger than 1.5 parts in



Fig. 1. Comparing the SC-cut crystal's temperature performance with that of two other common cuts, AT and BT.

10%/°C. This slope, coupled with the very high thermal gain made possible by a novel oven design, permits satisfactory operation in ambient temperatures from -55° C to 71° C.

Conventional AT and BT cut crystals all have problems with coupled modes, known as activity dips. The SC does not exhibit this problem. Over a wide range of curvature, activity dips do not seem to exist. Therefore, curvature is not a critical factor in the SC.

The short-term stability of the SC is typically better than that of AT or BT crystals, even though its Q is lower. Using the Allan variance method of measurement, typical values for one-second samples are:

$$\begin{array}{rrrr} AT & \sim & 5 \times 10^{-12} \\ BT & \sim & 2.5 \times 10^{-12} \\ SC & \sim & 1.25 \times 10^{-12} \end{array}$$

Many oscillator users are particularly interested in the performance of the SC under acceleration. One-hundred-percent testing of this parameter has shown that the SC when used in the 10811A/B Oscillator consistently has less than 5 parts in 10¹⁰/g frequency fluctuation.

For a detailed study of the SC crystal, the reader is directed to the bibliography.

Acknowledgments

The authors with to express their thanks for much of the early work of fabricating and testing these devices to Jerry Leach of the Andover Division, and to Henry Yoshida of Hewlett-Packard Labs. Also, our thanks to Don Hammond of Hewlett-Packard Labs for his encouragement to carry out this development program.

Bibliography

 R. Holland, "Nonuniformly Heated Anistropic Plates: II, Frequency Transients in AT and BT Quartz Plates," 1974 Ultrasonics Symposium Proceedings, pp. 593-598.
 E. EerNisse, "Quartz Resonator Frequency Shifts Arising from Electrode Stress," Proc. 29th Annual Symposium on Frequency Control, pp. 1-4, 1975.

3. J. Kusters, "Transient Thermal Compensation for Quartz Resonators," IEEE Trans. Sonics and Ultrasonics, SU-23, pp. 273-276, July 1976.

4. J. Kusters and J. Leach, "Further Experimental Data on Stress and Thermal Gradient Compensated Crystals," Proc. IEEE, Vol. 65, pp. 282-284, Feb. 1977.

 J. Kusters, C. Adams, H. Yoshida, and J. Leach, "TTC's - Further Developmental Results," Proc. 31st Annual Symposium on Frequency Control, pp. 3-7, 1977.
 A Ballato, E. EerNisse, and T. Lukaszek, "The Force-Frequency Behavior of Doubly Rotated Quartz Resonators," Proc. 31st Annual Symposium on Frequency Control, pp. 8-16, 1977.

7. J. Gaqnepain, J. Poncot, and C. Peqeot, "Amplitude-Frequency Behavior of Doubly Rotated Quartz Resonators," Proc. 31st Annual Symposium on Frequency Control, pp. 17-22, 1977.

8. A Warner, B Goldfrank, M. Meirs, and M. Rosenfeld, "Low 'g' Sensitivity Crystal Units and Their Testing," Proc. 33rd Annual Symposium on Frequency Control, pp. 306-310, 1979. 9. D. Hammond, C. Adams, P. Schmidt, "A Linear, Quartz-Crystal, Temperature Sensing Element," ISA Trans., Vol. 4, pp. 349-354, Oct. 1965.

 R. Burgoon and R Wison, "Design Aspects of an Oscillator using the SC Cut Crystal," Proc. 33rd Annual Symposium on Frequency Control, pp. 411-416, 1979.
 J. Kusters, M. Fischer, and J. Leach, "Dual Mode Operation of Temperature and Stress Compensated Crystals," Proc. 32nd Annual Symposium on Frequency Control, pp. 389-397, 1978.

12 J. Kusters and C. Adams, "Production Statistics of SC (or TTC) Crystals," Proc. 34th Annual Symposium on Frequency Control, pp. 167-174.

Charles A. Adams



With HP since 1962, Charles Adams has done quartz crystal R&D, managed acousto-optics and piezoelectric labs, and served as production engineer for the SC-cut crystal. Currently he supervises the crystal physics section of the precision frequency sources group. A member of the IEEE and author of a dozen papers on quartz and surface wave resonators, he holds a BA degree in physics and mathematics from McMurry College and has done graduate work at Stanford and San Jose State Universities. Born in

Cisco, Texas, Charles now lives in San Jose, California. He's married and has three children. His interests include hunting, backpacking, carpentry and cabinet-making, and working with a church youth group. He was a member of the U.S. National Guard for ten years and is an instructor in cardiopulmonary resuscitation for the American Red Cross.

two heaters. The specific oven temperature depends on the individual crystal and is about 82°C. Because of their power dissipation, the output amplifier and the oven controller are outside the oven.

