


HEWLETT-PACKARD JOURNAL

MARCH 1985



Articles

4 **HP Maintenance Management: A New Approach to Software Customer Solutions**, by Joseph L. Malin and Irving Bunton *No standards existed for computerizing plant maintenance. The solution? Constant user feedback for the designers.*

9 The Need for Plant Maintenance

11 **Development of a High-Performance, Half-Inch Tape Drive**, by Hoyle L. Curtis and Richard T. Turley *Close coordination with the host computer designers and compatibility with earlier drives were key considerations.*

13 LSI Simplifies Tape Drive Electronic Design

15 System Integration

16 **Write and Read Recovery Systems for a Half-Inch Tape Drive**, by Wayne T. Gregory *Besides the necessary erasing, reading, and writing functions, it is important to protect the data from accidental alteration or destruction.*

19 **Digital Formatting and Control Electronics for Half-Inch Tape Data Storage**, by Jimmy L. Shafer *Encoding and decoding data transparently in either GCR or PE formats require fairly complex operations. LSI circuits simplify some of the design problems.*

25 **Streaming Tape Drive Hardware Design**, by John W. Dong, David J. Van Maren, and Robert D. Emmerich *Two microprocessors are required—one for master control, the other for servo control.*

29 **Firmware for a Streaming Tape Drive**, by David W. Ruska, Virgil K. Russon, Bradford W. Culp, Alan J. Richards, and John A. Ruf *Support of queued operations keeps the tape streaming and handles all interactions with the user or the host computer.*

31 Authors

34 **Low-Cost, Highly Reliable Tape Backup for Winchester Disc Drives**, by John C. Becker, Donald A. DiTommaso, and Sterling J. Mortensen *Store up to 67 megabytes on a quarter-inch tape cartridge.*

36 A Design Methodology for Today's Customers

37 **Tape/Disc Controller Serves Integrated Peripherals**, by Craig L. Miller and Mark L. Gembarowski *A fixed disc drive or a quarter-inch tape drive or both can operate in a common environment.*

39 **Cartridge Tape Data Integrity Ensured at Five Levels**, by K. Douglas Gennetten *The drive has read after write, error correction, and media monitoring capabilities.*

44 **Controlling the Head/Tape Interface**, by Walter L. Auyer, Charles H. McConica, David J. Schmeling, and Mark E. Wanger *It's critical to data integrity and unit-to-unit interchangeability*

47 **Software Methodology Preserves Consistency and Creativity**, by Mark L. Gembarowski *Concepts used include hierarchy charts, structured programming, top-down design, structured walkthroughs, and structured analysis.*

In this Issue



Magnetic tape, as two of our authors point out on page 11, is associated with computers so closely that a picture of a reel-to-reel tape drive is often used to convey the image of a large computer. Although it's now considered too slow for on-line storage, tape continues to be used for off-line storage, disc backup, data interchange between computers, transaction logging, and software distribution. In this issue, two sets of articles deal with the design of a pair of new HP tape drives. Both drives are designed primarily for backing up on-line disc memories. (In a well-managed computer system, the continually changing contents of the system discs are periodically dumped to tape so that, should some catastrophe wipe out the data on the discs, the system's memory can be restored except for transactions that have occurred since the last backup operation.) Described on pages 11 to 31 is a half-inch, industry-standard, nine-track reel-to-reel tape drive, the HP 7978A. Designed to provide reliable, low-cost backup for fairly large systems, it can place up to 140 megabytes of data on a standard 10.5-inch reel of tape. The other tape drive described in this issue, the HP 9144A (pages 34 to 48), is a quarter-inch cartridge tape drive designed to back up smaller systems using Winchester-type disc drives or the equivalent. It can store up to 67 megabytes of data on a single tape cartridge. Although both design teams had high reliability and low cost as primary objectives, their designs diverge in interesting ways. Both drives use a recording density of about 10,000 bits per inch, but the half-inch drive records its nine tracks of data in a single pass of the tape, while the quarter-inch drive records 16 tracks of data one track at a time, going back and forth from one end of the tape to the other.

Our cover this month associates computers with maintenance, illustrating the subject of the article on page 4. HP Maintenance Management, a software package for the HP 3000 Computer, helps factory maintenance departments schedule work and keep track of spare parts. Its designers are proud of the fact that it was conceived by an HP customer and designed with exceptionally close customer contact to make sure that customer needs were met, there being no accepted standards for this kind of software. HP Maintenance Management uses the customizer technology first described in our April 1981 issue. The software can be customized to meet the needs of individual companies and is fully supported by HP.

