

NOVEMBER 1980


HEWLETT-PACKARD JOURNAL



Contents:

- 3 Patient Monitoring Enhanced by New Central Station**, by Timothy B. Blancke and Larry L. Nielsen *Critically ill patients get better care with a new patient information center that also monitors itself.*
- 11 High-Speed Raster Technique Provides Flexible Display**, by Robert L. Stettiner and George L. Adleman *Smooth waveforms on a raster-scan CRT display are made possible by a combination of higher scan rate and beam width modulation.*
- 15 Multi-Processor Architecture and Communications for Patient Monitoring**, by James M. Rueter *Four processors use three different communication links to share tasks and accomplish the overall monitoring function.*
- 19 Self-Test and Serviceability for Dependable Central Patient Monitoring**, by Jeffrey M. Goldberg *This system automatically analyzes itself for failures, reconfigures its units to continue operating when one unit fails, and notifies the operator of any malfunction.*
- 23 Firmware for a Patient Monitoring Station**, by Kim L. Hanna *A matched-filter design accurately analyzes electrocardiogram waveforms. Data is displayed and recorded using a combination of communication systems.*
- 29 An Interactive HP 3000/IBM Mainframe Link**, by Connie J. Ishida *This 3270 emulator software eases the design of distributed processing networks consisting of HP 3000 systems and IBM mainframes.*

In this Issue:



Critically ill patients recovering from surgery, heart attack, or serious illness are often placed in special intensive-care units where their vital signs can be watched constantly. In these units, electronic bedside monitors attached to each patient relay vital-sign information to a central station or patient information center at the nurses' station. The patient information center assists the hospital staff in three ways. First, it generates audible and visual alarms if preset vital-sign limits are exceeded. It's important that the central station announce these emergencies without generating too many false alarms—for example, when the patient moves. The second function of the patient information center is to display the vital-sign data.

By watching this data, the staff can detect problems before they reach the alarm stage. Trend plots may also be provided by the central station to help guide the patient's treatment. The central station's third function is to provide recordings of the electrocardiogram and sometimes of blood pressure. Especially important is a record of the few seconds just before an alarm, which shows what kind of irregularity led to the alarm.

Most of this issue is devoted to the new Hewlett-Packard 78500 Series Patient Information Centers. These new systems make many significant contributions to special patient care, including simplified operation, sharper displays, noticeably reduced false alarms, a waveform memory that stores patient data when several alarms occur simultaneously, three waveform recording choices for each patient, optional trend displays and improved pressure processing, and four levels of self test. Not only does this new equipment test and diagnose itself, but when some parts fail, others take over their functions to minimize the impact of the failure on system operation. The 78500 Series is compatible with all existing HP bedside monitors, and thanks to the decreasing cost of smaller, more powerful integrated circuits, it costs no more than HP central stations introduced ten years ago. The 78500 Series is compact and can be installed either in HP cabinets or in cabinets of the hospital's own design. Our cover shows the 78500 installation at Woodland Memorial Hospital in Woodland, California.

The article on page 29 discusses the design of a software package that allows HP 3000 Computers to communicate easily with large IBM mainframe computers. Most small-computer manufacturers offer similar programs, called "3270 emulators" because the large mainframe computer thinks it's talking to an IBM 3270 terminal control unit. IML/3000 represents a real contribution to this field. It takes care of the low-level details and makes it easy to set up HP-to-IBM program-to-program communications.

R. P. Dolan

Patient Monitoring Enhanced by New Central Station

Multi-microprocessor architecture and a new integrated display concept provide more patient data, fewer false alarms, internal self-tests and extensive recording capabilities in an easy-to-operate system.

by Timothy B. Blancke and Larry L. Nielsen

THE NEW 78500 SERIES Patient Information Centers (PICs), Fig. 1, provide a wide range of patient monitoring data at the nurses' station of patient-care units. The combination of a multi-microprocessor architecture and a new display technology makes it possible to display a flexible mixture of smooth waveforms, alpha-numeric, and graphics on a single cathode-ray tube (CRT). This saves valuable space at the nurses' station and presents all the information at a glance.

What PICs Do

Patient monitoring systems usually consist of two types of equipment—bedside monitors and central station monitors. Bedside monitors are connected to various transducers that measure the physiological status of the patient. For example, an extremely sick patient in a surgical intensive care unit might be monitored for:

- Heart rate, rhythm and wave shape (ECG)

- Arterial, pulmonary and venous blood pressures
- Blood flow (cardiac output)
- Internal body temperature
- Respiration rate
- Respiratory carbon dioxide.

As the patient progresses towards good health the number of monitored variables is decreased for patient comfort as quickly as practical because some of these measurements require the insertion of tubes or wires into the patient. Eventually, only the electrical heart functions (ECG) are measured using electrodes attached to the patient's chest. If the patient is ambulatory, a small portable telemetry transmitter may be used to send the patient's ECG to a central station for monitoring.

Many hospitals have several kinds of care units (e.g., surgical intensive-care unit, coronary-care unit, and progressive-care unit). These units use bedside monitors of varying degrees of complexity. A patient getting better



Fig. 1. The 78500-Series Patient Information Centers are advanced central monitoring systems designed for user convenience in a hospital environment. These modular microprocessor-based systems have increased reliability through comprehensive self-testing and can be easily expanded to meet future monitoring requirements.

might be moved from one kind of unit to another. A patient who has suffered a heart attack might spend five days in a coronary-care unit and then ten more days in a progressive-care unit.

The central station monitoring equipment receives information from the bedside monitors and uses it for three purposes:

1. Alarms. Some variables, such as heart rate, can change very quickly, putting the patient's life in jeopardy unless corrective action is taken within three or four minutes. Audible and visual annunciation of these events is provided by the central station. Often, false alarms are sounded—generally because the patient moves, causing noisy signals (artifacts) to be sent to the PIC. The number of false alarms can be significantly reduced with modern signal-processing techniques.
2. Observation of patient status. Many variables need to be observed to prescribe the most appropriate care for a patient. Some, such as the ECG wave, may also be used to check on the accuracy or the urgency of the automatic alarms. The ECG wave is watched to spot developing irregularities that are not drastic enough to cause an alarm, but need to be treated before they do.

While it is important to know the instantaneous vital signs to treat life-threatening conditions, it is also good medical practice to observe the trends in the patient's vital signs and to correct them before anything serious develops. This is sometimes called "giving the patient a smooth ride." Many trends are very slow, requiring several hours for a vital sign to change significantly. Trend plots of vital signs can be provided by the PIC to aid in guiding the patient's therapy.

3. Documentation. Most central stations provide a strip-chart recording of the ECG. When an alarm occurs, the recorder automatically provides documentation of the alarm event, and records the patient's ECG that caused the alarm. A waveform storage element delays the ECG wave for about ten seconds so that it is possible to record the momentary event that triggered the alarm. Manually requested recordings are often made every few hours and compared with earlier ECG recordings so that subtle changes can be detected as they develop.

In some cases, it is helpful to make a combined two-channel recording of the ECG and blood-pressure waves, because the combination gives an overall picture of the heart's electrical and mechanical activity.

Also, if the central station has the capability to display vital signs or trend plots, it is convenient to have a hard-copy device to provide a permanent record for the patient's files.

The purpose of all PICs is to provide a central location for observing and recording the physiological data and alarm information from a number of patients. If a PIC is to be effective it must fulfill this purpose for all types of hospital users. Physicians, nurses, clerks, and support technicians all play a role and must feel that it is easy to enter, retrieve, and use desired information. To be viewed as a friend and not as an adversary, the system must respond quickly to user inquiries while maintaining a constant surveillance and display of essential data and alarms. The PIC must be well-suited for the hospital environment by not requiring

too much floor space, special temperature and humidity controls, and special power. The system must be safe, quiet and aesthetically pleasing.

The 78500 PICs

Efficient and economical handling of the wide range of patient monitoring data is achieved by the design of two basic versions of the new HP 78500 Series PICs. The 78501A PIC provides economical monitoring, display and recording capabilities for up to eight patients. It is intended primarily for cardiac patients found in coronary-care, progressive-care and telemetry units. The 78502A PIC provides more extensive display and recording capabilities, as well as trend plots and hard-copy printing, but for only four patients. It is specially designed for higher-risk cardiac patients and those suffering with both cardiac and hemodynamic dysfunction. These patients are typically found in intensive-care and surgical-care units. When it is necessary to monitor more patients, up to four 78501A/78502A systems can be combined to share common recording and hard-copy facilities.

In addition to using much of the same hardware, both systems are designed to place the observation, annunciation and documentation functions at the user's fingertips. All normal observation and documentation functions are provided by the use of eight softkeys and two dedicated control keys. These keys are conveniently located along the right-hand portion of the CRT display. The softkeys are controls that have different functions at different times. Labels describing the current functions of all active softkeys are always shown on the CRT screen next to the keys. Only occasionally will the door on the display unit need to be opened to allow access to six additional dedicated controls.

Observation

The 78501A PIC monitors and compactly displays the ECG waveforms and heart-rate information for eight patients. The primary or standard display used for observation is shown in Fig. 2. This data format is displayed after power-on and whenever the **STANDARD** key is pressed. This format shows four seconds of real time ECG waveform, the accompanying heart-rate or alarm information, and, if




Fig. 2. Standard monitoring display for 78501A system shows cardiac data for up to eight patients.




Fig. 3. Standard monitoring display for 78502A system shows cardiac and hemodynamic data for up to four patients.

blood-pressure monitoring is active, the corresponding pressure parameter (systolic/diastolic pressure for the P1 pressure site and mean pressure for the P2 pressure site).

The 78502A PIC has many displays that are specially structured for the presentation and observation of data. Like the 78501A, the standard display (Fig. 3) is normally the resting display. Six seconds of real-time ECG waveform and a large-character display of the associated heart rate (and

VPB* rate if arrhythmia monitoring is available) are shown for each active patient. Whenever they are available, the systolic and diastolic blood pressure values are also given. Like the 78501A, the alarm information is given above the respective ECG waveform display area.

The display shown in Fig. 4a is the result of pressing the **FREEZE** key on the 78502A PIC. For this display and all 78502A PIC displays other than the standard display, the ECG waveform duration is reduced from six seconds to four seconds. This is done to free the right-hand portion of the display for other display purposes, while retaining the important real-time monitoring activity. Thus, the real-time waveforms and alarm information are constantly available for observation regardless of the display of optional information. The right-hand portion of the freeze display is used to display either the first or last four seconds of eight seconds of frozen ECG data. This allows study and/or documentation (strip-chart recording) of any abnormal heart beats that may have been observed by the user.

The remaining displays that are useful for observation are accessible by using the vital-signs (VS) softkeys. A vital-signs display is shown in Fig. 4b and is designed to give a listing of all electronically gathered parameters and associated alarms for the selected patient. In addition to providing a quick glance at all parameters, an in-depth look at

*VPB stands for ventricular premature beat, also called PVC—premature ventricular contraction.

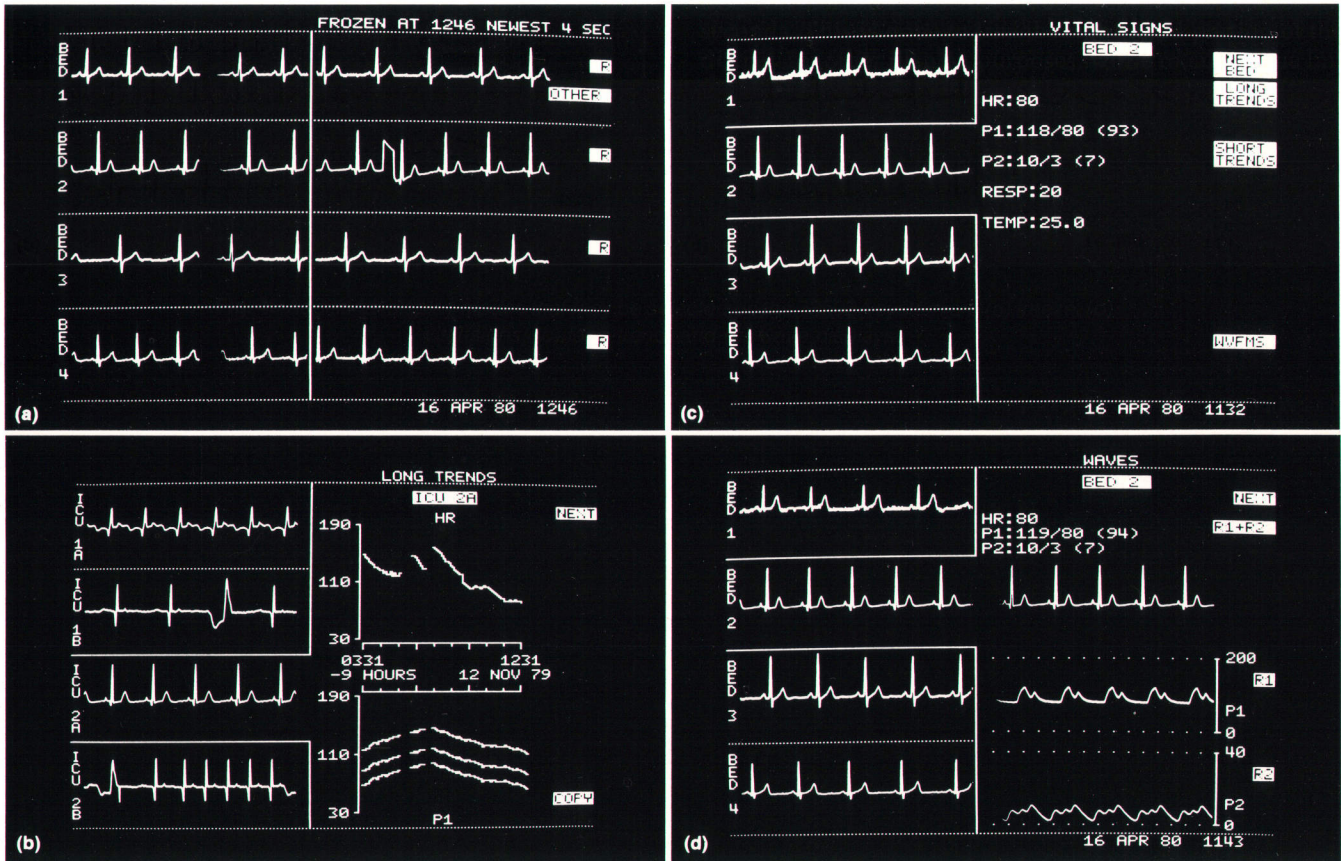


Fig. 4. Additional display modes that may be selected on the 78502A system are (a) freeze, (b) vital-signs, (c) long-trend, and (d) wave displays.

the past and the present data for the selected patient can be obtained via the softkeys within the vital-signs display. They effectively open the right-hand portion of the display into a series of detailed physiological displays that can be evaluated to guide the care of critically ill patients.

A long-trend display, shown in Fig. 4c, is representative of all trend displays. Either long trends (soft-configured for either 9 or 24 hours) or short trends (90 minutes) may be selected for observation and/or documentation.

The last display within the detailed physiological-display group is the wave display (Fig. 4d). In this display the ECG and up to two other real-time waveforms for the selected patient can be displayed.

Other displays that are used occasionally are the control displays. The first of these is shown after the dedicated control key is pressed. The cardio-control display is then obtained for a selected patient by pressing a softkey. This display is shown in Fig. 5 for the 78501A and is representative of both PICs. The softkey definitions for this display mode allow adjustment of the cardiotech alarm limits. When the patient's ECG is telemetered, controls for the waveform size and bed on/off functions are provided as illustrated. Another softkey turns alarms on or off for that patient.

Annunciation

The annunciation of alarms and inops (equipment malfunctions) is accomplished by (1) lighting an alarm lamp, (2) writing a detailed alarm message in the patient's common-display area, (3) sounding an appropriate alarm and (4) optionally closing an alarm relay to trigger any additional alarm indication such as a light in the hallway outside the patient's room. A typical 78501A PIC display with alarms and inops is shown in Fig. 6.

The alarm lamps (red, yellow and green) are located in the upper right-hand corner of the PIC display unit. More than one lamp will light if alarms of different levels exist for different patients. Three sounds are used to announce alarms. They vary, in order of decreasing priority level, from a continuous chime every second (red patient alarm), to a chime every two seconds for a period of fifteen seconds




Fig. 5. The cardio-control display mode is used to change the alarm limits for the monitored parameters.