Crystal Mode Suppression

The SC-cut crystal is capable of resonating in many different modes. The crystal we use is cut for a thirdovertone C-mode resonance at 10.0 MHz. The thirdovertone B-mode resonance is above this at 10.9 MHz. Below 10 MHz, the next mode is the fundamental A mode at 7 MHz, and below that are the strong fundamental B and C modes. In all overtone crystals, the fundamental modes will dominate if they are not suppressed. This leads to the common high-pass mode-suppression techniques, such as low-pass traps or tanks. The high-pass techniques are not sufficient to suppress any modes above the desired frequency, such as the third-overtone B mode at 10.9 MHz. One way to implement bandpass mode suppression is illustrated in Fig. 7. The capacitive shunt arm C3 is replaced with a series-parallel network. With properly chosen component values, the conditions for oscillation will occur only over a selected band of frequencies. Fig. 8 shows equivalent circuits for low, middle, and high frequencies. At low frequencies the series-parallel network looks inductive. With one shunt arm inductive and the other shunt arm capacitive, it is not possible to produce in-phase feedback for oscillation. At high frequencies the result is similar. Here again, the series-parallel network looks inductive. No oscillation is possible because of the mixed capacitive and in-



John A. Kusters

Jack Kusters was born in Racine, Wisconsin. After four years in the U.S. Air Force, he attended Loyola University of Los Angeles, graduating with a BSEE degree in 1964. He received his MSEE and Engineer degrees from Stanford University in 1965 and 1968. With HP since 1965, Jack did research in physical acoustics for 13 years, then served as project leader, section manager, and now production engineering manager for several Santa Clara Division product lines. Out of his work have come 17 papers and 20

patents, mostly on acoustic resonators. A member of the IEEE, Jack has served on the Frequency Control Symposium committee for four years. Now living in Cupertino, California, he's married, has four sons, and enjoys working with wood, computers, and the Boy Scouts. He's also an American Red Cross instructor in cardiopulmonary resuscitation and first aid.

ductive shunt arms. At frequencies inside the passband both shunt arms are capacitive and the circuit is capable of producing the proper phase shift for oscillation. If the crystal has a resonant mode inside the passband and the loop has enough gain, the circuit will oscillate.

Phase Noise

When considering the phase noise of an oscillator, it is convenient to think of the oscillator as having two parts, an oscillating loop section and a buffer amplifier section. If properly designed, the buffer amplifier section contributes



Fig. 6. Block diagram of the 10811A/B Oscillator. The oscillator loop includes a 10-MHz SC-cut quartz crystal.



Fig. 7. The SC-cut crystal in the 10811A/B Oscillator is mode-suppressed to allow only C-mode operation. This diagram shows a standard Colpitts-type oscillator (top) and the same oscillator with mode suppression.

only 1/f and white phase noise and the oscillating loop section contributes only $1/f^3$ and $1/f^2$ phase noise processes. It is the filter action of the crystal that provides the extra $1/f^2$ factor in the oscillating loop section.

The method of extracting the signal from the oscillating loop is very important in achieving good phase noise. Since the crystal is a very good filter, the crystal current is a very clean signal. A simple way to extract the signal cleanly is shown in Fig. 9a; this method was used in the 10544A/B/C Oscillator. A capacitor is placed in series with the crystal so that the crystal current flows through the capacitor. The voltage across the capacitor is proportional to the crystal current and is used to drive the buffer amplifier stage. All the oscillator loop noise (except crystal noise) is filtered by the crystal. The output voltage level is proportional to the capacitor impedance and the crystal drive current. For a good signal-to-noise ratio in the buffer stage, a large loop output voltage is desirable. One way to increase the loop output voltage is to decrease the capacitance. However, this becomes a problem, because as this capacitor approaches the value of the tuning capacitance, it severely restricts



Fig. 8. Equivalent circuits for the mode-suppression circuit of Fig. 7. Top: low and high-frequency equivalent circuit. Bottom: passband equivalent circuit.



Fig. 9. Phase noise is sensitive to the method of extracting the signal from the oscillator. (a) Method used in the 10544A/B/C. (b) Improved 10811A/B method.

tuning range. A compromise is required to balance the tuning range and noise requirements.

In the 10811A/B Oscillator, a new approach was taken. The circuit, shown in Fig. 9b, allows the capacitor to be made small without affecting the tuning. The crystal current is run through a common base stage, which feeds the crystal current to the capacitor. The amplifier isolates the capacitor from the tuning circuits since the oscillating loop sees the amplifier as an impedance of $(r'_e + r'_{bb}/\beta)$. This is small compared to the crystal resistance, so it does not affect the crystal Q significantly. Since the capacitor is no longer in series with the tuning circuit, it may be made quite low to get a large loop output voltage. The other factor affecting the voltage is the crystal current. A value of 1 mA was chosen, giving a crystal dissipation of 50 μ W.

Oven Power Consumption

An important oven design requirement was to reduce the normal operating power consumption below that of the 10544B, while still using a linear controller. The 10544B with a linear controller requires 4.5 W of oven power in an ambient temperature of 25°C. Because this oven uses a resistive heater, at 25°C about half of this power is dissipated in the pass transistor located on the outside of the package. This power is effectively wasted. The 10544A uses a more power-efficient design that keeps power consumption low in the control transistor by switching at 3 kHz. However, the switching does give spurious frequencies that are unacceptable in some applications. The 10811A/B Oscillator combines the best of both designs by eliminating the heater winding and using two control transistors to heat the oven with direct current only (Fig. 10). Heaters Q1 and Q2 are bolted to the oven cavity. R1, a small resistor, develops a voltage V1 proportional to the heater current. This voltage is used to limit the warmup current and provide feedback for the main oven control loop.

The thermal resistance between the oven cavity tempera-



Fig. 10. 10811A/B oven controller block diagram. Two heater transistors heat the oven using direct current only. Overall power consumption is less than two watts in still air at 25°C.

ture and the ambient temperature is 25% higher than that in the 10544B. This was accomplished by reducing the size of the oven assembly, thus allowing the foam insulation to be slightly thicker. The result of these efforts is an oven oscillator that draws less than two watts in still air at 25° C, compared to 4.0 W for the 10544B/C and 2.6 W for the 10544A.

Warmup Time

The warmup time for most crystal oscillators is mainly a function of the thermal transient behavior of the crystal. The warmup time for the 10544A/B/C Oscillator, which uses a BT-cut crystal, is specified as 20 minutes to within 5×10^{-9} of the 24-hour value for V_{cc} = 20V (8 watts). Because the SC-cut crystal is much less sensitive to thermal transients, the warmup time for the 10811A/B Oscillator is governed by the most temperature-sensitive electronic components and, of course, the thermal capacitance of the oven assembly. To minimize total thermal capacitance, and to improve heat transfer to the electronics, the oven assembly is as compact as possible. The oven cavity and lids are made of die-cast aluminum, while the oven cavity in the 10544B is made of copper. Aluminum has a lower thermal capacitance per unit volume than copper by a factor of 0.7. The most temperature-sensitive components are located on the two printed circuit boards, which are folded into the oven cavity. When the two lids are screwed in place, these components are surrounded by heated aluminum walls. For the above conditions, the 10811A/B warmup time is specified as 10 minutes.

Because power transistors are used for the heat sources the heater current I_H must be limited during warmup. This is accomplished by U1, R1, R4 and R5 in Fig. 10. When V_{cc} is applied to the oven, U2's output is approximately V_{cc}, tending to turn Q2 on. U1 sinks just enough base current from Q2 to make V1 = V3, and therefore limits I_H to V3 ÷ R1. Since V3 is a linear function of V_{cc}, this circuit transforms Q1 and Q2 into what appears to be a fixed heater resistance. This was done to allow the 10811A/B Oscillator to be used in applications where the 10544A/B/C Oscillator is used. The 10544A/B/C uses a 47 Ω resistive heater winding.

Temperature Coefficient

The reason for enclosing an oscillator in an oven is to reduce the effect of ambient temperature fluctuations on the temperature-sensitive components in the oscillator. A measure of the ability to do this is termed thermal gain. Thermal gain is defined as the ratio of the change in ambient temperature (TA) to the resultant change in temperature of the region in the oven that it is desired to control (T_C) . Thus thermal gain = $\Delta T_A / \Delta T_C$. Even though the SC-cut crystal has an improved temperature coefficient, it is still ten times more temperature-sensitive than any other component in the 10811A/B Oscillator. This means that if all of the components of the oscillator including the crystal were to experience the same temperature change, the resultant change in frequency due to the crystal would be ten times more than that due to any other component. It follows that the thermal gain to the crystal must be ten times greater than the thermal gain to the rest of the oscillator electronics for both areas to contribute equally to the oscillator's temperature coefficient.

The temperature coefficient of the crystal varies according to the difference between the actual crystal operating temperature and the turnover temperature of that particular crystal. Turnover temperature is the temperature at which the derivative of crystal frequency with respect to crystal temperature is zero, that is,

$$\frac{\mathrm{d}\Delta f/\mathrm{f}}{\mathrm{d}\mathrm{T}} = 0.$$

(See Fig. 1 on page 22.) It is desirable to set the oven temperature very close to the crystal turnover temperature. Most crystal oscillators use a variable resistor to set the oven temperature. A time-consuming process must be followed in which changes in the oscillator frequency are compared to turns of this potentiometer. This process continues until the frequency change per turn is small enough. At this point the oven temperature is considered matched with the crystal's turnover temperature.

In the 10811A/B, the combination of high thermal gain and the SC-cut crystal's lower sensitivity to temperature allows the oven temperature to be set merely by installing a fixed resistor. Avoiding the long temperature-set process makes the 10811A/B the first field-repairable HP crystal oscillator.

Thermal Gain

To achieve a high thermal gain to a particular region in a (continued on page 28)

Flexible Circuit Packaging of a Crystal Oscillator

by James H. Steinmetz

Design objectives for the 10811A/B Oscillator packaging scheme included greater electrical performance, decreased manufacturing cost, and easier assembly and repairability. This oscillator employs a new crystal technology with higher performance support electronics and a radically new oven design. The electronic packaging takes a new and radical approach, one of selectively stiffened flexible circuitry. The operation and basic layout of this oscillator is much the same as that of the earlier 10544A/B/C. A block diagram illustrating the basic similarities of these oscillators is shown in Fig. 1. The flex circuit forms the backbone of the new oscillator. It incorporates the three printed circuit boards, the I/O connector and the various interconnections. The oscillator and AGC board are enclosed within a die-cast crystal oven, thereby minimizing thermal effects on these critical circuits.



Fig. 1. Cross-sections of the new 10811A Oscillator and the older 10544A.

Oscillator Assembly

The use of flex circuitry packaging in the construction of the new oscillator makes it ideally suited for high-volume production. The production sequence is as follows. First, the flex circuits are obtained from a manufacturer in panels of five complete circuits per panel, as illustrated in Fig. 2. The panels are selectively treated with a removable solder resist in such areas as the I/O connector. Then the panels are loaded with components either by manual or automatic means. At this point, the panels are wave soldered and cleaned. The loaded and soldered panels are transferred to the electrical tester. This instrument gives a complete functional check to the circuitry and components. If any problems are detected, such as solder bridging or parts out of specification, they are repaired at this point. The panels are then transported to a machine that separates the individual oscillator flex circuits from the panel. From here, the individual oscillator flex circuits from the panel.



Fig. 2. Flex circuits come five to a panel and are loaded, soldered, and tested before being separated.

circuits are loaded into trays and transported to the final assembly operation.