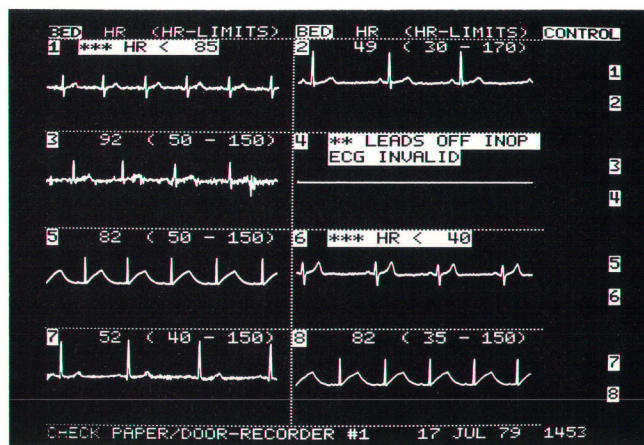


Fig. 6. When an alarm or malfunction (inop) condition occurs, appropriate messages are displayed. A display giving typical alarm and advisory messages for the 78501A system is shown.

(yellow patient alarm), and finally a continuous beep every two seconds (yellow inop). Only the highest-priority-level alarm will sound at any one time.

Documentation

In addition to the audible and visual alarm indications described above, a one or two-channel strip-chart recording of the alarm waveform(s) is automatically initiated. If all recorders are already busy recording earlier requests, a snapshot of the alarm waveform is stacked (temporary storage of the data within internal memory) until a recorder becomes available. At that time, the waveform(s), heart rate, time-of-alarm and alarm message are recorded as if they had been recorded at the time of the alarm.

Manual recordings can be made for all active beds and the record requests are queued in the order requested. The recordings may be activated by the R softkey within the standard, freeze and wave displays or at the patient's bedside. The recordings may be direct (real-time), delayed, timed-run, or continuous (real-time).

In addition to the single or dual-channel strip-chart recordings, a graphics printer may be used to provide permanent recordings for patient files. A copy of the vital signs display or any trend display may be obtained by depressing the COPY softkey on the appropriate display. Also, a periodically scheduled or manually requested trend report may be obtained for all long trends of a selected patient.

Instrument Building Blocks

A high degree of flexibility is provided so that systems can be configured to satisfy a wide range of patient-care needs. The basic set of instruments includes:

- 78510A Display Module
- 78511A Equipment Cabinet
- 78571A Single-Channel Recorder
- 78572A Dual-Channel Recorder
- 9876A Thermal Printer.

An example of a configured system is shown in Fig. 7. A 78510A and 78511A configuration can be used to accommodate either four heavily instrumented patients, or eight lightly monitored patients. In large care units, as many as

four 78510A/78511A combinations can be interconnected to share a documentation bank of a printer and up to three strip-chart recorders. On the other hand, as many as 24 patients (three 78510A/78511As) can share a single recorder. This is possible because the recorder burden is minimized by an algorithmic heart-rate processor that produces far fewer false-alarm recordings, and by a recorder memory facility that automatically stores alarm recordings when the recorder is already busy.

To conserve space at the PIC the equipment cabinets (which have no controls) can be located as far as 60 metres from the displays and recorders. In addition, the 78510A Display can be mounted in a variety of positions—on the wall, from the ceiling, or on a swivel mount.

The principal difference between the four and eight-patient systems is the configuration of processor subsystems within the equipment cabinet. The eight-patient system uses two identical bedside-interface subsystems (one each for four patients) while the four-patient system uses only one bedside interface and substitutes a trend subsystem for the second bedside interface. The four and eight-patient systems each use similar display and communications subsystems.

The heart of each subsystem is the common microprocessor board. This board contains a custom 16-bit silicon-on-sapphire (SOS) microprocessor, sockets for eight read-only memories (ROMs), and a custom HP-IB* interface chip which makes it possible for a processor to talk with other processors. By simply plugging in a different set of ROMs, a single processor board can be made to perform the processing functions for any one of the four subsystems.

The commonality of instruments and processor subsystems within the equipment cabinet has produced many benefits. Less hardware had to be designed. The use of common hardware decreases manufacturing costs and minimizes the number of spare parts that need to be stocked in the field for service. And, if a hospital's monitoring needs become more sophisticated, it is easy to upgrade a 78501A system to a 78502A on-site.

Signal Flow in the Eight-Patient System

In the 78501A two bedside interfaces collect the information from eight bedsides, digitize the data and analyze the ECG signals to produce heart-rate signals and alarms. This information is passed on to the display subsystem and the communication subsystem using the HP-IB. The display subsystem stores this information for delayed recordings, formats waveforms, alphanumerics and graphics for display, and generates a video signal for the 78510A display. The communications subsystem controls the flow of information on the HP-IB and interfaces the equipment cabinet with the recorders and the display's alarm lights and sounds.

Signal Flow in the Four-Patient System

Signal flow in the four-patient 78502A is similar to that in the eight-patient system, except for the trend subsystem functions. In this case the patient's data is also sent to the trend subsystem. The trend subsystem stores this information so that trend plots can be created for the display and the

optional printer. The same processor also derives systolic, diastolic, and mean blood-pressure values by using digital techniques. These results are more accurate than the pressure values derived by the analog processors in existing bedside monitors. These blood-pressure values are used in vital-signs displays, trend plots, and reports.

The Bedside Interface Subsystem

A single bedside-interface subsystem connects with a cable from each of four bedsides. Each bedside cable contains as many as fifteen analog signals representing physiological parameters which may include several blood pressures, ECG, heart rate, respiration, end-tidal CO₂ and temperature. Several other wires in the same cable carry status information, such as alarm signals. Each bedside also has its own ground wire to maintain a true differential-measurement system. This precaution is taken because significant common-mode ground-loop noise can exist between bedsides that are not supplied by the same local power-distribution-grounding system.

A major goal of the PIC is to make it possible for a nurse to be able to interchange a wide variety of bedside monitors without requiring any reconfiguration of the PIC. Thus, the first task of the bedside interface is to ascertain what kind of a bedside monitor it is connected to and then assign the correct labels to the variables it will receive. This is accomplished by using one of the bedside wires to connect the PIC to a resistor in the bedside monitor. The bedside interface measures the resistance value to identify the type of equipment at the bedside. One bedside interface can be connected to four different types of bedside monitors.

The eighty or so incoming physiological values and status signals are then sampled and digitized at appropriate rates by a 10-bit analog-to-digital converter. ECG waveforms are sampled every two milliseconds to maintain the 0-100 Hz diagnostic bandwidth. At the other end of the spectrum, slowly varying variables such as temperature are sampled every four seconds. All this information is labeled by the bedside interface processor and made available to the other processors through the HP-IB.

A second function provided by the bedside interface processor is the analysis of each patient's ECG wave to provide the corresponding value of heart rate. The heart rate is then compared with alarm limit values set at the PIC. Because the algorithm used to derive the heart rate is quite sophisticated (see article on page 23) the PIC produces accurate data and few false alarms. A field trial with ambulatory patients wearing telemetry transmitters in a progressive-care unit showed that false heart-rate alarms were reduced by a factor greater than five when compared with older analog cardiometers. This reduction in false alarms is significant to the nursing staff, which finds false alarms very frustrating. Another benefit is the reduced consumption of recorder paper, because every false alarm initiates an unnecessary recording.

Display Subsystem

The display subsystem contains most of the intelligent features perceived by the user. The main functions of the display subsystem are to

- Format data for presentation and generate the video signal for the 78510A Display

*Hewlett-Packard's implementation of IEEE Standard 488 (1978).

- Provide system control functions through softkey interaction with the user, and
- Provide a data delay and storage capability for waveform recordings.

Display formats are provided to meet various medical needs. Each format is a mixture of waveforms, graphics, and alphanumeric. It is the formatting power of the display subsystem that makes it possible to present a wide variety of quickly comprehensible information in a limited display area.

The alphanumeric capability of the display also makes the implementation of softkey controls possible. Only eight softkeys are needed to implement the following control functions for four or eight patients:

- Select display frames
- Initiate recordings and operate the printer
- Set cardioteach upper and lower alarm limits
- Reset alarms
- Adjust the amplitude of the ECG waveform from telemetry receiver
- Initiate/terminate trend plots.

The logic to display a suitable set of control labels for every situation is contained within the display processor. Before arriving at a decision concerning the appropriate control labels to be displayed, the processor must also consider various factors which are not constant, such as

- Which bedsides are active?
- What data is available from each bedside?
- Current alarm situation for each patient?
- Availability of one and two-channel recorders?
- Availability of printer?

No labels are provided for inactive controls. Although a formidable amount of logic is required to accomplish softkey control, the small number of user controls and the responsiveness of the display make the system seem friendly and easy to learn. The integration of the control

function with the information display also saves space by concentrating all user activity into one area.

The display subsystem has two output subsections—Superaster for waveforms, and a conventional 300×260 picture-element bit map for alphanumeric and graphics. Waveforms cannot be satisfactorily displayed by a bit map of this size. The medical community expects waveforms to appear as smooth curves, and not to exhibit a “connect-the-dots” roughness. The Superaster display technique was created to overcome this problem. Briefly, the Superaster wave generator uses a 1200-line vertical raster combined with a custom video smoothing technique to produce smooth waveforms of constant brightness, regardless of the slew rate of the waveform. This technique is more fully described in the article on page 11.

The Superaster and bit-map outputs are combined to produce the video for the 78510A display. This makes it possible to present smooth waveforms, write copious alphanumeric data, and plot graphics information anywhere on the screen.

The last function of the display subsection is to manage a 16K-word memory that is used as temporary storage for waveform and hard-copy data. ECG waveform data for each patient is continually stored temporarily to provide a delayed (typically by eight seconds) waveform output for recording. The delay is used to capture a snapshot of any ECG abnormalities leading up to an alarm. The delay is also used by the nurse to record ECG abnormalities seen on the display. The recording delay is longer than the amount of time that a particular wave appears on the screen, giving the nurse plenty of time to capture an interesting event. If the recorder happens to be busy, the display processor will freeze the waveform and save it until a recorder becomes available. If necessary, all patients can have frozen waveforms at the same time.

Since the four-patient 78502A PIC has the same amount

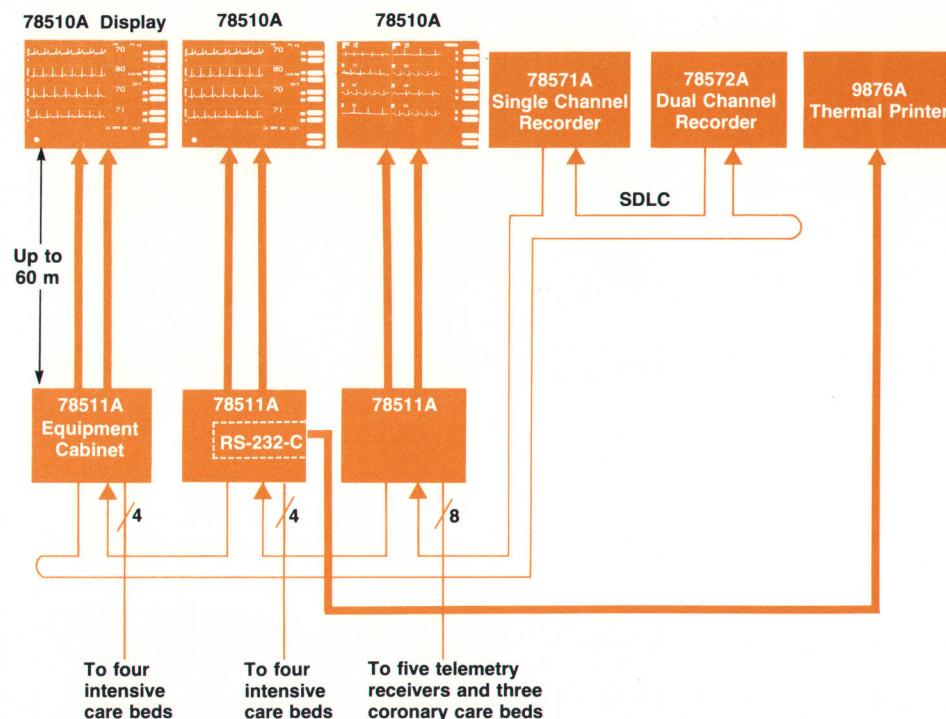


Fig. 7. A typical system application for eight intensive-care patients and eight progressive-care patients sharing one central station location. Two 78502A systems and a 78501A system share two recorders and a hard-copy device. The two-channel recorder is used to record ECG and blood pressure data simultaneously.

of memory but only half the number of patients of the eight-patient 78501A PIC, memory is available for some extra functions.

First, it is possible for the four-patient PIC to store two waveforms (e.g., ECG and blood pressure) for each patient. Second, it is possible to freeze eight seconds of all four patient ECGs for both display and recordings. Finally, a trend plot or vital signs tabulation can be dumped from the screen into temporary storage when a hard copy is desired. This feature allows the user, after requesting a hard copy, to change the display without having to wait until the hard copier has completed its task.

Communications Subsystem

The keystone of the PIC architecture is the communications subsystem. This subsystem performs three functions. First, it allows the bedside interface, display, and trend processors within a single equipment cabinet to communicate among themselves. The second function of the communications subsystem is to allow these processors to communicate with the controls, lights and sounds of the 78510A Display. The third function is to allow several equipment cabinets to share one or more recorders and a printer. In the extreme case, one communications subsystem could be controlling three other internal subsystems, its own display link, and, through the synchronous data link control (SDLC) loop, an additional three equipment cabinets, three recorders and a printer.

System configuration functions are performed automatically whenever power is turned on or if a self-test button is pressed. The automatic configuration feature is important to the user because it adds a "fail-soft" feature. If one of the system elements should fail, the user need only press a self-test button to reconfigure the system with the failed element eliminated. In many cases the built-in real-time self-testing modes will automatically detect a failure and reconfigure the system without any user interaction.

HP-IB Function

HP-IB was chosen for the interprocessor link within an equipment cabinet. The high transfer rate (in this case approximately 200K bytes/s), and the ability to send the same message to several listeners simultaneously make HP-IB an efficient communications medium. The integration of the interface functions by the custom SOS interface chip relieves the processors of time-consuming bus-management tasks, leaving them more time for processing physiological information.

The communications subsystem controls HP-IB communication in a two-step process. First, during an initialization phase, it polls all the HP-IB addresses to determine which processors are present and it then builds a table of the types of messages each processor wants to receive. After the message needs are known, the communications subsystem repetitively polls each internal processor to see if it wishes to talk, setting up the appropriate listeners in response to the message type the talker wishes to transmit.

Communication with the 78510A Display

The 78510A Display contains the alarm lights and sounds as well as the softkeys and other controls used to operate the system. The activation of lights and sounds is done by

messages from the equipment cabinet to the display, while the controls in turn send key-push messages from the display to the equipment cabinet. Since the data rates are slow and only two locations need communicate, this link is simply accomplished by use of two UARTs (universal asynchronous receiver/transmitters).

SDLC Function

The final system communication need is to interconnect as many as four equipment cabinets with as many as three recorders and a printer. The equipment cabinets may be located up to 60 metres from the recorders and printer. Flexible two-way communication paths are needed because any equipment cabinet can use any recorder or the printer, and all the recorders and the printer may be operating simultaneously. At first glance HP-IB would seem to be a suitable medium, but in this application the required distances are three times the maximum allowed by HP-IB. In addition, serial communication is attractive because a smaller, simpler cable can be used.

SDLC* was chosen over a UART serial implementation because SDLC can accommodate communications between many instruments and a controller separated by hundreds of feet. Chips that implement the addressing, control, and error-checking functions of SDLC are available from several vendors.

By connecting the instruments in a loop and sequentially passing control around the loop, it is possible for any instrument to communicate bidirectionally with any other loop member. When it is designated as controller, a communications subsystem can address one or more recorders and the printer to send or receive device control information and data. When its communication needs are complete, the communications subsystem passes control to the next loop member.

Since several equipment cabinets might wish to use the same recorder or the printer at the same time, some means must be provided to distribute the resources equitably among the equipment cabinets. This function is performed by a master equipment cabinet which keeps a list of all loop members, queues all recording requests and makes sure that all high-priority alarms from all equipment cabinets are recorded before low-priority alarms. The master also maintains a queue for printer requests.

Any communications subsystem within the loop can be the master controller. During the initialization period the first to come on line assumes this job. Thus instruments can be added to and deleted from the loop and the entire system will automatically reconfigure itself.

The SDLC loop operates at an instantaneous rate of 333,000 bits/s. When the recorders and printer are all busy the average data rate is about 50,000 bits/s. Thus, in the extreme case, a communications processor spends about 15% of its time dealing with data on the SDLC loop.

Trend Subsystem

The trend subsystem consists of a standard microprocessor printed circuit board (with trend-processing ROMs) and a special memory board. Since the trend memory can contain as much as 24 hours of patient data, it is important

*Synchronous data link control, developed by IBM.

that this data not be lost as a result of a power failure. The trend memory board contains 6K words of CMOS RAM, which requires very little power in the standby mode. The CMOS memory is connected to a battery backup power supply which will hold patient data for at least 24 hours after a power failure.