The final assembly consists of installing the flex circuit into the crystal oven, soldering the crystal leads to the board and screwing the oven covers into place. The oven control board is then folded over the top of the oven assembly and separated from it by a thermal insulation spacer. At this point, the heaters for the oven, which are two Darlington transistors in TO-220 packages, are folded down from the oven control board and fastened to the oven assembly. Two leads from the oven's thermistor are then soldered to the oven control board. In this configuration, the assembly is slid into the outer housing and the I/O connector is fastened into position. Another thermal insulator is installed on top of the oscillator assembly and the outer cover is installed. These final assembly sequences are illustrated in Fig. 3.

The 10811A, described above, has an edge-card connector for the package I/O. The B model is electrically the same as the A except that it has filtered power and coaxial signal connections.

The interconnection of the new B model oscillator I/O was significantly improved over that of the 10544C by employing a sculptured flexible circuit (SFC) interconnect. This interconnect is soldered to the various connections on the outer cover and runs to a simple connector on the oven control board. The A version's edge-card connector and its associated flexible interconnection are removed by a simple cutting operation at the point where the flex leaves the oven control board. This technique allows for the double use of the same flex circuit assembly, thereby allowing increased production volume and reduced unit costs. A comparison between the 10544C and the 10811B is shown in Fig. 4.

Benefits

The use of flex circuits in this new oscillator has resulted in a substantial decrease in the labor required for assembly. Board area has been reduced since some pads and connectors are no longer required. This design also allows a complete electrical



Fig. 3. Final assembly of 10811A/B Oscillator.

check of the flex assembly before final assembly. Repairability has been greatly simplified, since the oscillator can remain electrically assembled and still be folded out so that it is accessible for troubleshooting.

Flex circuitry made it possible to include an edge-card connector for input/output. This is accomplished by folding the flex circuit over the edge-card stiffener board (onto the component side) before the lamination of the flex circuit to the stiffener board. This eliminates the need to solder or otherwise connect to a package I/O, Furthermore, the flex circuitry simplifies the interconnection of the TO-220 heater transistors for the oscillator's oven. These interconnections are actually fingers of flex circuitry that extend out from the oven control board. At the final assembly stage, the transistors, which have been previously soldered to the flex fingers, are folded in a 180° arc and mounted to the oven.

A side benefit of the flex circuitry is that capacitors can be formed by matching circular pads on both sides of the circuitry separated by the base dielectric. As illustrated in Fig. 5, this capacitor looks like a component pad without a hole. Capacitors can be formed in this way on rigid printed circuit boards too, of course, but since rigid boards are thicker the capacitor pads must be considerably larger.

Design Problems

A problem that occurred early in the development of this new oscillator was the flex circuit's tearing within the interconnect region. This can be an inherent limitation with a flex circuit constructed of a polyimide material because of its relatively low shear strength. The problem was rectified by using larger radii at the interconnection interfaces to the stiffener boards and improving the techniques used to blank out the flex circuit during its manufacturing. Another problem encountered was the cracking of conductor traces in the pad area of the flex fingers that interconnect the TO-220 heater transistors. This problem was solved by increasing the pad area encapsulated within the flex circuit laminate, in effect creating its own strain relief.



Fig. 4. The sculptured flexible circuit interconnect of the 10811B (left) is a significant improvement over the older 10544C (right).

The most difficult problem was ultimately traced to stray capacitance between two overlapping traces.* Because of the inherent sensitivity of the oscillator's electrical design, any changes in capacitance within critical parts of the circuitry can result in a frequency change. This problem had not been encountered in previous designs, since the base material was considerably thicker: 1.57 mm versus the new base thickness of 0.03 mm. The theoretical model for such a situation can be predicted from the capacitance equation, C=0.08858KA/t, where C is capacitance in picofarads, A is area in cm², t is dielectric thickness in centimetres, and K is dielectric constant. It can be seen that the smaller the dielectric thickness, the greater the capacitance. Therefore, the solution was to decrease the capacitor area, since it was not practical to change any of the other parameters. Flex circuit technology allowed for the local necking down of the trace in the

*An overlapping trace is simply a trace on one side of a circuit board that crosses another trace on the other side.



Fig. 5. Stray capacitance between overlapping traces on the thin flexible substrate is reduced by necking down the overlapping traces. When needed, capacitors are easily formed (circular area without hole).

region of the crossovers, as illustrated in Fig. 5. The neck-downs are approximately 0.13 mm wide instead of the standard 0.38 mm.

Acknowledgments

Richard Liszewski contributed to the design and production implementation of the flex circuit. Jerry Curran and Dave Gottwals contributed to production implementation of the 10811A/B. Credit for the original design concept for the oven housing and flex circuit goes to Bob Wilson.

Reference

1. J. Steinmetz, "Flex Circuitry Packaging of a Crystal Oscillator," National Electronic Packaging and Production Conference, February 1980.



James H. Steinmetz

Jim Steinmetz is a product design supervisor with HP's Santa Clara Division. He joined HP in 1973 after receiving his BSME degree from the University of California at Davis. His responsibilities at HP have included special machine design, manufacturing engineering, package design for displays, and product design for Rubidium frequency references and crystal oscillators. He's authored two papers on flex circuitry and is named inventor on a pending patent on a device package interconnect scheme. He's a registered profes-

sional engineer in California and a member of ASME, NSPE, and CSPE. Jim is married, has a son, and lives in Los Altos, California, not far from Palo Alto, his birthplace. His special interests include rebuilding wrecked cars, remodeling his home, and traveling.

structure, such as the crystal in the oven cavity, two loop gains must be considered: the controller loop gain (partially electrical) and the mechanical gain. The controller loop gain can be defined as the change in heater temperature resulting from a change in thermistor temperature divided by the change in thermistor temperature, assuming the thermistor is separate from the heater. This gain can easily be made very large, and in the 10811A/B Oscillator it is about 10⁵. The mechanical gain is defined as the change in ambient temperature divided by the resulting change in crystal temperature. The mechanical gain is a function of the thermal properties of the structure, the location of the two heaters, and the thermistor location. The difference between mechanical gain and thermal gain is that thermal gain is dependent upon both the controller loop gain and the mechanical gain and can be no larger than the smaller of these two. If it can be assured that the controller loop gain is more than ten times greater than the desired thermal gain, the thermal gain approximately equals the mechanical gain. This is the case in the 10811A/B, since the controller loop gain is designed to be 100 times the desired minimum thermal gain of 1000. Therefore, with the controller defined, it then was necessary to locate the heaters, thermistor, and crystal such that the mechanical gain to the crystal, $\Delta T_{ambient} / \Delta T_{crystal}$, would be greater than 1000.

Normally, when one heater winding is the only source of heat for an oven, the crystal and thermistor must be mechanically rearranged on a trial-and-error basis until the desired mechanical gain is obtained. If later in the design a mechanical change occurs, this may change the thermal behavior of the oven enough that the above process would have to be repeated. With two heaters, adjusting the ratio of the power dissipated in the two heater transistors achieves the same result as physically moving the thermistor or crystal. With the first prototype 10811A/B Oscillator, several trials were required initially to locate the thermistor, crystal, and heaters in the aluminum oven housing. After that, it took several hours to fine tune the oven to a thermal gain of approximately 10,000. This fine tuning was done by adjusting the ratio of resistors R2 and R3 (Fig. 10). Later, when more oscillators were constructed using a mechanical design identical to the prototype, the same ratio of R2 and R3 was used. No further attempts were made to optimize the thermal gain of these other ovens. The thermal gain to the crystal for all of these ovens was measured to be greater than 1000.

Mechanical Aspects

A combination rigid-flexible printed circuit board concept is used in the 10811A/B Oscillator (see page 26). The printed circuit board is constructed of a flexible five-layer KaptonTM and copper laminate cemented to a standard G10 fiberglass stiffener. The stiffener is required only for mechanical support since all traces and pads are contained in the flexible laminate. Since the stiffener boards are initially part of a single large sheet, it is possible to leave small cut-away tabs temporarily connecting the individual boards to the remaining support board. Later on, these tabs are cut with a pair of side cutters to release the boards. While still connected to the support board, the individual boards are loaded with components and wave-soldered in the conventional manner.

During the assembly procedure, no hand wiring is required to interconnect any of the four rigid boards. The edge connector, the oven controller board, and the two boards in the oven are all inherently interconnected with this rigidflexible printed circuit board concept.

Acknowledgments

We would like to thank Mike Fischer, R&D section manager for the Precision Frequency Sources group, for his valuable technical consulting.

References

1. R. Burgoon and R. Wilson, "Design Aspects of an Oscillator Using the SC Cut Crystal," Proc. 33rd Annual Symposium on Frequency Control, pp 411-416, 1979.

2. R. Burgoon and R. Wilson, "Performance Results of an Oscillator using the SC Cut Crystal," Proc. 33rd Annual Symposium on Frequency Control, pp 406-410, 1979.

J. Robert Burgoon



Rob Burgoon is project leader for the 10811A/B Oscillator. He received his BSEE degree in 1968 from California State Polytechnic University at San Luis Obispo and his MSEE degree in 1970 from California State University at San Jose. With HP since 1970, he's served as design engineer and project leader for several precision frequency sources group products. He's a member of the IEEE, has authored four papers and articles, and is listed as an inventor on two pending patents. Born in Norfolk, Virginia, Rob is married, has a daugh-

ter, and lives in San Jose, California. He speaks Norwegian and Spanish and enjoys volleyball, racquetball, and genealogy.



Robert L. Wilson

Bob Wilson received his BS degree in electrical engineering and computer science in 1973 and his MSEE in 1975 from the University of Santa Clara, Joining HP in 1975, he designed a precision crystal oscillator oven, worked on the 10811A/B project, and designed a new aging system for the 10811A/B. He's now production engineering supervisor for precision frequency sources. Bob has co-authored two papers on the SC-cut crystal oscillator and is named inventor on one patent and one pending patent. Bob grew up in Corvallis and Coos Bay, Oregon: He lives in Santa Clara, California and enjoys scuba diving, sailing, and making furniture.

SPECIFICATIONS HP Model 10811A/B Crystal Oscillator

FREQUENCY STABILITY (see definition of terms): LONG TERM (Aging Rate): <5×10⁻¹⁰/day after 24-hour warm-up. See Note 1. <1×10⁻⁷/year for continuous operation.

SHORT TERM: Refer to tables below.

Time Domain	Stability	Frequency Domain Stability				
Averaging Time 7 [seconds]	Stability σ _γ (τ)	Offset from Signal f [Hz]	Phase Noise Ratio $\mathcal{L}(f) \left[dBc/\sqrt{Hz} \right]$			
10^{-3} 10^{-2} 10^{-1} 10^{0} 10^{1} 10^{2}	$\begin{array}{c} 1.5 \times 10^{-10} \\ 1.5 \times 10^{-11} \\ 5 \times 10^{-12} \\ 5 \times 10^{-12} \\ 5 \times 10^{-12} \\ 1 \times 10^{-11} \end{array}$	10 ⁰ 101 102 103 10 ⁴	- 90 - 120 - 140 - 157 - 160			

ENVIRONMENTAL SENSITIVITY:

- TEMPERATURE: $<4.5\times10^{-9}$ over a -55° C to 71° C range. $<2.5\times10^{-9}$ over a 0° C to 71°C range
- OPERATING: -55°C to +71°C.