Design for Service

From the earliest design stages special consideration was given to a service strategy. This plan provides for two levels of service—board-exchange and component-level repair. Either level can be implemented by Hewlett-Packard or hospital service personnel.

Board exchange is accomplished without the need for sophisticated test equipment through the use of built-in self-test software. About 25% of the firmware in each processor is used only for self-test. A small amount of test circuitry is also built in to facilitate the firmware-driven tests. An example is a video-sync detector which can be read by the display processor. If the sync generator should malfunction, this test indicates that the problem is in the equipment cabinet, rather than the display.

The self-test function checks:

- All digital communication links
- Video sync
- ROM and RAM
- Display lights and sounds
- Microprocessor functions.

If an error is found, a specific error code is displayed and printed on the recorder. In addition, test patterns are generated for the display and recorder. These tests point to which instrument is at fault, which subsystem is at fault, and often which printed-circuit board is defective. Thus, it is generally possible to quickly perform board-exchange-level repair without requiring any test equipment.

Component-level repair is accomplished through the use of a specially-designed analyzer, the 14451A Service Board. This analyzer can be added to any processor subsystem. It then takes over control of the subsystem and provides a series of special test programs which exercise the circuitry in a known way. An oscilloscope and a digital signature analyzer are then used to probe circuit nodes. The results seen at any node can be compared with a reference waveform picture or signature for that node. The reference is provided in the component-level service documentation. The nodes are probed until a device with proper inputs but an incorrect output is found. The details of the service strategy are covered in the article on page 19.

Fail-Soft Redundancy

When a failure occurs, it is often possible to operate the viable parts of the system and avoid a total system "crash." This is done by breaking the system into subgroups and defining which remaining groups are viable. For instance, if the communications processor fails, HP-IB communications will cease. However, the display could still function if the display subsystem were to take over the job of HP-IB controller. This reconfiguration is automatically accomplished by the built-in self-tests. When the display subsystem decides that the communications subsystem is inactive the display subsystem activates a special fail-soft program which causes it to take over control of the HP-IB

communications with the bedside interface subsystem. Then data can again flow from the bedside interface subsystem to the display subsystem. The display will also show a message which says "Backup-No Recorder No Control."

Similar backup modes have been provided for the loss of other internal subsystems; and, on a larger scale, for the loss of equipment cabinets, displays, recorders, and the hard-copy printer. In each case, it is the combination of the built-in self-tests and the automatic reconfiguration process that makes the backup modes possible.

Acknowledgments

The patient information centers were a very large project and it is impractical to list here all of the people who made contributions. We thank Tom Horth and Lew Platt for their useful suggestions and help by encouraging the application of softkey and Superaster concepts. Larry Hoffman, Doug Ritchie, and Tony Matheson deserve thanks for creating the new styling of the instrument modules and the console. Many other technical contributors are acknowledged in the accompanying articles.

Besides the engineering department, the advice of three other groups was used to fine-tune the project objectives. A

Timothy B. Blancke



Tim Blancke has a BEE degree awarded by Cornell University in 1954. He came to HP in 1968 after working on maritime cathodic corrosion prevention systems and precision ac voltage and phase angle calibration standards for several years. Since then he has worked on medical signal conditioners and multi-channel recorders, and is now a section manager responsible for various patient monitoring products. During his military service Tim was a first lieutenant in the U.S. Army Signal Corps. He is a member of the IEEE and the Association for Advancement of

Medical Instrumentation. A native of Montclair, New Jersey, Tim now lives in Concord, Massachusetts. He is married, has two children, and enjoys swimming, tennis, music, and boating along the nearby Atlantic coastline.

Larry L. Nielsen



Since joining HP in 1966, Larry Nielsen has contributed to or led the development of several medical instruments for training, diagnosis, and monitoring. Now he is a project manager for future central stations. He is the author of three papers on bioelectronic systems and is named as a co-inventor on two patents. Larry is a native of Laramie, Wyoming and attended the University of Wyoming where he earned the BSEE and MSEE degrees with a biomedical electronics option in 1965 and 1966, respectively. He is married, has two children, and lives in Burlington, Mas-

sachusetts. His outside interests include playing bridge, woodworking, gardening, and improving his home.

SPECIFICATIONS

HP Model 78501A and 78502A Patient Information Centers (for more details, see data sheet 5953-1383 for 78501A, 5953-1384 for 78502A)

78510A DISPLAY

SIZE: 169 mm H x 240 mm W - 5%
WAVEFORM DISPLAY (from 78511A)

	TIME	TRACE LENGTH
78501A	4 s ± 5%	100 mm ± 5%
78502A	6 s ± 5%	150 mm ± 5%

ECG SENSITIVITY RANGE (non-telemetry): Full scale equals ± 2.5 V ± 10% (1.28 V/cm ± 10%)

ECG SIZE FOR TELEMETRY UNITS: 1, 4, 1.2, 1, 2, 4, which corresponds to 0.625, 1, 2.5, 5, 10 mV full scale (× 250, × 500, × 1000, × 2000, × 4000 the voltage on the patient)

PACE PULSE VARIATION (0.2-5 ms width): Not greater than 2:1

ECG HIGH FREQUENCY CHARACTERISTICS: Deflection not greater than 3 dB down from midband value over frequency range from dc to 50 Hz. (Low frequency filter is determined at bedside)

HEART RATE

RANGE: 0.15 to 3 V corresponds to 15 to 300 bpm; indication is 0 below 15 bpm. 78501A and 78502A use heart rate from hardwired bedside below 80 bpm, when available

ACCURACY: ± 3 beats ± 2% of HR for constant rate input

ALARMS

UPPER LIMITS: 20-240 bpm in steps of 5 bpm

LOWER LIMITS: 15-235 bpm in steps of 5 bpm

DELAY: 5 s or less for rate alarms after rate violates limit

78571A and 78572A RECORDERS (see data sheet 5953-1387 for complete information)

PAPER SPEED: 25 mm/s and 50 mm/s ± 2% (50-60 Hz)

FREQUENCY RESPONSE: Not greater than 3 dB down at 100 Hz (per AHA specification for real-time recording)

LINEARITY: ± 1% for input waveforms within chart margins and response range specified

1 SECOND MARKER: Printed as dots on top margin of chart paper. Spacing is 1:0:01 s

BATTERY BACKUP TIME: 1 min (for system time)

WAVEFORM CHANNEL 1: ECG always, except (for 78502A only) when used to record P1 and P2

WAVEFORM CHANNEL 2 (78572A only)

78501A: One of up to three pressure waveforms (choice done at system setup and same for all beds)

78502A: Pressure, respiration or CO₂ waveform, depending on recording soft key

Alarm recordings—ECG and pressure (if available and configured)

PAPER: HP Chemical Thermal is recommended

78571A: 50 mm ruled (63 mm total width), roll or Z-fold

78572A: Dual 50 mm ruled (120 mm total width), roll or Z-fold

9876A THERMAL PRINTER (see data sheet 5953-1025 (40) for complete information)

RESOLUTION: 560 dots across 185-mm wide plotting field; 3 dots/mm horizontally and vertically. Four 9876A dots are used to plot one data point from 78502A bit map.

PAPER: Perforated, Z-fold, black chemical-thermal image. Depending upon environmental heat, humidity, and light, will fade with time. For archival use, photocopy is necessary

PAPER SIZE

U.S.A.: 216 mm W x 279 mm L

METRIC: 213 mm W x 297 mm L

TRENDS (78502A only)

SHORT TRENDS SAMPLING: Vital-sign data is sampled every second (when bed alarms are on); one data point for every active parameter is accumulated by the trend processor every 12 seconds. At least two valid data points are averaged every minute and stored for use in short trends

LONG TRENDS SAMPLING: Once every six minutes (for 9-hour trend) or every 16 minutes (for 24-hour trend) the most recent averaged data for short trends will be averaged or minimum/maximum values determined to update long trends

RESOLUTION

VERTICAL: 83 points

HORIZONTAL (Time): 91 points

BATTERY BACKUP (Physiological data—trends and alarm limits): 24 hours

SCALES

HR	30-190 bpm
VPB	0-40 VPB/min
Respiration	0-40 R/min
CO ₂	0-80 mmHg or 2-10 kPa
Temperature	29-45 C
Pressure	30-190 mmHg or 0-32 kPa

OTHER WAVES DISPLAY (78502A only): Pressure waveforms are displayed with a graticule whose scale is derived from the bedside monitor range switch

ENVIRONMENTAL REQUIREMENTS:

TEMPERATURE RANGE: 0 to 55 C, operating; 40 to 70 C non-operating

POWER: 100, 110, 120, 220 or 240 volts - 5%, 10%, 44-66 Hz

PRICES IN U.S.A.:

78501A Patient Information Center: \$19,600

78502A Patient Information Center: \$17,300

T04 Trend Option: \$3000

78571A Single-Channel Recorder: \$3850

78572A Dual-Channel Recorder: \$4950

MANUFACTURING DIVISION: WALTHAM DIVISION

175 Wyman Street

Waltham, Massachusetts 02154 U.S.A.

panel of ICU/CCU nurses from local Massachusetts hospitals and several physician consultants represented the medical user's viewpoint. A multinational panel of sales engineers represented the interests of customers in the many

countries that HP serves. A panel of service engineers critiqued the installation, service, and reliability concepts of the patient information centers.

High-Speed Raster Technique Provides Flexible Display

by Robert L. Stettiner and George L. Adleman

THE DISPLAY FOR THE 78500 SYSTEM provides the operator with smooth high-uniform-brightness waveforms that are superimposed upon a field of alphanumeric or graphic data. The smoothness eases the operator's job of detecting abnormalities in ECG or other medical waveforms, while the uniform brightness provides the visibility of larger CRTs in a conveniently compact unit. The high flexibility of placement for waves and alphanumeric/graphic data allows clear labeling of waveforms from critically ill patients.

The 24-by-16.9-cm display is generated upon a 30-cm conventional CRT using raster deflection. Waveform smoothing is done by using a high raster frequency to overlap raster lines and beam-width shaping on individual lines. The low deflection power required for the raster technique enables the system display unit to operate reliably without a cooling fan and associated dust filter.

High-Speed Raster Design

The raster display is very similar to that of the ubiquitous computer terminal except that the display is scanned vertically, from bottom to top, instead of the more usual left to right. The scanning rate also differs; it is approximately 70 kHz instead of the usual standard (~15 kHz) or double (~30 kHz) television raster rates.

Examination of Fig. 1 can lead us through the problem of waveform display. Fig. 1a shows the type of display that would be generated by turning on the video to make a dot at the exact location where an imaginary waveform overlaid upon the screen would cross a vertical raster line. The separation of these dots proves objectionable to physicians and nurses who are trained to recognize medical conditions by studying smooth, direct-written electrocardiograms.

A continuous picture can be made, as shown in Fig. 1b, by turning on the video between two points representing two successive values of the wave coinciding with two successive raster lines. Many displays of this type have been manufactured for use in the medical field. The major drawback of the technique is that the relatively square corners on the displayed trace create a waveform appearance that is both artificial and different from the waveforms the operators are familiar with. This makes pattern recognition difficult. Even increasing the number of vertical lines on the screen to 2000—a raster rate of over 100 kHz—will not remove the artificiality of the waveform. Combining higher raster rates with synthetically-thickened lines does resolve the corner perception problem; however, overly high raster rates are required. The combination of moderate raster rate and beam width shaping produces an economically feasible solution.

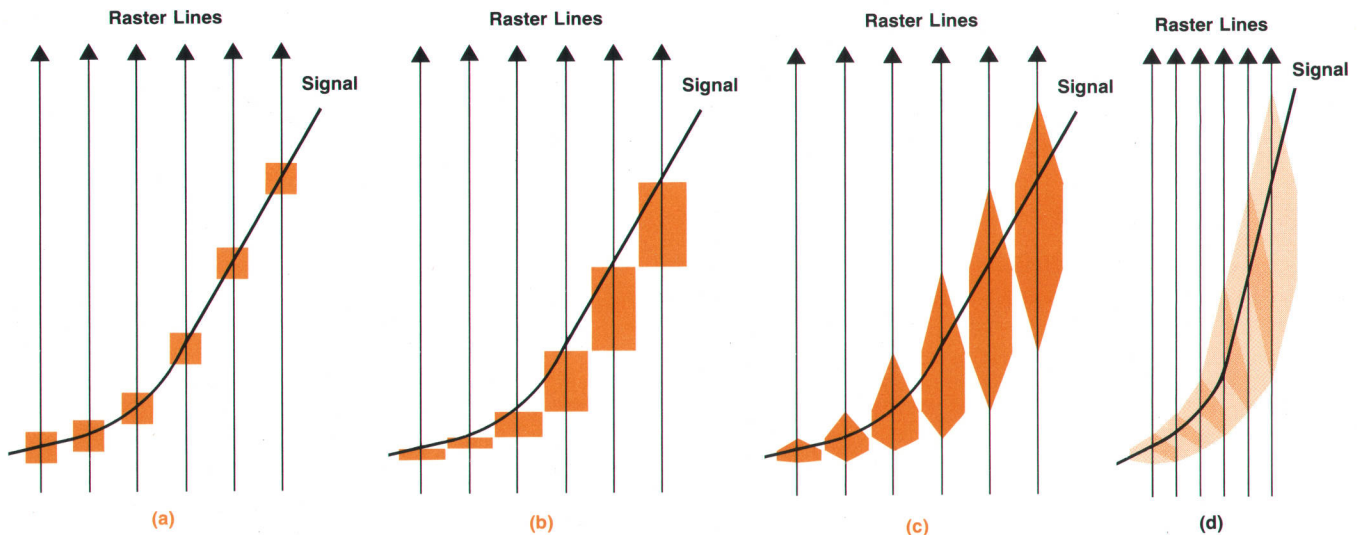


Fig. 1. Various approaches for depicting a signal waveform using a raster scan display (the horizontal spacing between raster lines is greatly exaggerated to clearly show the beam shape on individual lines). (a) Dots are displayed where each raster line coincides with the waveform. (b) Video is turned on between each two successive intersections of the raster scan lines with the waveform. (c) Same display as for (b) except that the video is modulated smoothly on and off instead of abruptly switched on and off. (d) Same as (c) except that the horizontal spacing is reduced to show how the traces actually overlap.

Fig. 1c and Fig. 1d show a continuous-appearing waveform generated by modulating the beam width from zero to full width between adjacent corners, instead of cutting the beam off sharply. This modulation of beam width occurs naturally in standard low-resolution electromagnetic deflection CRTs with intensity modulation. Continuous modulation—as shown in Fig. 1c—is not a necessity because the modulation can be done in steps. Four steps are a very good approximation to the smooth function, while three or even two steps may be acceptable, depending on subjective factors. Fig. 2 shows a normal and an expanded view of a waveform segment from the actual 78510A display, which uses four-step modulation.

Implementation of the Display

Fig. 3 is a block diagram of the display. The 305-mm (12-in) CRT has 1200 vertical raster lines (5 lines per mm). The beam width is 0.3 mm at cutoff and up to three times wider at higher brightness. Raster rate is 69.4 kHz for a period of 14.4 μ s, with 2.7 μ s used for fast (vertical) retrace and the remaining 11.7 μ s used for display on each raster line. Alphanumeric and general graphics are stored in a memory map of the screen and added to the separately generated waveform video and sync pulses.

The waveform video could be generated by any technique that results in video similar to that shown in Fig. 1c. Our system generates the video by comparing the waveform sample values to the current CRT beam position. Analog comparison techniques were chosen over digital because of cost benefits.

The waveform sample value is loaded into the memory by the microprocessor during the fast retrace time that immediately follows a load request. The appropriate sample is read from the memory to a digital-to-analog converter during the 11.7 μ s display time. During the next fast retrace, the output of the digital-to-analog converter is transferred to

Basic Deflection Techniques

CRT beams can be deflected by either electrostatic or magnetic fields. Electrostatic-deflection tubes are generally found in such high-frequency applications as oscilloscopes, primarily because the field energy required to deflect the beam is some three orders of magnitude lower than that required for an electromagnetic-deflection tube. Electromagnetic-deflection tubes are generally less expensive because of their common use by the television industry. These tubes also have a higher beam efficiency, (i.e., a higher percentage of electrons leaving the cathode arrive at the phosphor screen). This efficiency is often 95% for an electromagnetic versus 5-10% for a wide-angle electrostatic CRT.

CRT displays are grouped into two main types—directed beam and raster. The directed-beam display, commonly found in oscilloscopes, linearly deflects the beam to follow the signal pattern in the vertical direction while often using a constant deflection rate in the horizontal direction to represent time. For this type of display, electrostatic tubes are used to minimize deflection power. When high brightness and tight focus are mandatory, some high-speed, directed-beam displays use electromagnetic deflection with the disadvantage of high power consumption (typically more than a kilowatt).