- STORAGE: -55°C to +85°C. LOAD: $<5\times10^{-10}$ for a $\pm10\%$ change in 50-ohm load. $<5\times10^{-10}$ for a $\pm25\%$ change in 1-kΩ load.
- POWER SUPPLY: <2×10⁻¹⁰ for a 1% change in oscillator supply voltage.
- RIPPLE: <-90 dBc spurs from 10 mV rms ripple on oscillator supply voltage at 100 Hz.

OVEN: $<1 \times 10^{-10}$ for 10% change in oven supply voltage. GRAVITATIONAL FIELD: $<4 \times 10^{-9}$ for 2g static shift (turn-over).

- MAGNETIC FIELD: < -90 dBc sidebands due to 0.1 millitesla (1 gauss) rms at 100 Hz. HUMIDITY (typical): 1×10^{-9} for 95% RH at 40°C.
- SHOCK (survival): 30 g, 11 ms, 1/2 sinewave.
- ALTITUDE (typical): 2×10^{-9} for 0 to 50,000 ft. WARMUP: 10 min. after turn-on within 5×10^{-9} for final value, at 25°C and 20 Vdc. See Notes 1 and 2.

ADJUSTMENT:

COARSE FREQUENCY RANGE: >±1×10⁻⁶ (±10 Hz) with 18 turn control. ELECTRONIC FREQUENCY CONTROL (EFC): ≥1×10⁻⁷ (1 Hz) total, control range -5 Vdc to +5 Vdc.

OUTPUT:

١

FREQUENCY: 10 MHz.

VOLTAGE: 0.55 ±0.05 V rms into 50 ohms. 1V rms ±20%, into 1k ohms.

HARMONIC DISTORTION: Down more than 25 dB from output.

SPURIOUS PHASE MODULATION: Down more than 100 dB from output (discrete sidebands 10 Hz to 25 kHz).

POWER REQUIREMENTS:

OSCILLATOR CIRCUIT: 11.0 to 13.5 Vdc. 30 mA typical. 40 mA max.

OVEN CIRCUIT: 20 to 30 Vdc; turn-on load is 42 ohms minimum. Steady-state power drops to a typical value of 2 0W at 25°C in still air with 20 Vdc applied.

CONNECTORS:

10811A: Mates with CINCH 250-15-30-210 (HP 1251-0160) or equivalent (not supplied). 10811B: Solder terminals and SMB Snap-on connectors. Mates with Cablewave Systems, Inc. #700156 or equivalent (not supplied).

SIZE: 72 mm × 52 mm × 62 mm (2-13/16 in × 2-1/32 in × 2-7/16 in, 14 cu in). WEIGHT: 0.31 kg (11 oz)

PRICES IN U.S.A.

Quantities	10811A	10811B		
1-4	\$800	\$900		
5-9	\$768	\$864		
10-24	\$736	\$828		
25-49	\$672	\$756		

DEFINITION OF TERMS:

- LONG-TERM FREQUENCY STABILITY is defined as the absolute value (magnitude) of the fractional frequency change with time. An observation time sufficiently long to reduce the effects of random noise to an insignificant value is implied. Frequency changes due to environmental effects must be considered separately
- TIME DOMAIN STABILITY $\sigma_{\rm V}(\tau)$ is defined as the two-sample deviation of fractional frequency fluctuations due to random noise in the oscillator. The measurement bandwidth is 100 kHz
- FREQUENCY DOMAIN STABILITY is defined as the single sideband phase noise-tosignal ratio per Hz of bandwidth (a power spectral density). This ratio is analogous to a spectrum analyzer display of the carrier versus either phase modulation sideband.

See "NBS-Monograph 140" for measurement details.

NOTES:

- 1. For oscillator off-time less than 24 hours,
- 2. Final value is defined as frequency 24 hours after turn-on.
- MANUFACTURING DIVISION: SANTA CLARA DIVISION

5301 Stevens Creek Boulevard Santa Clara, CA 95050 U.S.A.

New Temperature Probe Locates Circuit Hot Spots

Use it with any general-purpose digital multimeter and some HP oscilloscopes to get readings directly in degrees Celsius.

by Marvin F. Estes and Donald Zimmer, Jr.

AST, ACCURATE TEMPERATURE measurements are needed in a wide variety of thermal design, diagnostic, and testing applications. Hewlett-Packard's new Model 10023A Temperature Probe (Fig. 1) is designed to provide these measurements. With the probe, surface temperature measurements are read directly in degrees Celsius on any general-purpose digital multimeter (DMM) having an input impedance of 10 megohms or more. The pencil-like probe tip easily accesses small components and a press-to-read switch makes measurements easy.

The probe is a self-contained temperature-to-voltage transducer with a forward-biased diode chip bonded to a small ceramic substrate in the probe tip. A calibrated, linear output of 1 mV/°C is assured by individually characterizing each diode in a precision thermal reference bath. An integrated-circuit resistor network is then laser trimmed to match each diode to its electronic compensating circuit.

The use of Hewlett-Packard integrated circuits permits the entire electronics assembly, including the battery, to be packaged in the probe barrel. A standard dual banana plug output connector is compatible with most digital voltmeters including the built-in DMMs on Hewlett-Packard's Option 034/035 1700 Series Oscilloscopes.