Raster displays are more familiar because of their use in television. The beam follows a fixed path, covering the face of the tube in a series of parallel lines. The picture is generated by modulating the beam intensity and hence the picture brightness at the appropriate places on the CRT face. One advantage of this technique is that when the end of a raster line is reached, the energy stored in the electromagnetic deflection field can be reversed in polarity by low-loss resonant energy transfer. The power required is typically 10 watts with only 1-10% of that amount used in the associated solid-state drivers.

Cost, brightness, and power are factors that make the electromagnetic-deflection raster a highly desirable display technique.



Fig. 2. (a) Photo of a normal waveform as shown on the 78510A Display. (b) Expanded view of waveform in (a) showing video modulation required to display waveform.

the appropriate position in an analog shift register. A direct implementation of an analog shift register with data being transferred down a line of sample-and-hold stages would suffer from the accumulation of charge transfer errors. Our approach uses a set of four sample-and-hold stages that are sequentially loaded with the latest sample value. Thus, each stage holds a value until it is updated. The samples are put into proper time relationship by analog multiplexers which unscramble the sample-and-hold operations. Thus, the four samples—A,B,C, and D—are switched to comparison circuit block locations corresponding to their time sequence. Resistor dividers are used to generate the inter-

mediate values between A and B, and between C and D. These values represent points 25%, 50%, and 75% along the line between the sample pairs. Fig. 4 shows the step-modulated video on a single raster line as a function of the four intersections (A,B,C, and D) between signal and raster lines. During the next visible display period, the eight values—six fractional and B and C—are compared to a ramp. The comparator outputs are combined by fast logic to form four digital video signals. The outputs closest to the outer samples A and D are combined to form a 1/4-brightness signal. The outputs halfway between the outer sample pairs are combined to form the 1/2-brightness signals. The remaining two comparator pairs are combined to form 3/4- and full-brightness signals. The four separate gray levels are combined with their counterparts from the other three waveform processors in the system. The resulting four digital signals are fed to a digital-to-analog converter which converts them to an analog video signal. This digital-to-analog converter is weighted in a nonbinary fashion to provide gamma correction (the exponential relationship between grid-to-cathode voltage and brightness) for the CRT.

Graphic Memory

The graphic bit map is relatively simple to implement. The face of the CRT is mapped into a 260-row-by-300-column memory block. The memory is loaded in 10-bit words by a microprocessor during a fast retrace period. The memory is read out in bit-serial form during the 11.7- μ s display time for each line. Slight deviations from standard television display practice are made to accommodate the high speed and high number of raster lines. Because of the high raster speed, each dot read out of the memory must be read onto the screen in 45 ns. The conventional (500 ns) dynamic RAMs are multiplexed 20:1 at their output to allow them to operate at comfortable speeds. The multiplexing is performed with a Schottky parallel-in/serial-out

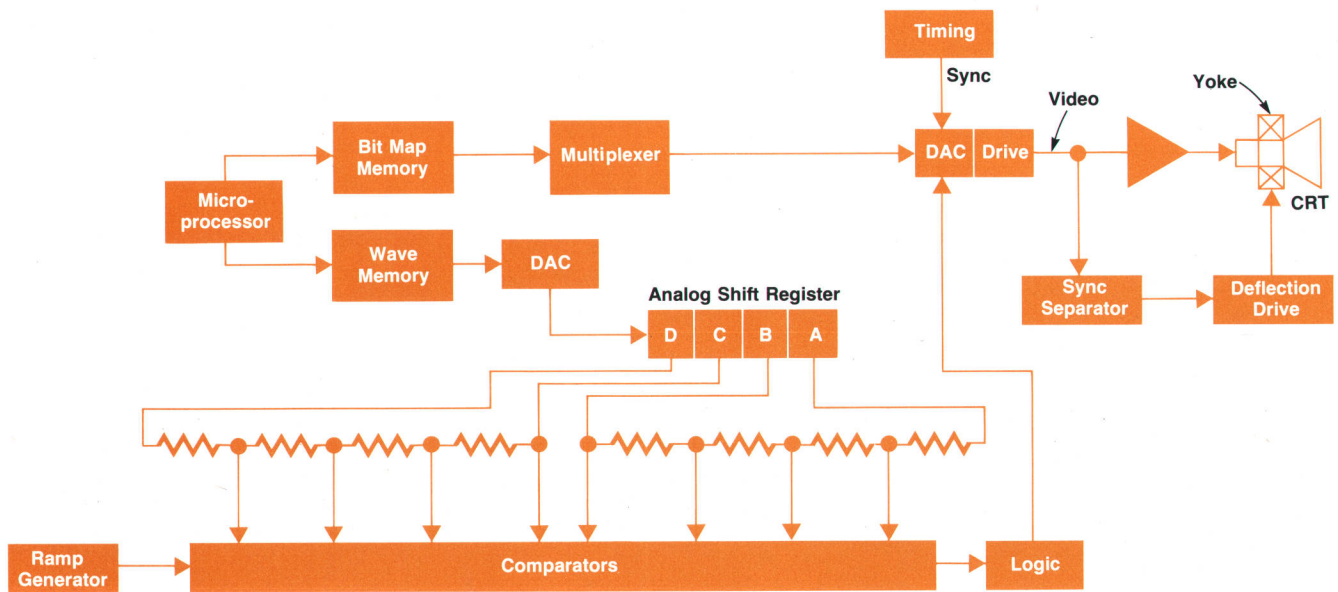


Fig. 3. Block diagram for 78510A display.





Fig. 4. The smoothly modulated approach shown in Fig. 1c can be approximated by the varying the video beam intensity in four distinct steps. The step-modulated video for one raster line is shown.

George L. Adleman



George Adleman is a native of Boston, Massachusetts and was awarded the BSEE degree in 1970 by the University of Massachusetts. He joined HP in 1975, initially worked on the 78500 display, and now is a project leader for a bedside product display. Before coming to HP he did peripherals engineering and designed automatic transistor test equipment. George is a member of the IEEE and is a co-inventor on a patent for smoothing a raster display. He is married, has two daughters, and lives in Brockton, Massachusetts. During his leisure time George enjoys cooking,

bicycling, woodworking, and silversmithing.

Robert L. Stettiner



Bob Stettiner came to HP in 1966 and has been the project manager for a number of instruments in the 78300 series, some earlier medical products, and most recently for the 78510A Display. His prior work experience includes development work for radar systems and automatic test equipment. Bob has authored two papers on displays, is a co-inventor on a patent for phase interferometers and a patent for a medical display, and is a member of the Society for Information Display. Bob has a BEE degree awarded in 1956 by the City College of the City University of

New York. A native of New York City, he is married, has two sons in college, and lives in Lexington, Massachusetts. Bob has a number of outside interests that include swimming, cross-country skiing, cooking, bridge, music, science fiction, and history.

Sample Selection Filtering and the Bandwidth Problem

The 78510A display has five raster lines per millimetre of CRT horizontal distance. The represented sweep rate for the displayed waveforms is 25 mm/s which is derived from accepted medical use of displays and paper chart recorders for electrocardiograms. Thus, the display presents 125 samples per represented second of waveform. The Nyquist Theorem shows that only signals with a bandwidth less than 62.5 Hz can be presented with this sampling rate. Other practical considerations reduce the signal bandwidth to less than 40 Hz, at which point some significant information is lost to the viewer. In clinical cardiac monitoring, the most significant data that would be lost is the pulses from cardiac pacemakers, which are less than 2 ms long. The entire 78500 system has a sampling rate of 500 Hz and can reproduce a 100-Hz-bandwidth signal on a strip chart recorder. Have we created a display that is a severe limit on the system?

Conventional analog or digital filtering techniques can be used

to approach the 62.5-Hz figure with an expense commensurate with the closeness of approach. Sample selection filtering allows the 78510A to show a waveform with an apparent bandwidth of better than 100 Hz although the word "bandwidth" is not directly applicable to a system that has aliasing and phase distortion.

A sample to represent a period of 8 ms is selected from four 2-ms samples taken during that time. The selection criterion is to maximize the peak-to-peak amplitude of the displayed wave. Each of the four samples is compared to the average of the preceding four samples and the sample selected is the one that has the greatest absolute difference from the average. No artificial signals are generated and the resulting waveform, whether seen on the CRT or written out on a strip chart recorder, is visually almost indistinguishable from a wave made out of the actual 2-ms samples.

cascode design to drive both the grid and cathode of the CRT. By cutting the swing of each transistor output stage to half of that required for single-ended drive, lower-power faster transistors can be used.

Acknowledgments

Walter McGrath, Bill Reed, and Dick Stanley provided a

physical package for the display electronics. Rick Beebe developed the sample-selection filtering technique. George Elenbaas, Dave Erickson, Bill Kole, and Mark Wendell developed the circuitry for the high-speed raster display unit. John Scampini and Tonny Wong developed circuitry for generating the waveform and alphanumeric video and interfacing the video generators to the display control microprocessor.

Multi-Processor Architecture and Communications for Patient Monitoring

by James M. Rueter

THE MAIN INTELLIGENCE of the 78501A and 78502A Patient Information Centers is contained within the 78511A Equipment Cabinet. This cabinet can be remotely located from the display and recorder to provide a compact equipment arrangement for the medical staff. Remote placement of the cabinet can also allow easy access when failures occur. The equipment cabinet for the 78501A system contains four processors. Each performs some of the monitoring tasks required and each communicates with the other processors to accomplish the overall monitoring function. Three different communication schemes are used for communications between the processors within the equipment and to other instruments. Each scheme is optimized for the particular needs of the interface it serves.

Each processor is an HP MC5 microprocessor made with HP's CMOS silicon-on-sapphire (SOS) process. There were

several reasons for choosing the MC5 in this application:

1. An average instruction time of around 1.2 μ s gives excellent processing capability.
2. A rich I/O instruction set improves interfacing efficiencies.
3. The majority of the data to be handled is 10 bits wide, hence a 16-bit microprocessor can manipulate data by use of single instructions and single registers.
4. Excellent software development tools are available.

Multi-Processor Architecture

Even with the processing power of an MC5, the total task could not be performed by a single processor. Fig. 1 shows the block diagram of a 78501A Central Station with connections to eight bedside monitors. Each processor uses the HP-IB* to communicate with other processors within the

*Hewlett-Packard's implementation of IEEE Standard 488 (1978).

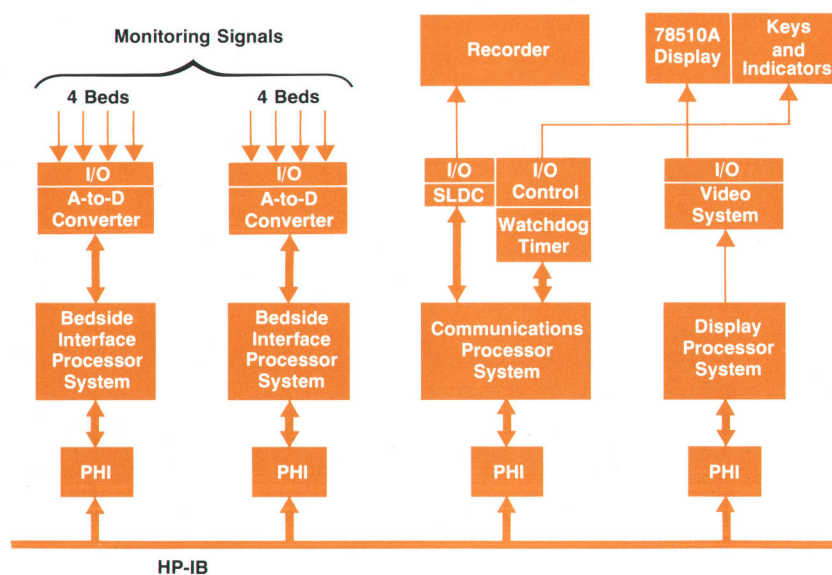


Fig. 1. Block diagram of a 78501A Patient Information Center connected to eight bedside monitors. Communication among the processors in the system uses the HP-IB.

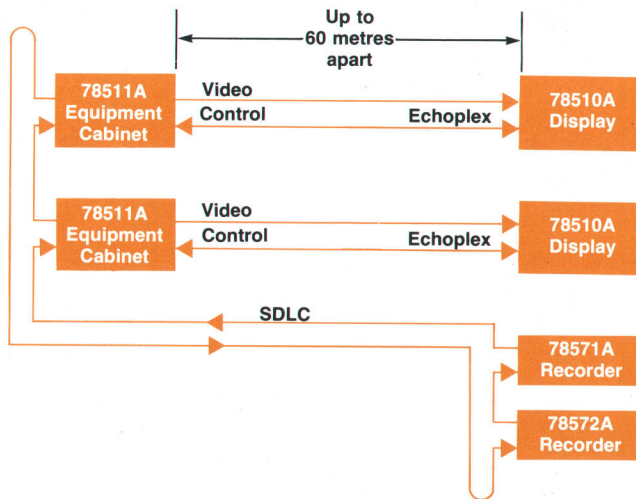


Fig. 2. Three types of communication links are used in a 78502A system: HP-IB for communication inside the 78511A equipment cabinet, a private serial line (echoplex) between each equipment cabinet and its display, and a loop using SDLC to interconnect equipment cabinets and recorders.

equipment cabinet.

Each of the bedside interface processors digitizes the physiological signals from four bedside monitors. These processors check the bedside for status and alarm conditions. This information is passed on to the other processors. The communications processor controls the normal operation of the HP-IB and manages the communications with the strip chart recorders and the keyboard of the 78510A Display. The display processor receives display information from the HP-IB and formats the CRT screen properly.

If any of the processor systems fails, the remaining processors continue to operate and communicate and make some of the monitoring functions available. The loss of one bedside interface processor will cause the monitoring of four of the eight beds to stop. The loss of the communications processor will cause the display processor to take over the task of controlling the HP-IB. The communications processor has a watchdog timer which is continually reset by software during normal operation. If the watchdog timer is not reset, signaling a fault, the communications processor is automatically prevented from exercising control of the HP-IB. If the communications processor malfunctions, display of the ECG waveform and other parameters continues, but the controls on the display and the strip chart recorder will not function. If the display processor fails, no meaningful display is provided, but alarm and manual recordings (using the keys on the display) remain operational.

The backplane in the 78511A Equipment Cabinet has a common connection to the bottom of all circuit boards in the card cage, except for the four specialized display boards. The connection of each processor's bus to its associated circuit board(s) is made by ribbon cable connections on the top of the boards. This scheme provides a flexible arrangement so that new processor systems can be developed and added in the future without changing the design of the original hardware. Only the firmware to allow

the new processor to communicate and interact intelligently will have to be changed. (Currently only about 2/3 of the space is used in an average equipment cabinet.)

Equipment Cabinet Communication

The processors within the 78511A must have a fast, efficient means of communication to maximize the processing time available for monitoring tasks. The HP-IB was picked as the means to ensure efficient communication for the following reasons:

1. HP-IB is a standard communications interface whose design and protocol have been thoroughly thought out and tested.
2. Many existing LSI circuits are available for the operation of the interface. These are available at modest cost and relieve the processor from performing many of the low-level details involved with controlling communication. IC drivers designed specially for the HP-IB are also readily available.
3. The test equipment needed during development to monitor and debug communications is widely available.
4. In the future, new processors of different manufacture or architecture can easily be added because almost all microprocessors have a compatible HP-IB interface IC.
5. Most of all, HP-IB allows a high data rate (up to one megabyte per second) among multiple processors.

Each MC5 processor uses an HP-designed HP-IB interface IC called PHI. PHI is designed to be compatible with the MC5 and uses the same CMOS silicon-on-sapphire process. PHI allows the MC5 to operate in response to commands sent to it by an HP-IB controller or to manage the HP-IB as controller. The complex logic of PHI relieves the processor of many of the details of HP-IB operation. The speed of communications is improved by the use of the eight-word inbound and outbound FIFO (first-in, first-out) buffers built into PHI. They buffer data so that communication can continue for a period of time without processor intervention.

Communication between Instruments

There are two types of two-way communication links between instruments in a standard 78501A or 78502A system. They are shown in Fig. 2. The first is a private serial line between each equipment cabinet and its corresponding display. This link tells the display which front-panel indicators to light and which sounds to generate. Information from any front-panel key pushes is returned on a reverse path. The second communication link uses synchronous data link control (SDLC, developed by IBM) in a loop configuration. This loop allows up to four equipment cabinets to control up to three recorders.

The display is designed to be as simple as possible to provide a package that takes up minimal space in the central station area. This also lets more of the circuitry be located in the remote equipment cabinet where it can be more easily serviced.