Measurement accuracy of the temperature probe is $\pm 2^{\circ}$ C from 0°C to $\pm 100^{\circ}$ C decreasing linearly to $\pm 2^{\circ}$, -4° C at -55° C and to $\pm 4^{\circ}$ C, -2° C at $\pm 150^{\circ}$ C. This accuracy specification is traceable to the U.S. National Bureau of Standards. For applications requiring relative rather than absolute



Fig. 1. Model 10023A Temperature Probe gives fast indications of surface temperatures on many general-purpose digital multimeters and HP 1700 Series Option 034/035 Oscilloscopes.



Fig. 2. Simplified electrical analog of the problem of measuring the temperature of a transistor with a probe.

measurement of similar temperatures, the probe has a short-term repeatability of ± 0.3 °C.

By using a temperature sensor with very low thermal mass and a geometry that places it very close to the measurement surface (approximately 0.25 mm or 0.010 in), the 10023A can make measurements quickly and closely track heating or cooling in devices operating under varying load conditions. This design also permits temperature measurements with very low thermal gradient errors.*

Very fine wires, approximately 0.10 mm in diameter, are used for connecting the diode sensor to the electronic compensating circuits to achieve very high thermal isolation. This thermal isolation reduces the tendency for the probe tip to act as a heat sink or cooling fin and change the measured surface temperature. Minimum disturbance of the operating environment by the probe is particularly important when accurate temperature measurements of small electronic components are needed.

The electrically isolated probe tip permits measurements *HP Application Note 263 contains typical performance characteristics.



Ceramic Chip in Probe Tip

Fig. 3. Simplified schematic of the 10023A Temperature Probe.

of non-grounded components such as the case of a power transistor with its collector common to the case. A very low tip-to-ground capacitance of approximately 0.5 pF avoids electrical loading when thermal probing in high-speed circuits.

The temperature probe is calibrated at the factory and does not require periodic calibration. There are no internal adjustments. In the event of tip damage, the probe is easily repaired with a new, precalibrated tip and matching compensating network.

Tip Heat Characteristics

The tip is the most important part of a temperature probe, especially one designed to measure the temperature of small objects like transistors, diodes and IC packages. Therefore the HP 10023A temperature probe's development centered around the development of a tip.

Three heat characteristics are particularly important. One is the thermal capacity of the probe tip. If the capacity is too large, even if there are no heat leaks, the probe will take a long time to reach a stable temperature when measuring the temperature of a small transistor. To keep the thermal capacity low, the 10023A tip is constructed of a piece of alumina 2 mm in diameter by 0.25 mm thick, and the diode chip $(0.25 \times 0.25 \times 0.02 \text{ mm})$ is bonded directly to the back of the alumina. The alumina contributes essentially all of the thermal capacity. A small amount is added by the black plastic barrel and the wires that connect the diode chip to the circuitry.

Another important characteristic is the thermal resistance from the measured body to the sensor or diode chip. This is important for two reasons. First, too much thermal resistance, like too much thermal capacity, produces too long a time constant when a measurement is taken. Second, this resistance can affect the accuracy by acting as a heat divider with the thermal resistance from the heat sensors to ambient. The tip resistance was minimized by reducing the thickness of the alumina substrate to 0.25 mm.

The third important heat characteristic is the thermal resistance from the heat sensor (diode chip) to ambient. This is important because it can act as a heat divider with the thermal resistance from the measured body to the sensor and thereby affect accuracy. This resistance effect was minimized by using #38 AWG copper wires from the diode to the probe circuits and by using a plastic barrel to hold the assembly together.

These various thermal capacities and thermal resistances are shown in simplified electrical analog form in Fig. 2. It is obvious that in the steady state R_1 should be low and R_2 should be high to provide good accuracy. Also in the steady state, R_A of the source should be much less than R_2 or the temperature of the transistor will be lowered by thermal loading. C_1 should be much less than C_A or the time constant will be affected. One can also see that if a cold temperature probe is touched to a hot transistor the temperature of the transistor will be lowered until the heat source of the transistor can charge capacitor C_1 . There are two time constants, R_1C_1 and R_AC_A .

Probe Circuit Design

The circuit of the 10023A Temperature Probe (Fig. 3) is very simple. It consists of a bridge with the sensor diode chip as one leg. The temperature is measured by sensing the voltage across one of the middle resistors.

Consider the simple diode equation

$$i_0 = K_1 e^{-K_2 V D/T}$$

where K_1 and K_2 are constants, i_0 is the forward current through the diode, and V_D is the voltage across the diode. Now, if i_0 is constant then $T = K_3 V_D$, where K_3 is a constant. Errors arise because we do not have a perfect current source and because the simple diode equation is not exact. The error curve of the voltage out of the temperature probe is shown in Fig. 4.

The temperature probe's three main circuit components are the power supply (battery), the alumina substrate with the microcircuit, and the probe tip, which contains the temperature sensing diode on the small ceramic chip. During one phase of manufacture the temperature-vs-current characteristics of each individual diode are measured. Using this information plus the value of the voltage regulator in the microcircuit, the resistor values are calculated and laser trimmed. The microcircuit and probe tip are now a calibrated matched pair. Since these two components, the probe tip and substrate, contain all the circuits that affect calibration of the probe and they contain no user adjustable or variable components, the probe does not require periodic calibration.

All of the probe electronics are on a printed circuit board that measures 10 by 61 mm. Power comes from a zinc/ silver-oxide battery located on the printed circuit board. The battery, which has a nominal voltage of 1.5V, is a



Fig. 4. Typical 10023A error curve with possible unit-to-unit variations (shaded band).