To minimize complexity in the display and still allow two-way communications between the display and the equipment cabinet, a simple asynchronous serial interface called echoplex is used. Using a few wires, echoplex can transmit information over a long distance. Two-way communication takes place without need for a microprocessor

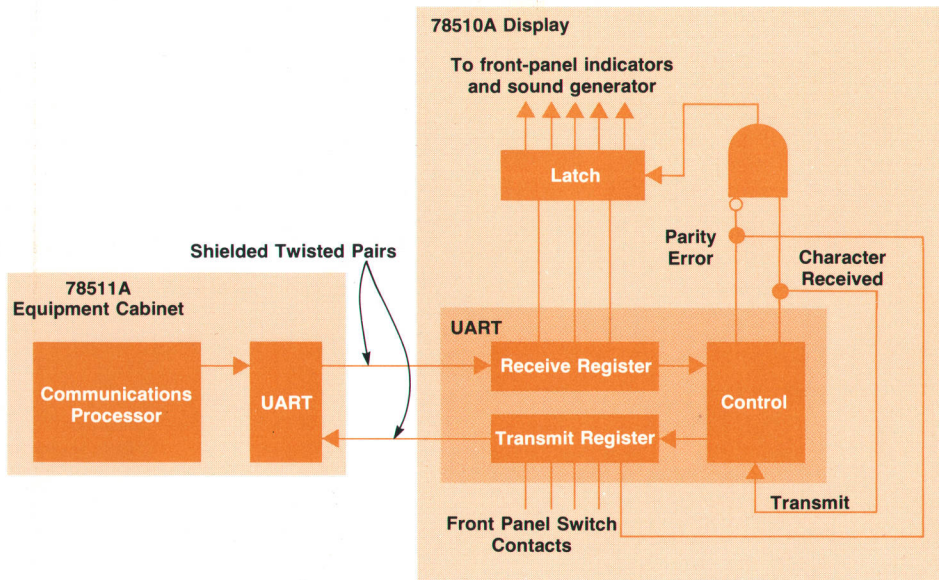


Fig. 3. Block diagram of the link between the 78511A Equipment Cabinet and the 78510A Display. This asynchronous serial interface is called echoplex.

or complicated logic in the display. A block diagram of this link is shown in Fig. 3. Using a universal asynchronous receiver/transmitter (UART), the communications processor sends a word (seven data bits + one parity bit) serially to the display. The UART in the display receives that single word. The UART has an output indicating that a character has been received and is available in its receive register. If the received parity is correct, then the received data is latched to drive front-panel indicators and the sound generator. The pulse that indicates that a character was received is connected to the transmit control terminal to output a single character back to the equipment cabinet. This character is formed by the transmit register inputs, which record any key pressed on the display's front panel. The received parity is also sent back as one of the data bits to indicate that correct data is reaching the display.

With this arrangement, every time the communications processor sends a character, it receives a character (hence the term "echoplex"). The display is interrogated on a regular schedule to determine if a key has been pressed or not. Care must be taken to send a valid character to the display each time. Even if no new front-panel indication is desired on the display, the exchange will cause data to be relatched.

A self-test feature tests the display circuits independently. A switch causes the UART's inputs to be connected to its outputs and a key push now causes a character to be sent and received to exercise the front-panel indicators.

This simple communication link allows numerous control signals to be exchanged between two instruments when one of them has little intelligence. They can be separated by a considerable distance, since a single differential transceiver at each end of the link can provide good noise immunity.

Multiple Monitors

In many installations there are multiple sets of monitoring equipment at the central station. In these cases, it is desirable to share a single strip recorder or several strip recorders among several display/equipment cabinet pairs. When multiple recorders are present, recordings can be

made concurrently.

The communication method used in such a system must have several characteristics:

1. Many instruments must be able to exchange information.
2. Long distances between instruments must be acceptable, such as equipment cabinets located up to 60 metres from the recorders at the central station.
3. Relatively high data rates must be possible to allow several recordings to be made at the same time.

SDLC provides these features and also the advantages of reliable error detection and readily available integrated circuits to control the basic operation of SDLC communication. SDLC is a bit-oriented serial communications protocol. Each message has an eight-bit address field, an eight-bit control field and a variable-length data field. Each message contains a sixteen-bit cyclic redundancy check (CRC) that forms the end of each message and is used to detect errors that occur during transmission. Each message is bracketed by a special flag pattern, which is used to synchronize messages. This pattern is 01111110. Fig. 4 depicts the format of a message.

To differentiate between flags and a pattern that would be identical to a flag within a message, bit stuffing is performed on the message when it is transmitted. After the flag is sent, a zero is stuffed (inserted) after any five consecutive ones in a row. This eliminates the possibility of any false flags being transmitted. When the end of the message is reached, the trailing flag is sent without bit stuffing. The receiver recognizes the unique flag pattern preceding messages and "unstuffs" all zeros after five consecutive ones until it recognizes the trailing flag.



Fig. 4. Format for SDLC message. The 01111110 flag pattern is used to synchronize messages.

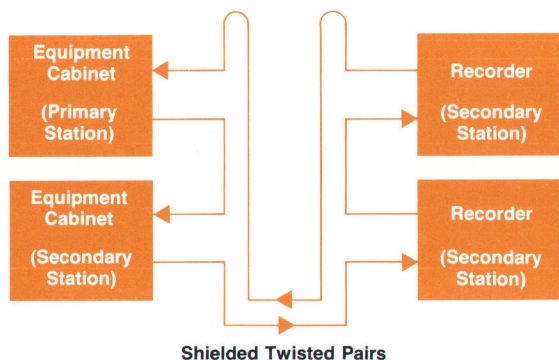



Fig. 5. The SDLC loop interconnects the equipment cabinets and recorders in a 78500-Series system. The loop is formed with a twisted pair of wires in a shielded cable.

Fig. 5 shows SDLC in a loop configuration. One equipment cabinet is the primary station which controls the link and breaks the loop. All other instruments on the link are secondary stations that respond to commands from the primary station. When a transmitted message flows around the loop each secondary station inspects it to see if the address contained in the message matches its own address. If the address matches, the secondary station receives and processes the message. The secondary station can respond when the loop is available by sending down the loop a message with its own address. The primary station receives all messages and recognizes the address as that of the secondary station.

The SDLC signals are sent differentially over a twisted pair of wires in a shielded cable to reduce problems of external interference and common-mode voltages. Each instrument is connected in sequence to the instrument preceding it and the loop is completed by another pair of wires in the same cable.

The use of synchronous communication requires that all stations have a synchronized clock to receive the serial data properly. To eliminate the need for a separate signal, the clock information is encoded with the data using NRZI encoding. NRZI code requires that a zero be represented by a transition in the transmitted signal and a one be represented by no transition. The transition caused by a logical zero is used to synchronize a counter in the receiver. This

James M. Rueter



Jim Rueter earned a BEE degree in 1966 from the Georgia Institute of Technology and joined HP that same year. He has designed various types of bedside monitors and acted as project manager for the arrhythmia monitoring system and the processor hardware of the 78501A and 78502A. Jim is a member of the Association for Computing Machinery and a native of Savannah, Georgia. He plays tennis and squash and enjoys bicycling when he is not pursuing his studies for a master's degree in computer engineering as an honors co-op student at Stanford University. Jim is married, has three children, and lives in Boxborough, Massachusetts.

counter is running at an integer multiple (15 in this case) of the expected bit time so that the data can be sampled near the midpoint of its bit cell (when the count is 8). Synchronization at the beginning of each message is assured since the preceding flag contains two zeros. Frequent synchronization also occurs during a message since the transmitted signal cannot have more than five ones in a row due to bit stuffing. As a consequence of using NRZI, the two wires for the differential signal can be reversed with no difference in operation.

Acknowledgments

A great many people contributed to making the 78501A and 78502A Patient Information Centers a reality. My thanks to all these people and especially to members of the design team who worked with me. Bill Shannen and Don Brown designed the packaging and interconnection schemes. Kerry Newcom was the designer of the processor board and numerous design aids. The power supply resulted from the effort of Doug Johnston. Dan Finlayson was responsible for the SDLC and printer interfaces. The bedside interface was designed by Tremont Miao and Mike Kramer.

Self-Test and Serviceability for Dependable Central Patient Monitoring

by Jeffrey M. Goldberg

SELF-TEST IS THE ABILITY of an instrument to determine without user intervention whether its circuits are functioning correctly and reliably. Serviceability is the implementation of fault isolation procedures and the addition of necessary hardware and software to aid in tracking down and repairing equipment failures. To achieve reliability and fast servicing for the 78500 Patient Information Centers (PICs), a four-point service plan was developed. It consists of:

1. Automatic detection of equipment failures using built-in tests to isolate the faulty instrument quickly
2. Isolating the defective board, enabling board replacement within one hour, using minimal additional tools
3. Optional component-level repairs typically requiring a service analyzer board and some additional troubleshooting hardware such as a volt-ohm meter, oscilloscope and digital signature analyzer
4. Verification that the fault has been corrected in the instrument using internal self-test.

The goal was to produce an instrument that would retain much of its functionality, even in the presence of an instrument failure. Following this design philosophy, the PIC continuously monitors itself and is immediately forced into a backup (partial operation) mode whenever the instrument's ability to continue functioning in its current mode is in question. To do this, each processor immediately stops what it is doing and performs a self-test using the kernel expansion testing technique. This involves starting with a known functional resource (in our case, the microprocessor) and using it to test surrounding resources which, if operational, may be used to test even more resources. This allows the processor to test associated hardware without hanging itself on the defective component. It then uses the information gathered from its internal diagnostics to reconfigure the system by severing all diseased sections. Processors that go out of control even with the precautions taken are caught by hardware. This hardware, called a tickle circuit, must be periodically serviced by a processor, otherwise it will trigger a general system reset sequence.

Although many of the features of self-test and serviceability are transparent to the operator unless a malfunction occurs in the system, they continuously reassure the user that the instrument is functioning properly in the medical environment. Self-test enables timely detection and correction of faults when they occur, rather than allowing them to go undetected and potentially disrupt patient care. When a malfunction is discovered, the PIC immediately localizes the problem and channels the ECG and other vital monitoring functions to the remaining operational hardware. It then informs the operator of the loss of any monitoring capability, and the extent of the damage. A permanent record of the fault condition, including date and time, is

preserved for service personnel. These features, combined with built-in modularity to assist board replacement, speed and simplify the otherwise very complex repair task.

Four levels of increasing self-test and diagnostic capability are available in the 78500 PIC series (see Fig. 1). Each implements only the required degree of fault detection and error reporting necessary to perform its portion of the total task. Before proceeding into a detailed description of the

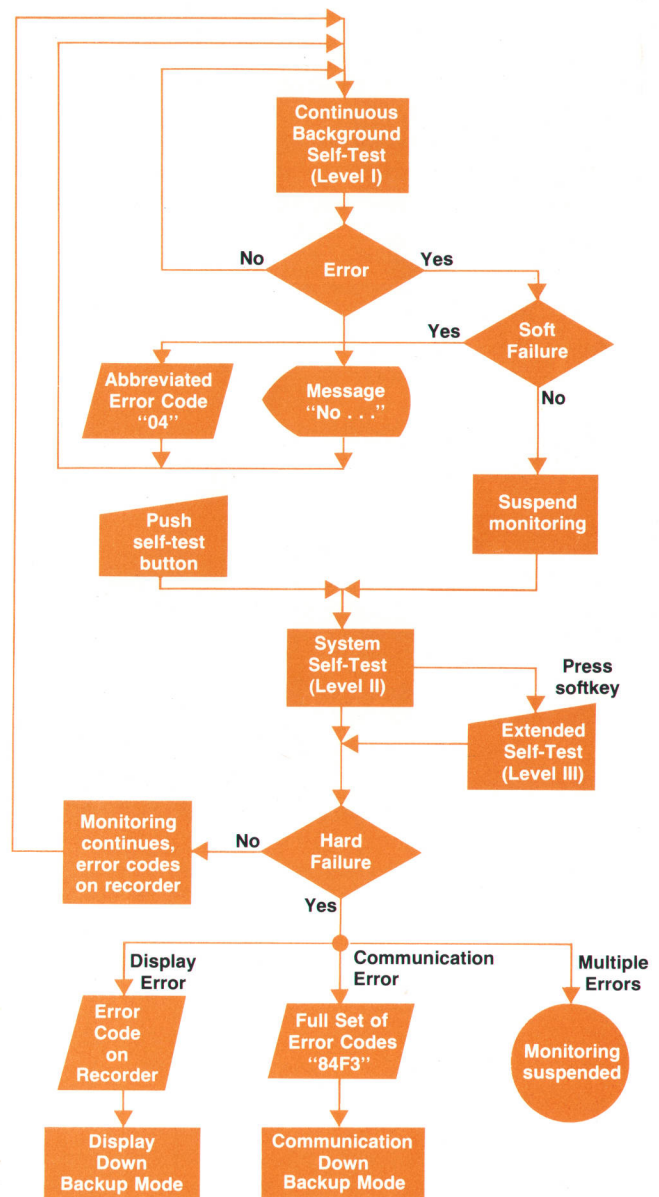


Fig. 1. Flow chart of the 78500-Series self-test sequences.

operation and implementation of the four levels it may be of some benefit to review the hardware architecture of the PICs from the serviceability aspect. The 78500 PIC is a distributed microprocessor system composed of three basic building blocks. The communications subsystem is responsible for coordinating information transfer among the internal subsystems and controlling all external communication links to devices such as recorders, softkeys, and hard-copy devices. The display subsystem is responsible for all graphics and real-time wave processing including monitoring patient parameters. The bedside interface subsystem is responsible for sampling data from a group of four patients and making the data, as well as derived parameters, available to all the other subsystems. In a 78501A PIC, there are two bedside interfaces. In a 78502A PIC, there may also be a trend subsystem (an optional block that provides trending and additional pressure processing). Each subsystem not only carries its share of the workload, but also has the capability to take over some of the responsibilities of one of its neighbors. Thus, if the communications host processor, which normally controls the internal communications bus, fails during operation, the system will reconfigure itself with the display subsystem taking over this function. Similarly, if the display subsystem should fail so that the patient's waveforms are not being displayed, the communications subsystem can redirect the waveforms manually selected by use of the softkeys to the recorder.

The individual subsystems are all tied together by an internal communications interface bus (HP-IB*). The hardware employed to implement this protocol also allows a failing subsystem to store an error code on the bus. This code is available upon request from the system controller. Each processor board has a light-emitting diode (LED) to indicate a failure in its associated subsystem. The activation of any one of these process malfunction lights (PML) illuminates the externally-visible instrument malfunction indicator. This signal is detected by the auto-restart circuit, which in turn originates a global reset to all boards in the system, starting the self-test cycle. Power startup and pressing the self-test button also create the reset signal.

Level I: Continuous Self-Test

An automatic fault monitor constantly surveys the system's operation seeking evidence of incorrect performance. In case of a problem, and depending on the severity, a warning message or an error code number may be generated, or a complete shutdown may result. In any case, an instrument malfunction lamp and an audible reminder are activated whenever the instrument hardware or monitoring operations are in question.

Error codes are four-digit hexadecimal numbers that tell service personnel the likely cause of the problem. The meaning of all error codes can be found in the service manual. These codes list the subsystem, printed circuit board and sometimes even the component that has failed. These codes fall into two categories, hard and soft. Errors detected in functions critical to patient care are classified as hard; the others are classified as soft. Hard errors curtail further use of the subsystem in which they are detected, while soft errors allow the use of the hardware to continue

while informing the operator of the condition. A consistent faulty conversion of patient data into digital form is one example of a hard failure. A stuck bit illuminated on the display screen is an example of a soft failure.

A number of messages are intended for the user, either to identify the problem or suggest corrective action. For instance, "Clean air filter" or "Check paper/door on recorder #1" are self explanatory. Backup warning messages and system error codes inform the operator that the system has lost some of its monitoring capability and is functioning in one of its backup modes. "Backup—No recorder or control" is a prime example.

The primary function of continuous (Level I) self-tests is to detect failure modes while the instrument is in operation. A mandatory requirement is that this process be automatic and not allowed to interfere with normal monitoring operations in any way.* Thus, the fault checker is implemented as a background task. In the PIC, background tasks acquire a slot of run time every 32 ms from an interrupt-driven operating system. This occurs during a period of time when the processor would otherwise be idle. This provides for repetitive checks of processor ROM, RAM, and most resources critical to patient care. In addition, all control and data links are monitored for a loss in communication. Thus if the arrhythmia computer goes down, or if someone trips and pulls out the control cable to the PIC, or if the recorder runs out of paper, the system immediately brings these problems to the operator's attention. Of course, for this process to be completely transparent to the patient-monitoring processes, hardware devices cannot be reinitialized or taken off-line during testing. RAM locations can only be checked nondestructively, one location at time, with interrupts temporarily disabled during modification. More stringent requirements, such as software interlocks, were required for testing memory storing long-term data (such as trends, where the data must be kept intact during power interruptions).

Level II: Power-On and User-Initiated Self-Test

Diagnostics built into the system perform comprehensive system checks of the hardware, software drivers and all communications links. These checks are made at any of the following three times—when power is first supplied to the instrument, a few seconds after the detection of a faulty condition found by the Level I self-test, and by request of the user whenever the SELF-TEST button is pressed. This self-test module is responsible for all reporting of error codes and reconfiguring the system into its various backup modes whenever necessary. External indications of Level II self-test include error messages on the display and a recorder test strip annotated with error/status codes, giving the operator a hard copy of all reported faults for future reference. This redundancy makes retrieval of error codes always possible. If one mechanism fails, the other can be used to determine the code.

User-initiated self-test consists of seven seconds of software simulation tests and dedicated, noninterruptible, self-terminating hardware tests. All hardware and software initialization is performed at this time.

*Hewlett-Packard's implementation of IEEE Standard 488 (1978).

*Implementation of this design was greatly enhanced by the use of the multi-priority operating system that was incorporated into the monitoring software.

Table I

Backup modes are responses of the 78501A and 78502A to internal failures that allow as much system operation as possible under the circumstances. Examples of these backup modes and the operation possible with them are:

Loss of display processor	Alarm and manual recordings possible from central station and bedside.
Loss of communication processor	Standard display provides ECG and pressure. No cardio control or initiation of recording possible but bedside alarms can be used.
Loss of arrhythmia system	Internal cardioteach activated.
Loss of recorder	Display and alarms are active.
Loss of power	Battery backup maintains control settings (clock keeps running in recorder).

The following functions are performed simultaneously in all host processors during user-initiated self-test. First, a reset is issued to most hardware devices including the processor. This hardware signal is the initiator of this self-test cycle. The processor performs its dedicated self-test using the kernel-expansion technique described previously. The processor initializes all remaining devices and places the instrument into a known state. After the ROM, stack, RAM and supporting devices are thoroughly checked, the processor exercises the interrupt mechanism to insure that it is functioning. Any failures that the self-test detects are stored using a corresponding four-digit hexadecimal error code. If the failure is nonrecoverable, the affected subsystem is severed from the remaining system. By this time, each subsystem has identified itself and has tuned in to listen on the HP-IB for the first communication attempt among subsystems. Assuming the communication subsystem is functional, it will send out an initial test message individually to each of the eight possible subsystems. It will then poll their status and ability to communicate. If it finds that it is unable to communicate with all of the subsystems, it will assume that it is at fault and die gracefully so that the display

subsystem can take over as communications manager. As long as portions of the PIC remain functional, the instrument will continue to operate by dynamically configuring itself with the hardware that is presently operational and available (see Table I).

Level III: Extended Self-Test

Internal synthetic data processing, sequential activation of all alarm and malfunction paths and generation of a display test pattern are done to check out many of the aspects of video generation. All this is performed in addition to Level II tests. The operator uses Level III tests to compare performance with previous operational performance, or against standards described in the service manual. Extended self-test suspends the monitoring function by its very nature, and therefore is only entered at the user's request. Pressing any one of the softkeys after Level II self-test has started, but before monitoring has resumed, enables this mode. Pressing the **STANDARD** key resumes normal monitoring.

For the convenience of the user and service personnel, PIC status and a complete log of all errors are recorded (see Fig. 2) and displayed in this level. When a failure occurs, this feature facilitates fast isolation to the faulty subsystem and boards without special service tools.

Level III diagnostic software is called extended self-test because it adds additional user-interpretive testing to Level II diagnostics. Commands are sent sequentially to each of the processors to activate the red-alarm and instrument-malfunction indicators. Information pinpointing the fault can be derived from the activation sequence. For instance, if the red-alarm indicator fails to activate, but only during the second period of this sequence, this is good evidence that the problem is local to the second subsystem. Thus, the problem exists in the display subsystem, not in the global red-alarm drive circuitry. After each subsystem has had the opportunity to alternately activate the red-alarm and instrument-malfunction indicators, the communications processor will activate the yellow and green alarms in sequence, and then turn on all alarms and control indicators for one to two seconds. The display subsystem will then be commanded to put up a waveform test pattern (Fig. 3, used to determine any peculiarities in the display generation), a




Fig. 2. Stair-step and ramp waveforms are recorded to verify recorder response and linearity. A log of any errors found during extended self-test is also recorded. Errors are indicated by a change in the last three zeros of the status code number(s). An example of a recorder test with no errors is shown above.




Fig. 3. During the extended self-test the above test pattern is displayed on the 78510A Display to aid in determining if there are any irregularities in the display generation. The status codes for the system are displayed on the right. Any errors are indicated by a change in the last three zeros of the appropriate status code number.

bit-map alignment pattern, and a complete list of status codes for the entire instrument. Finally, a prompting message of what to do next is displayed.

Level IV: Component-Level Fault Detection

The fourth and final level of diagnostics deals with component-level repairs made on the site. The majority of the hardware and software used to troubleshoot the instrument in this mode is supplied in the service support package and runs with minimal support from the instrument itself. This allows Hewlett-Packard to implement component level repairs without increasing the cost of the instrument to those installations that plan to support their instrument primarily through Hewlett-Packard's board replacement program.

Using the 14451A Service Circuit-Analyzer-Board (Fig. 4) from the service support package, additional facilities may be exercised in a given subsystem to aid in finding the failure. These facilities include additional and more extensive diagnostics than are provided in the built-in self-tests, circuit excitation programs for probing logic paths with an oscilloscope and digital-signature analyzer, ability to take control of and test all data and control buses, observation of the execution of all testing modules through the output of test identification numbers to the 14451A display, and ability to view all error codes produced by a subsystem. Level IV provides for component-level repair when board exchange is impractical. It is especially useful for hospitals that maintain their own equipment, where minimum replacement board inventory is a must. However, it can also be used to speed up the fault isolation and correction verification on board replacement repairs.

The service circuit-analyzer-board software can be subdivided into automated diagnostic programs and excitation programs. The former checks out a specified board, precisely instructing the operator when operator interaction is required. Diagnostic tests provide a list of error codes de-

scribing every fault condition found on the board during testing. Excitation programs, on the other hand, stimulate the hardware on the board for signal probing and digital signature analysis. Service personnel can then track down the problem using documented service procedures and signal comparisons.

Besides the extensive service circuit-analyzer-board software programs provided to test each of the subsystems, the 14451A provides direct access to all microprocessor buses to detect opens and shorts. The analyzer can directly access the built-in firmware. It can monitor the execution of all test software and latch the test identification number of each test run. If the system should ever hang up without indicating an error, information as to the last test module executed can easily be retrieved for examination. In addition, there are other hooks built into the PIC that let the 14451A and other troubleshooting tools monitor the operation of the instrument and seize control over its operation whenever necessary for servicing. In effect, this dual effort transforms the simple service circuit-analyzer-board into a powerful bus and logic analyzer.

Acknowledgments

The serviceability and fail-soft strategy implemented in Hewlett-Packard's 78500 Patient Information Center was based on the collective insights of many individuals. The author wishes to thank all the individuals involved for generously giving their support, time and suggestions during the project. Special acknowledgment should go to Scott Bonnet, Bob McHenry, Jeff Barnes, and the marketing staff for their continued and unshakable belief in promoting a




Fig. 4. The 14451A Service Circuit-Analyzer-Board allows service personnel to exercise additional facilities in the system to aid in identifying a failure.

truly serviceable product, Larry Nielsen and Ken Fuchs for their implementation of initialization/self-test's shell mainline, Doug Johnston and Joe Kobrenski for their initial work with the PHI chip, Jack Bailey for his mystical abilities in inventing a universal memory test module, and Mike Kramer, Jack Pirkle, and Al Scarpone, for their extensive work on the service circuit-analyzer-board—providing service personnel with extensive factory-like testing capabilities out in the field. We all owe a debt of gratitude to Paul Johnson for his insights as chief production engineer and from the many late nights spent on personal time to improve the product. And finally, special thanks to Bob Stettiner and Jim Rueter, whose leadership and extensive hardware background helped me through my most difficult moments during product development.



Jeffrey M. Goldberg

Jeff Goldberg is a native of Boston, Massachusetts and attended Tufts University in Medford, Massachusetts where he was awarded the BS degree in electrical engineering in 1977. Jeff came to HP that year and has worked on the design of the power supply for the 78510A Display and the self-test and serviceability software for the 78500 Series PICs. He has taught FORTRAN programming and is a co-author of a paper on computer simulation of ventricular tachyarrhythmias. Jeff has held various offices in the local scuba club, is single, and lives in Burlington, Massachusetts. During leisure hours he enjoys skiing, scuba diving, tennis, and racquetball. Jeff is also an electronics enthusiast and recently built a home computer for his father's birthday present.

Firmware for a Patient Monitoring Station

by Kim L. Hanna

FIRMWARE DEVELOPMENT for the 78500 Series Patient Information Centers (PIC) was a major portion of the total development effort. The 78501A firmware consists of three sets of code and the 78502A has four sets. Although there is general commonality between the two products, there are some major differences. Code for the 78501A processors consists of 9K words of monitoring and 3K words of initialization/self-test code for each processor. Monitoring code for the display and interface processors of the 78502A expanded to 13K words with an additional 9K words of monitoring code for the trend processor. Each of the processors uses 3K words of read/write memory (RAM). In addition, the interface and trend processors have RAM with battery support for non-volatile storage of certain control information and trend plots.

The code was developed using an assembler resident in an HP 3000 system. A number of structured constructs—IF-THEN-ELSE-END IF, LOOP-EXIT WHEN-END LOOP, and CASE—are provided by the assembler. These constructs, in conjunction with extensive equivalencing (hardware-independent coding), increased the code reliability and decreased development time significantly. A relocating linking loader, also resident in the HP 3000, was used to create absolute code from a number of relocatable assembly files. The absolute code was then either loaded into RAM using a debugger stored in a read-only memory (ROM) or programmed into programmable ROMs.

Several aspects of the firmware are covered in this article. The first is a description of the interface processor software for the cardiometer (heart-rate calculator and alarm

generator) which has increased patient-alarm reliability and reduced the occurrence of false alarms. The HP-IB* communication system linking the processors within one PIC is described next. A description of communication among several centers for recorder and hard-copy sharing follows. Finally, the operating system, which is resident in each of the processors, is discussed.

Cardiometer and VFIB Detection

One of the major contributions of the 78500 Series is improved analysis of the electrocardiogram (ECG) waveform. ECG monitors usually contain a cardiometer (tach) for heart-rate (HR) calculation from the patient's ECG. This rate is compared with user-selected low and high limits to generate audible and visual alarms. Previous tachs were quite sensitive to the muscle noise (artifact) generated by patient motion. This noise often caused a falsely high rate that exceeded the high-rate limit. Additionally, previous tachs did not always generate an alarm during ventricular fibrillation (VFIB, a life threatening condition in which heart contractions are uncoordinated and no blood is pumped). The 78500 tach significantly reduces false high-rate alarm frequency and improves VFIB detection reliability.

The key to tach false-alarm reduction is a QRS-matched filter. This filter is a fifteen-sample finite-impulse-response filter whose impulse response shape approximates the shape of a normal heartbeat waveform (QRS complex, see Fig. 1). The filter, therefore, has maximum absolute output when similarly-shaped waveforms are input. Other wave-

*Hewlett-Packard's implementation of IEEE Standard 488 (1978).




Fig. 1. (a) Typical QRS waveform for one heartbeat. A matched filter processes fifteen 8-ms samples from the ECG waveform. If the samples approximate the values shown in (b) the filter has maximum absolute output.

components (e.g., the T wave) produce reduced outputs.

Fig. 2 shows the tach block diagram. The patient's ECG is sampled every 2 ms in the interface processor. Fast-transition, high-amplitude components are attenuated by a slew-rate-initiated clamping algorithm to reduce the amplitude of pacemaker artifacts and the probability of counting these artifacts as beats. Two adjacent 2-ms samples are averaged in this process; the result is a train of 4-ms samples. Next, a four-pole, 30-Hz, infinite-impulse-response Butterworth filter removes unnecessary high-

frequency components of the signal, producing 8-ms samples in the process. A 1.25-Hz high-pass filter removes any dc offset from the signal. Finally, the clamped and filtered ECG waveform is passed through the QRS-matched filter.

The beat detector identifies QRS complexes in the processed ECG waveform by comparing the maximum absolute waveform value that has occurred since the last heartbeat to a threshold value. If this absolute value exceeds a threshold value, a heartbeat is considered to have occurred. The threshold, illustrated in Fig. 3, decreases from the average QRS wave amplitude to half that value at a rate determined by the average QRS amplitude and average heartbeat interval. An inhibitory period (200 ms minimum), during which no heartbeat will be detected, follows each beat. This reduces the susceptibility to T waveform counting. The inhibitory period is an inverse function of the high-rate limit, with lower high-rate limits giving longer inhibitory periods.

When a beat is detected, the average beat interval and average QRS height are updated. The beat interval, averaged over several beats, is used to calculate the heart rate for alarm limit comparison, heart-rate display, trending, and recorder annotation.

The second major addition to current tach practice is more reliable detection of ventricular fibrillation. VFIB is characterized by a sinusoidal ECG waveform resulting from uncoordinated heart-muscle activity. The VFIB detector (Fig. 4) uses the sinusoidal characteristics of the waveform to detect the presence of the condition and generate audible and visual alarms.

The average period of the ECG waveform for one second is first determined as follows:

$$T = 2\pi \sum_{1s} \frac{|V_j|}{|V_j - V_{j-1}|}$$

where T is the number of sample points in an average period and V_j is the value of the jth sample. This period is then used to adjust a notch filter, the notch being located at the average frequency corresponding to the average period determined above. This filtering is accomplished by adding

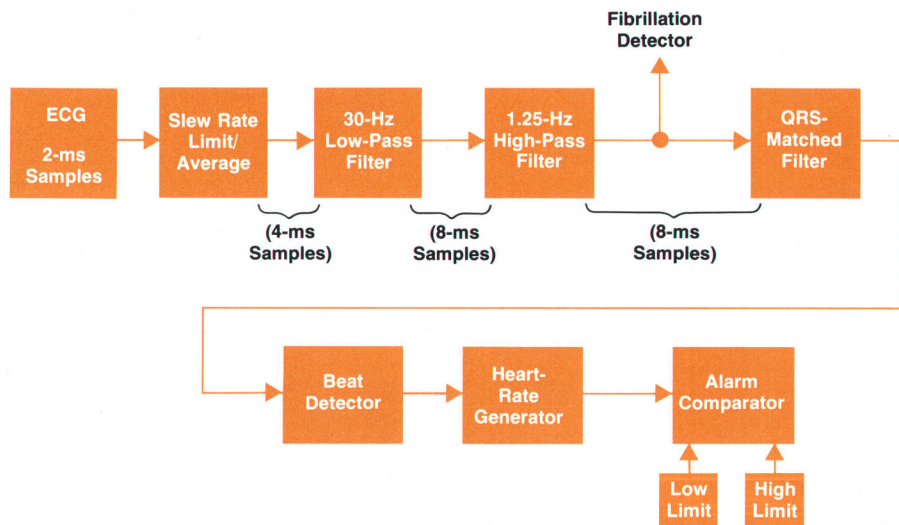


Fig. 2. Block diagram of the cardi tachometer.

two samples separated in time by the average half period:

$$(\text{Filter output})_j = V_j + V_{j-T/2}$$

The unfiltered signal and filter output are summed over a one-second period to calculate the "filter leakage fraction" as follows:

$$\text{Filter leakage fraction} = \frac{\sum_{1s} |V_j + V_{j-T/2}|}{\sum_{1s} |V_j| + |V_{j-T/2}|}$$

The leakage fraction will be zero for an ideal sinusoid input. Because the VFIB waveform is not an ideal sinusoid, a higher leakage fraction is used as the threshold for detection of this condition. Additionally, the average frequency must be in the range of physiologically occurring VFIB (2-9 Hz). If the leakage fraction and frequency criteria are met for three consecutive seconds, the VFIB alarm is sounded.

This ventricular-fibrillation detection algorithm and the improved heart-beat detector with decreased susceptibility to muscle-motion artifacts together form a very effective basic ECG monitor.

Data Communication

Communication among the processors in one equipment cabinet is accomplished using HP-IB and HP's custom PHI interface. HP-IB is a desirable communication method because a high data-transfer rate is possible, several simultaneous listeners to one message can be designated, and a very good peripheral chip (PHI) is available. The HP-IB controller function is normally assumed by the communications processor, but will be taken over by the display processor in the event of communication-processor failure. The controller routes messages based on a routing table (Table I) that is initialized at power-on.