Marvin F. Estes



Marvin Estes received his BSEE degree in 1965 from Purdue University and his MSEE in 1966 from Clarkson College of Technology. He's also done further graduate work at North Carolina State and North Colorado State Universities. Before coming to HP, he worked on magnetic effects and capacitive energy storage with the U.S. National Aeronautics and Space Administration. A member of IEEE, he's the author of a paper on coupled striplines and is named inventor on a patent on a new way of producing artificial diamonds. As an HP circuit designer for eight

years, he's developed a trio of voltage and temperature probes, the latest being the 10023A. Born in Warren, Pennsylvania, Marvin is married, has four children, and lives on five acres in Black Forest, Colorado with a variety of dogs, chickens, bees, cats, rabbits, and goats. Besides animals, he's also interested in electromagnetic applications of random walk techniques.

Donald Zimmer, Jr.



Don Zimmer received a BSEE degree from the University of Minnesota in 1966 and a BA degree in chemistry from the University of Colorado at Colorado Springs in 1976. He joined HP in 1969 and worked as a CRT production engineer for five years and as a development engineer for two years before becoming a process development engineer in thick-film R&D, Later, as a member of the thick-film product engineering group, he developed the thick-film microcircuits for the 10023A Temperature Probe, and he's now doing integrated circuit R&D. A native

of Minnesota, Don is married and has two children. Among his interests are skiing, tennis, jogging, and reading.

low-cost, commercially available type commonly used in hearing aids and wristwatches. It provides an average of 50 hours of probe operation.

Since the probe electronics are on a printed circuit board housed in a plastic case, the components on the board are

SPECIFICATIONS HP Model 10023A Temperature Probe

MEASUREMENT RANGE: -55°C to +150°C.

OUTPUT: 1 mV/°C

SHORT-TERM REPEATABILITY: ±0.3°C (minimum of 48 hr). ACCURACY: ±2°C from 0°C to 100°C, decreasing linearly to +2°C, -4°C at -55°C and

+4℃, -2℃ at +150℃

MAXIMUM VOLTAGE AT TIP: 600 V (dc + peak ac).

TIP CAPACITANCE TO GROUND: approx. 0.5 pF.

THERMAL RESPONSE: <3 s to settle within 2°C of final reading (liquid measurement for a 100°C temperature change.

DMM INPUT R: ≥10 MΩ.

OPERATING ENVIRONMENT (probe tip to approx 13 mm (0.5 in) from probe tip): Temperature, -55°C to +150°C; altitude, to 4600 m (15,000 ft); vibration, vibrated in three planes for 15 min. each with 0.38 mm (0.015 in) excursion, 10 to 55 Hz.

OPERATING ENVIRONMENT (probe body): Temperature 0°C to 60°C (battery limitation); humidity (non-condensing), to 95% relative humidity at +40°C; altitude and vibration same as those for probe tip.

OVERALL LENGTH: Approx. 1.4 m (53 in).

WEIGHT: net 85 g (3 oz); shipping, 312 g (11 oz).

BATTERY LIFE: Approx. 50 hr (varies with ambient temperature).

LOW-BATTERY INDICATION: Probe output indicates approx. -70°C on DMM. First indication of a low-battery condition is a decreasing indication of 1° to 2°C/minute with probe tip at a constant temperature.

ACCESSORIES SUPPLIED: One replacement battery (1420-0256), one sliding lock collar (10023-23201), and one probe tip cover (00547-40005)

PRICE IN U.S.A.: Model 10023A Temperature Probe, \$150. Replacement Tip (includes precalibrated tip and matching compensation network. P/N 10023-60001), \$65.

MANUFACTURING DIVISION: COLORADO SPRINGS DIVISION 1900 Garden of the Gods Road

P.O. Box 2197 Colorado Springs, Colorado 80901 U.S.A.

susceptible to static discharge. The conventional method of protecting them is to offer a ground path for the static charge. Because the 10023A is a floating instrument and is not grounded, a different method had to be devised. By use of three capacitors on the circuit board and a conductive coating inside the plastic housing, the static charge is drained and lowered through the probe and instrument to prevent damage to the temperature probe and the attached instrument.

Acknowledgments

We wish to thank Larry Gamill for starting this project. Rick James for finishing it, and Stan Lang for his many helpful suggestions. We are also grateful to Don Skarke for his help during the final stages of the project.

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

HEWLETT-PACKARD JOURNAL

MARCH 1981 Volume 32 . Number 3

Technical Information from the Laboratories of Hewlett-Packard Company Hewlett-Packard Company, 1501 Page Mill Road Palo Alto, California 94304 U.S.A.

Hewlett-Packard Central Mailing Department Van Heuven Goedhartlaan 121 1180 AM Amstelveen The Netherlands Yokogawa-Hewlett-Packard Ltd., Suginami-Ku Tokyo 168 Japan

CHANGEOFADDRESS: To change your address or delete your name from our mailing list please send us your old address label. Send CHANGEOFADDRESS: changes to Hewlett-Packard Journal, 1501 Page Mill Road, Palo Alto, California 94304 U.S.A. Allow 60 days.

0200032503&&&HARR&

CHAYO

PENSACOLA

R JULIAN A HARRIS HAYO ELECTRONICS O BOX 2807

Bulk Rate US Postage Paid Hewlett-Packard Company

HP Archive

This vintage Hewlett-Packard document was preserved and distributed by

www.hparchive.com

Please visit us on the web!

On-line curator: John Miles, KE5FX jmiles@pop.net