The data transfer cycle, outlined in Fig. 5, is initiated each 32 ms when the controller sends a sync tap message to all processors. This message triggers all time-dependent

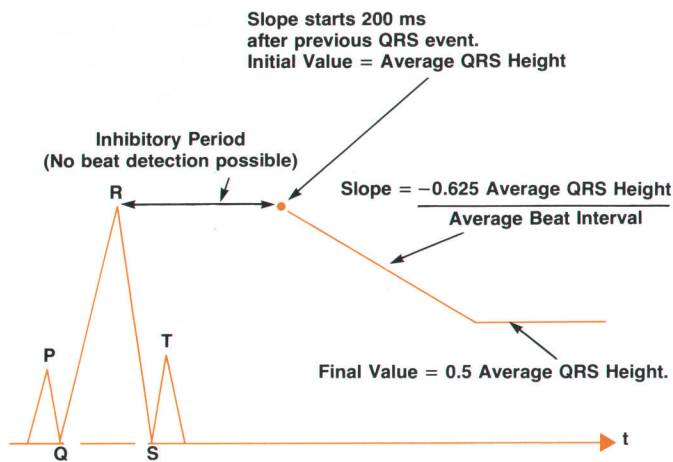


Fig. 3. Graph of the cardi tachometer threshold for detecting heart beats. The inhibitory period imposed after a beat is detected reduces the possibility of counting the T portion of the waveform as a separate beat.

Table I
Example Routing Table

R denotes receiver of message, L indicates that the communication processor is a listener when recording waves.

Message	Communication Processor	Display Processor	Interface Processor	Trend Processor
ECG	L	R		
HR Parameter		R		R
Record Request	R			
Interface Commands				R

Note: All receiver entries are eliminated for an absent or malfunctioning processor. For example, if trend is absent, system initialization will remove entries from its column.

activities in each processor. The sync taps act as a time reference so that each processor in the system operates on a common time base. Next, the controller sets up each of the processors in turn as a talker to output data to other processors. For each message, the processor is set up as a talker with the controller as listener for a two-byte transfer that contains the data ID. The data ID is examined and listeners are set up as determined by the controller's data-routing table. A data transfer is then enabled which is terminated by the talker sending the EOI (end or identify) message. The sequence of transferring the ID, setting up listeners, and transferring data continues until that processor has no more data to output. The remaining processors are handled in the same way until all have output data.

Initialization of the communication system consists of determining which processors are present and functioning correctly and initializing the routing table with permanent listeners for each of the data messages. For example, ECG data is output by the interface processor and received by the display processor for formatting on the screen. The data-routing table is modified dynamically during normal operation whenever users of data are added or deleted. For instance, the communications processor is set up as a listener of ECG data while recording the ECG waveform.

The PHI chip has a number of capabilities that make operation of this scheme easier and less time critical. First, it uses eight-word, first-in-first-out input and output buffers (FIFOs) that give a grace period to listeners and talkers before the processor must wait due to HP-IB handshake hold-up. A processor that is set up to listen must recognize this fact, and then start data reception. This process starts when the controller issues a listen command. This creates an interrupt in the receiving processor, but the sending processor is not held up until both its outbound FIFO and the receiver's inbound FIFO are full.

Another valuable PHI feature is the handshake abort. The PHI generates an interrupt whenever a read operation is attempted from an empty inbound FIFO or a write operation is attempted into a full outbound FIFO. This interrupt is used to move the program counter back one location before the next instruction is executed, which effectively repeats the instruction. Thus, a very tight loop for inputting or outputting data can be used without needing to check for FIFO status before each operation, resulting in a programmed data-transfer rate of 200K-bytes/s. This rate is twice as fast as is possible when the status is checked be-

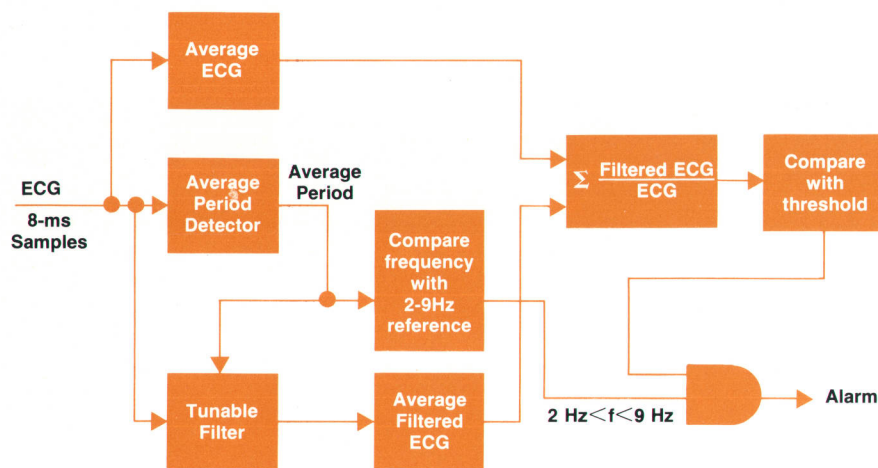


Fig. 4. Block diagram of the VFIB detector.

fore each operation.

Recorder and Hard Copier Sharing

One of the 78500 Series contributions is a greater ability to share recorders among a number of patients. A single recorder can be shared by up to 32 patients (8 maximum previously) resulting in much lower cost. Also, single and dual-channel recording capabilities can be mixed and shared, minimizing the effects of recorder failure.

To accomplish recorder sharing, all PICs (up to four) and the shared recorders are connected in a serial communication loop. Data is transferred using a bit-oriented data protocol which was chosen because:

1. A small inexpensive cable can be used
2. The protocol allows for destination addressing of messages
3. The speed is sufficient for transfer of 6250 bytes/s
4. Arbitrary data is allowed (no special characters are used as delimiters).
5. Transmission distances of 60 metres are required.

The protocol codes delimiters (flags) as a series of six ones, and insures that normal data never has six ones by stuffing a zero whenever five ones are detected in the data stream. The stuffing and destuffing, as well as cyclic-redundancy-check (CRC) error generation and checking, are handled by an interface chip and thus are transparent to the processor. A typical data message consists of an opening flag, address (of the destination unit), data, CRC value, and closing flag.

When powered on, one of the PICs becomes master of the serial loop; the others are slaves. The master has the responsibility for active device detection, allocation of devices to various centers, time/date acquisition (from a recorder) and distribution (to centers), loop integrity determination, and 32-ms clock distribution. Active device detection is accomplished by directing a message to the address of every potential instrument each second and waiting for a reply, which must arrive within 200 μ s for the device to be detected.

Every 32 ms a sequence of events occurs on the serial loop. The sequence starts with the master center sending a sync tap message to all devices using the "all units" address. Next, the master sends any recorder allocation messages (e.g., granting or releasing a recorder). Depending on the sync-tap number (0-31), a census tap, time tap, time

distribution, self-test, or status tap is performed with any appropriate replies. If the master has recorders assigned to itself it outputs waveform, annotation, and motor control messages to its recorders. When finished the master passes control to the first slave center on the loop. This slave sends any pending recorder allocation messages, then deals with the recorders assigned to it. Then it passes control to the next slave which performs similar operations. The last slave in the loop passes control back to the master.

Reception of a message addressed to a unit is detected by the unit's interface chip and a processor interrupt is generated. Because of the delay in responding to an interrupt in the PIC it is not possible to guarantee that the processor will be able to read the first data word before a second word has arrived and destroyed the first (word arrival period is 24 μ s). So, to avoid message loss, any message to a PIC is preceded by a wake-up message whose purpose is to generate an interrupt in the receiving processor so that it will be able to read the real message without losing any words. The wake-up message alerts the receiving processor about 100 μ s before it must start reading the data.

Operating System

Each of the processors has a number of competing tasks to perform at any given time. Such tasks as HP-IB data transfer, waveform sampling, and serial-loop data handling must be given very high priority, and are performed as interrupt routines. Other tasks, such as waveform display updating, HP-IB data preparation, and waveform analysis must be completed within each 32-ms period. However, there are other tasks that can be performed on a time-available basis after these tasks are completed. Some of these tasks are parameter sampling, display updating, recorder allocation, and trend-plot formatting. If several of these lower-priority tasks exist at a given time, a mechanism must exist both for determination of what is to be done first and for communication between the various activities. Instead of providing a unique means of processor allocation in each processor, the same operating system is used in each processor.

To accommodate the various tasks that need to be performed simultaneously, the operating system allows a number of processes to be active at any given time. Each process has an associated priority, and the operating system insures that the highest-priority ready process is running. A




Fig. 5. The HP-IB data communication sequence (a) is initiated every 32 ms. The data transfer details (b) apply to any noncontroller (e.g. display, interface processor) on the HP-IB.

lower-priority process is suspended when a higher-priority process is initiated. Process initiation is accomplished using a subroutine provided by the operating system. INITIATE (START_ADDRESS, PRIORITY, PARAMETER) is used to initiate a process with a given priority. Any currently-running process can initiate another process, and the initiated process then runs to completion on a time-available basis independent of the initiating process.

Two mechanisms are provided for process communication and coordination, critical regions (see page 28) and semaphores. Critical regions provide a mechanism that insures that only one process at a time is "in the region." Critical regions are used extensively to insure exclusivity without counting. Semaphores, on the other hand, count the number of signal operations (similar to an exit for a critical region) and allow immediate resumption of processes that wait on the semaphore if the number of unused signal operations is greater than zero. Semaphores are signalled when a message that must be processed by a lower-priority process arrives. In cases where more than one message may arrive before processing of the first is complete, the processing loop waits on the semaphore and will therefore run once for each message that arrives. Semaphores are implemented with the subroutines WAIT (semaphore #) and SIGNAL (semaphore #). These subroutines respectively decrement and increment the semaphore count. An active process, therefore, may be in one of several states: running, ready to run when processor time is available, waiting for a critical region to become empty, or waiting for a signal operation on a semaphore.

The operating system is very efficient and occupies only

about 350 words of ROM. It also requires five words of RAM for each process and two words of RAM for each semaphore or critical region. Approximately 100 μ s is required for the typical operating system subroutine call.

In addition to flexibility of process scheduling and freedom from the problems associated with other methods of scheduling, the operating system allows use of some very powerful debugging techniques. When no other processes are active, normal operation conducts as the lowest-priority task a RAM verification test and a ROM check-sum test. In normal operation, therefore, these tasks are performed on a time-available basis.

During development, however, there are two other tasks that can be substituted for these memory tests. The first executes a loop that creates a frequency-counter-compatible output pulse for each millisecond that the processor executes the code. By connecting a frequency counter, the percent of remaining unused time can be determined from the measured frequency: 10 Hz equals one percent. The second and most powerful use of the lowest-priority task is the running of a debugging program that interacts with a terminal to provide real-time memory examination and modification without interrupting normal real-time system operation. This examination and modification are done simultaneously with normal operation. This capability has provided some very powerful debugging techniques. Derived values and states of various stored parameters can be examined in real time and the response to uncommon or infrequent situations can be tested. This synergistic use of an operating system and a resident debugger running as the lowest-priority task has proved very

Critical Regions for Process Communication and Coordination

The concept of critical regions as a method for resource sharing and process communication was used in several areas in the 78500 Patient Information Center series.

A critical region has two possible states—empty and full. Two operations are possible with respect to a critical region—entering and exiting. Execution of the entry operation when a region is empty will make the region full, and the entering process will continue immediately. However, execution of the entry operation when the region is full will cause the entering routine to be suspended until the region is empty (caused by an exit operation from another process). Entry and exit operations must be protected to insure that one and only one entry is allowed before each exit.

The critical region concept was implemented in the operating system that is a part of each of the processors. Two subroutines are provided: `ENTER_REGION (REGION_NUMBER)` and `EXIT_REGION (REGION_NUMBER)`. All regions are initialized to empty at power-on. A process calling `ENTER_REGION` will have control returned to the instruction after the call either immediately (if the region was empty) or later when the region becomes empty.

This facility is used for several purposes. First, mutual exclusion is implemented using critical regions. In the display processor, the character-processor code converts ASCII-format codes into bit-map patterns for characters and graphics. This code is protected by a critical region because it is non-reentrant and is called by processes of various priorities. Before executing the character-processor code, a process must enter the character-processor critical region and be suspended if the character processor is being used by another process. Mutual exclusion is also used to protect data structures from multiple conflicting access.

A second major use of critical regions is for periodic timing. A routine that must be run at periodic intervals is coded as a loop whose execution is triggered by a timing routine that leaves the critical region at the desired rate. For example, updating the parameters on the screen at periodic intervals is handled in this way. If the processor becomes highly used, such as during display changes, this implementation has the feature that the screen update code will only be executed once, even though the timing routine may have left the critical region several times. This implementation also removes timing considerations from the periodic function code.

A critical region was also used to insure that recorder allocation among a number of PICs is handled properly. Each of the PICs has a number of processes that can request and stop recordings. A recording is started when a record request message from the HP-IB is placed in a circular receive buffer. Each PIC maintains a local list of pending record requests and recordings in progress. The master center maintains the highest-priority pending request from each unit.

The master also assigns recorders to individual centers (allocation process). To insure that the local request lists and the master's list agree whenever recorders are assigned, a global critical region (Fig. 1) was implemented to protect all local and master lists. So, before any list can be changed, the global region must be entered by the modifying process.

The master, in addition to allocation duties, manages this critical region. It accepts requests to enter the global region via the serial loop, and acknowledges these entry requests when the region is empty. The requesting center may then modify its local request list (a message is passed to the master, updating its request list) or relinquish a previously assigned recorder. Finally, an `EXIT_REGION` message is sent to the master allowing either another unit or the allocation process to enter the global region. In this way, the local and master request lists will be the same whenever recorders are allocated or released.

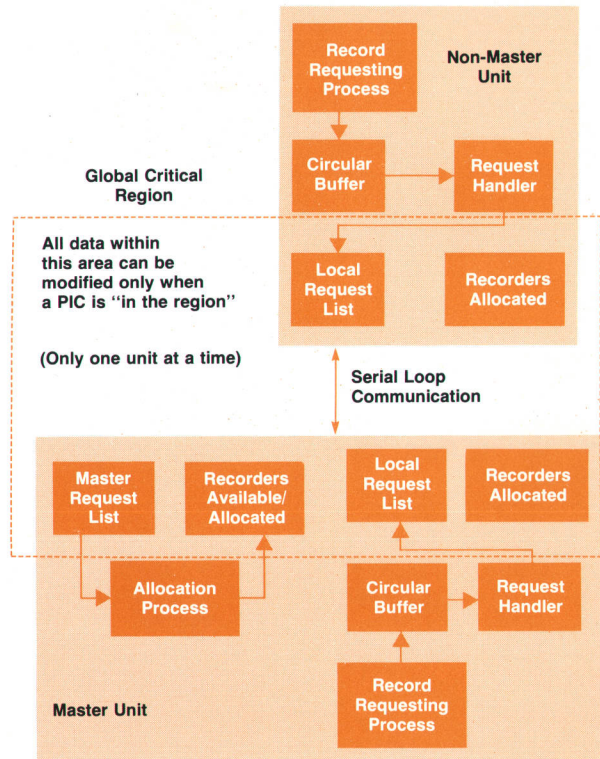


Fig. 1. Diagram defining the concept of global critical region.

powerful and time-saving.

Acknowledgments

The PIC software is the result of the efforts of many people. Ken Fuchs and Larry Nielsen participated in all phases from initial definition to final release, Richard Tang-Kong developed the interface between the PIC and bedsides, Frank Richichi, Bob Martin, and Joe Kobrenski developed the SDLIC and HP-IB software, Roger Fairbanks and Doug Hill did system QA, Les Bank and Dan Finlayson developed recorder and hard-copy firmware, Jeff Goldberg developed the initialization and self-test system, Evan Deardorff developed the operating system, Sandra Hand developed the tach, Pat Phillipps developed the graphics package, Don Lovegrove established a number of the development philosophies used throughout the project, and Paula Kelly and Maria Furnari were our software librarians.

Kim L. Hanna



Joining HP in 1972, Kim Hanna worked on respiratory diagnostic and monitoring equipment, HP-65 medical programs, and system design. He was the software project manager for the 78500 PIC and is working now on bedside monitors. A native of McAllen, Texas, he attended Rice University in Houston, Texas where he earned the BAEE (1971) and MEE (1972) degrees. Kim is named as a co-inventor on a patent for a patient monitor. He lives in Andover, Massachusetts, helps Explorer Scouts interested in electronics, plays racquetball and tennis, hikes and skis

cross-country, water skis, and enjoys scuba diving. Kim recently went on a white-water rafting trip. "The river won," he says.

An Interactive HP 3000/IBM Mainframe Link

More than a 3270 emulator, IML/3000 provides high-level intrinsics for interactive data interchange, and an inquiry and development facility that makes HP 3000 system terminals look like terminals attached to the IBM mainframe.

by **Connie J. Ishida**

UNTIL RECENTLY, AN HP 3000 USER could communicate with an IBM mainframe system only in batch mode, using MRJE or RJE. Since IBM mainframe systems originally developed as batch machines, with interactive access added later, batch access to IBM systems was an important first step. The next step, interactive access, is provided by a new HP 3000 software product called Interactive Mainframe Link (IML/3000). Using IML, the communications path to the mainframe is by emulation of an IBM 3270 cluster control unit with attached devices using binary synchronous communications (BSC) protocol.

The IBM 3270 information display system is a popular family of IBM terminals. A remote 3270 system consists of a cluster control unit with up to 32 attached CRT display terminals or printing devices. The 3270 operates on a leased line to the remote IBM host and up to 32 control units may be multidropped off the same leased line. Because of the popularity of the 3270 terminal, 3270 remote communications protocol has evolved as a standard for interactive access to IBM (or compatible) mainframe systems. A 3270 remotely attached to an IBM mainframe may be used in a mainframe session under Time Sharing Option (TSO), or in conjunction with a teleprocessing monitor to enter and retrieve data from the mainframe. IBM software packages CICS (Customer Information Control System) and IMS DB/DC (Information Management System Data Base/Data Communications) are teleprocessing monitors that interface between the 3270 terminal and the mainframe application program and data base.

IML/3000 gives the HP 3000 user access to mainframe sessions, data bases, and applications programs. IML has two modes of interactive access. The major mode of access is provided by a set of statements, or intrinsics, which are callable from an HP 3000 user program. The intrinsics provide high-level user access to the data being communicated to the IBM host. The intrinsics allow program-to-program communication between HP 3000 systems and IBM mainframes.

The other mode of IML access is provided by the inquiry and development facility (IDF) of IML/3000. This mode allows a user at an HP 264X block-mode terminal connected to the HP 3000 to log on to the IBM system as if the HP terminal were a 3270 terminal attached to the mainframe. IML/3000 takes advantage of the HP 3000's intelligent network processor (INP), using the INP to perform most of the 3270 control unit functions and handle the BSC protocol. INP is a microprocessor-based communications controller that relieves the HP 3000 of the burden of such functions as

protocol processing and 3270 data stream encoding and decoding.

Design Objectives

IML is designed to provide HP users with interactive access to their IBM mainframes, but it is designed to be more than just another 3270 emulator. Major design objectives were to provide customer ease of use, to minimize the impact on HP 3000 system performance by fully using the intelligent network processor (INP), to complement existing HP data communications products, and to make a contribution in the area of distributed data processing.

The intrinsics can be used within an HP application program in conjunction with the products V/3000 and IMAGE/3000 to create a distributed data base between the HP system and the remote host. If a request for information is entered from a V/3000 screen, the requested item will be searched for in the local IMAGE/3000 data base. If the item is not found locally, then the IML intrinsics can be used to form an inquiry to the IBM host to retrieve the information from the remote data base. IML makes the host's response available to the HP 3000 program, which may then update the local data base and relay the response to the requestor through V/3000. The HP customer decides how much or how little of the host data base to distribute. Some of the host application program functions might even be distributed to the HP application program. The HP application programmer's concerns focus on data base synchronization without the added burden of protocol or control unit data stream encoding and decoding.

IML can be used with an existing host application or an entirely new application can be developed on both the host and the HP system. In such a case IML provides a means for program-to-program communication. The 3270 screen image can be used as a buffer for passing messages. The screen panels can be simplified, since there is no longer a need to give operator instructions or to format the screen visually.

IML/3000 fulfills the need for interactive access to IBM mainframes but goes beyond this to provide a method for distributed data processing between IBM mainframes and HP 3000 systems. HP distributed systems software, DS/3000, can be used at the same time as IML/3000 to combine HP-to-HP and HP-to-IBM communications.

IML/3000 Organization

IML consists of several modules: a set of intrinsics, the IDF, a manager program, a monitor, and the INP software, which includes the control unit and protocol driver (see

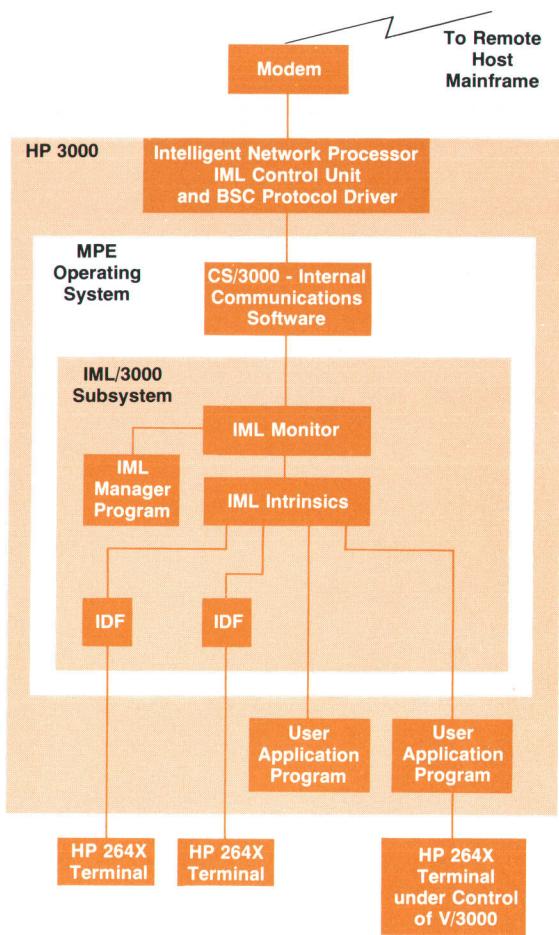


Fig. 1. IML/3000 consists of several modules and uses the HP 3000's intelligent network processor (INP) to minimize the impact on the HP 3000 system. The intrinsics are statements that are callable from HP 3000 user programs. The inquiry and development facility (IDF) allows a user at an HP terminal connected to the HP 3000 to log on to the host IBM system like a 3270 terminal.

Fig. 1). IML acts as a 3270 BSC control unit with up to 32 attached devices, where a device is either a user program or an IDF terminal. Additional control units may be supported by running IML on additional INPs. The application programs and IDF terminals read and write data into their screen images by calling the IML intrinsics. The monitor serves as a multiplexer for all transmit and receive requests from the intrinsics. The monitor communicates to the INP control unit and BSC protocol driver. The IML manager program allows an IML manager to define a configuration file for the control unit, start and stop IML and the communications subsystem trace facility, display the current status of IML, and acquire and release terminals for IDF use.

Intrinsics

The IML intrinsics provide access to IBM host systems for HP application programs written in COBOL, FORTRAN, BASIC and SPL. The IML manager uses security statements in the IML configuration file to define access to the intrinsics for users and programs. The intrinsics simplify user access to the host data. For example, the user can call

READFIELD, specifying a field number, without knowing exactly where that data is within the screen image. Ease of use for the HP application programmer was the goal, giving the program access to the data without a need to be concerned with 3270 or BSC protocol. The user may think in terms of field numbers, for which READFIELD and WRITEFIELD are provided, or in terms of locations within the screen, for which READSCREEN and STREAM3270 are provided. The TRAN3270 intrinsic is used to send the modified data within the screen to the host and the RECV3270 intrinsic is used by the program to accept the new host data.

The IML intrinsics can be used with standard I/O or in no-wait I/O mode. The HP 3000 Multiprogramming Executive intrinsics IOWAIT and IODONTWAIT have been adapted slightly for use as IML intrinsics.

The IML/3000 intrinsics are:

OPEN3270	Used to start communications with the host. A device is opened by allocating the internal screen buffer and the host is told that the device is now active.
WRITEFIELD	Writes the data specified by the calling program into the specified unprotected field.
STREAM3270	Starting from a given position this intrinsic moves across the screen entering data in fields it encounters. The stream of data may include special characters to indicate cursor movement and other functions within the screen panel.
TRAN3270	Indicates that the user program wishes to transmit the modified data in the screen to the host. This intrinsic interfaces with the IML monitor to accomplish the transmission.
RECV3270	The user program receives a screen that has been newly modified by the host. The internal screen image was already updated, so this intrinsic merely acknowledges the user program's acceptance of the new screen data.
READFIELD	Reads data from the specified field for the user program.
READSCREEN	Reads part or all of the internal screen image.
SCREENATTR	Returns information about the screen, such as cursor position.
FIELDATTR	Returns information about the specified field, such as the attributes of the field and the length of the data within the field.
ERR3270	Returns the error message associated with the specified intrinsic error number from the IML message catalog.
VERS3270	Returns the current version of the IML/3000 software running on the system.
RESET3270	Equivalent to pressing the reset key, which enables the keyboard. The 3270 keyboard is disabled after a transmit key is pressed.
CLOSE3270	Closes this device, informs the host that

the device is powered off and terminates IML use for the user program.

The following intrinsics are for use with no-wait I/O:

IOWAIT	Waits for completion of a previous no-wait I/O request.
IODONTWAIT	Checks for completion of a previous no-wait I/O request. If there is a completion, the user is informed, otherwise the intrinsic returns without waiting for any completion.
ABORT3270	Aborts an outstanding no-wait TRAN3270 or RECV3270 request.

Inquiry and Development Facility

The IML/3000 inquiry and development facility (IDF) was written as an application program, so that it interfaces to the 3270 data through the IML intrinsics in the same manner as any user program would. IDF uses the IML intrinsics to read the data from the internal screen image and composes the data stream with the necessary escape sequences to display the screen on an HP 264X block-mode terminal. When the user hits the ENTER key on the terminal, IDF reads the screen and then uses the IML intrinsics to update the internal screen image. IDF then calls TRAN3270 to send the data to the host. IDF makes use of the no-wait I/O intrinsics to be interrupted either by the user at the terminal or by the receipt of a response from the host.

The 264X softkeys are used to implement the 3270 program function, program attention, and clear keys. A 3270 has either 12 or 24 program function keys, which are implemented on a single 264X softkey. When this key is struck, IDF reads the screen and then displays a small menu to ask the user which program function key is desired. There are other differences between the HP and IBM terminals that preclude the use of IDF in a dedicated data entry mode. The purpose of IDF is to provide a tool that easily allows a standard HP block-mode terminal to be used to log on to a remote IBM system. IDF can be used for program development on the host or for casual inquiry to host programs or data bases.

IML/3000 Manager

The IML manager module is used by the IML manager, a person who is responsible for constructing the IML configuration file. This file specifies the IML environment and the users, programs, and HP terminals that can access the IML intrinsics and IDF. This file may also specify HP terminals that are automatically acquired for IDF when IML is started. The manager can start and stop the subsystem, turn communications subsystem (CS) trace on and off, and acquire other terminals for IDF use. The concept of having an IML manager does not require that the manager have access to the console. However, there is also a console type command, IMLCONTROL, which can start and stop IML and turn CS trace on and off. The manager can display information about devices and line(s) in use. The IML manager module and IMLCONTROL are designed to answer the varying needs of customers for control of the IML subsystem.

IML/3000 Monitor

The IML monitor is the master controller for the entire subsystem. When the IML subsystem is started, the IML monitor is created as a process. It remains active in the system until the subsystem is terminated. The IML manager program serves as a pathway for the IML manager to speak to the monitor.

The IML monitor multiplexes requests within the subsystem. It interfaces to the INP code through the communications subsystem (CS). The monitor also interfaces to the IML intrinsics TRAN3270 and RECV3270. The manager module passes requests (START, STOP, TRACE, etc.) to the monitor. When the subsystem is started, the monitor calls the CS intrinsic COPEN to send the IML download file to the INP. After the INP memory has been successfully downloaded and initialized, a positive response is returned through CS to the monitor, and IML is started. Since up to 32 devices can be in use at one time, the IML monitor takes care of queueing TRAN3270 requests for the INP.

Intelligent Network Processor Software

The 3270 display does not scroll; it functions only in page mode. The 3270 screen format was designed for data entry and allows the IBM application programmer to design screen panels and define fields anywhere within the panels. Fields within 3270 screen panels are delimited by special characters called attribute characters. IML maintains an internal character-by-character image of the screen for each device on the control unit. This screen image is passed to the INP software, which accesses the screen image to encode and decode the outbound and inbound data streams. The screen image is passed back from the INP to the HP 3000 where it is maintained in an extra data segment. The extra data segment also contains the field offsets to locate the attribute characters within the screen panel.

The INP download file for IML consists of three major processes. There is a physical driver/editor which is responsible for placing outbound characters onto the line and for accepting and analyzing inbound characters. A BSC protocol data stream is transferred between the physical driver/editor and the second process, the line control monitor. The line control monitor deals with BSC protocol and transfers 3270 data streams (less the BSC protocol) to the third process, the control unit. The control unit process interfaces with the IML monitor on the HP 3000 to obtain necessary screen images. The control unit constructs outbound 3270 data streams from the screen image and also decodes the host commands and orders within an inbound 3270 data stream. Thus the INP takes care of all the protocol and the screen image is returned to the 3000.

Acknowledgments

IML/3000 has involved a number of people over the last few years. The original design team included Spencer Frink and Chris Sauer. Ken Klingman and John Alderete served as project managers before me. Others involved in completing and releasing IML are Larry Goldman, Bob Carlson, Steve Bitondo, Ian Campbell, Bart Burstein, Bob Bagshaw, and Phil Sih. Jim Candlin, our section manager, helped bring this project from start to finish through all the ups and downs. I would like to thank all of these people for their contributions to IML. I really appreciate all of their efforts.



Connie Jean Ishida

Connie Ishida is project manager for IML/3000. She joined HP's General Systems Division in 1977 as a development engineer in data communications, having spent the previous three years designing operating system firmware for a business equipment manufacturer. Born in Denver, Colorado, Connie received her BS degree in computer science from Colorado State University in 1974. She's active in sports, including frisbee, racquetball, basketball and dance, and enjoys sewing, crafts, cooking, and reading science fiction. She and her husband, who also works in

HP's computer operations, live in Santa Clara, California.

Interactive Mainframe Link/3000

HP Model 32229A 3270 Emulator Software

FEATURES

- Programmatic access and update: Allows user-written programs in COBOL, COBOL II, BASIC, FORTRAN, or SPL to access or update, in an on-line fashion, files or data bases on the host mainframe, through host application programs.
- All IML/3000 program statements ("intrinsic") are high-level procedures for maximum programmer productivity.
- Provides for a smooth transition to distributed processing with the HP 3000 from centralized host processing; allows off-loading of teleprocessing function from the host mainframe to the HP 3000.
- IML/3000 intrinsics may be used in HP 3000 programs with the intrinsics of any other HP 3000 software product—IMAGE/3000, V/3000, KSAM/3000, DS/3000—to develop very powerful and flexible applications.
- Inquiry and Development Facility: allows direct access to CICS, IMS, TSO, or other teleprocessing or timesharing applications on the host from HP terminals, without requiring programming on the HP 3000.
- Simple keyboard commands allow the user to switch between a local HP 3000 session and remote access to the host mainframe through Inquiry and Development Facility.
- IML/3000 uses the Intelligent Network Processor (INP) hardware interface to the HP 3000 to offload some of the communications overhead from the HP 3000 CPU.
- Up to 32 HP programs using IML/3000 intrinsics, or printers and terminals using Inquiry and Development Facility (IDF), may communicate with the host mainframe at a time over a leased communications line at speeds up to 9600 bits per second, subject to normal performance considerations.
- Can share multidrop leased lines with existing bisync IBM 3270 control units.

HOST COMPATIBILITY

- IML/3000 is compatible with:
 - Mainframe teleprocessing/timesharing software: CICS, IMS DB/DC, TSO.
 - Standard versions of BTAM, TCAM, and VTAM.
- IML is currently not compatible with CMS (Conversational Monitor System).
- Note. To be used with IML/3000, host applications must use IBM 3270 screen sizes of 480 or 1920 characters and transmit data blocks not exceeding 2048 characters. Your Hewlett-Packard Systems Engineer will qualify your host operating environment prior to ordering of IML/3000 to determine whether IML/3000 compatibility has been verified in an environment similar to yours.

HP 3000 REQUIREMENTS

- Requires Intelligent Network Processor (INP) hardware interface for the HP 3000.
- Terminals supported by Inquiry and Development Facility (IDF): HP 2640B, 2640N, 2640S, 2645A, 2645N, 2645S, 2647A, 2648A connected to Asynchronous Terminal Controller (ATC) or Asynchronous Data Communications Controller (ADCC); HP 2645A 2647A, 2648A terminals connected by Multipoint Terminal Software (MTS/3000).
- Terminals using Inquiry and Development Facility require 8K memory option (12K for Multipoint) and display enhancement board.
- Printers supported by IDF:
 - HP-IB based systems: 2608A, 2631A.
 - Non-HP-IB based systems: 2608A, 2613A, 2617A, 2619A.

PRICE IN U.S.A.:

32229A License to use, \$4200.
32229R License to copy, \$1680.

MANUFACTURING DIVISION: DATA COMMUNICATIONS OPERATION

19447 Pruneridge Avenue
Cupertino, California 95014 U.S.A.

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


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

HEWLETT-PACKARD JOURNAL

NOVEMBER 1980 Volume 31 • Number 11

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., Sugunami-Ku
Tokyo 168 Japan



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