-R.P. Dolan

What's Ahead

In April we'll have articles on the design of the HP 2392A Terminal, a low-cost, ergonomically designed video display terminal, and the HP 7550A Plotter, an extra-high-speed eight-pen XY plotter designed to do a lot of plotting in a short time without operator assistance.

HP Maintenance Management: A New Approach to Software Customer Solutions

Suggested by an HP customer and designed with extensive customer feedback, this HP 3000 software helps cut the cost of equipment maintenance.

by Joseph L. Malin and Irving Bunton

INDUSTRIAL MAINTENANCE COSTS, it has been estimated, are increasing at a rate of 10% to 15% per year. Many maintenance managers state that their departments spend 40% to 50% of their time on nonmaintenance activities.

HP Maintenance Management, an applications software package for HP 3000 Computers, is designed to alleviate these problems by helping maintenance managers increase machine uptime, reduce spare parts inventories, and increase personnel productivity. It is a standard product that is fully supported by Hewlett-Packard and can be customized to meet the needs of individual customers.

HP Maintenance Management represents a coming together of creative techniques on many levels of the computer software development process. In designing a product for which there were no existing standards, the HP Maintenance Management team went directly to customers for information on how the product should work. We listened closely to customer requirements, focused on solving customer problems, and were determined to get HP Maintenance Management up and running as quickly as possible. This approach represents a new direction for HP software.

How It Works

The typical plant maintenance operation keeps track of the work it is doing by means of work orders. A work order is initially a request for service. This request is written down, usually on a standard form, and then is used by the maintenance department. A maintenance planner updates the work order with the specific instructions for handling the request, the location of the equipment to be serviced, and other information needed to complete the service. In the process of doing the work, the responsible technician further updates the work order by listing the spare parts used and the hours of labor, making comments about the cause of breakdown, and noting the action taken to do the repair. At the completion of the work, the work order is filed as part of the permanent record of the maintenance department.

HP Maintenance Management takes this process several steps further. The work order is computerized and stored in a data base on disc. The HP Maintenance Management system establishes links between the work order and records of labor and spare parts. The work order can be further subdivided into tasks to include more detail about the work. Placing this information on a computer improves

the standards, completeness, and visibility of maintenance information without adding undue complexity.

Labor Planning

HP Maintenance Management addresses the problem of labor planning by providing a means of allocating skills to a task. Craft records are associated with each task; these record the skill code, number of people required, and work time associated with the task (Fig. 1). By using on-line and batch reports, the maintenance planner can analyze the demand for the available technicians and schedule work order starts to balance the load. As a result, the technicians' time is used more effectively. More time becomes available for preventive maintenance work, and the technicians are taken out of a "fire-fighting" mode.

Preventive Maintenance

Preventive maintenance planning is done in HP Maintenance Management through the use of standard work orders and tasks (Fig. 2). These serve as templates for the general preventive work orders. All of the templates are initially loaded into the system. A special batch job then generates actual work orders from them, as if requests for preventive maintenance had been received. A new work order is scheduled a selected time interval after the present one is either started or completed. This flexibility meets the needs

Craft Code	Lead Person	People Needed	Estimated Labor Hours (hhh:mm)
1111	MALIN	1	2:00
1112	BUNTON	2	4:00

Fig. 1. Once a work order is entered, skills are allocated to each task by entering the data on this screen.

of most customers.

Spare Parts Inventory

The local automobile muffler shop, with its interminable racks of new and used mufflers, pipes, gaskets, etc. is perhaps the most familiar example of how a maintenance operation depends on the availability of spare parts. The muffler shop's problems of having the right part for the job is magnified tremendously in a typical large manufacturing plant. The spare parts manager must ensure that the spare parts needed for the most common repairs are on hand. The parts on hand have to be allocated among competing work orders. Costs have to be minimized by not stocking rarely used parts, obsolete parts, or parts that can be predictably ordered well in advance of their use. In many shops this knotty problem is approached in a heuristic (read "seat of the pants") manner, with mixed results.

HP Maintenance Management solves this problem with an automated inventory control system. The basic element of this system is a part, just as in the work order system it is the work order. For each part in inventory, one or more stocking locations can be recorded. The system then maintains an up-to-date record of which parts are on hand and in what quantities (Fig. 3). Just using the spare parts inventory system to record what parts are in stock can speed up the maintenance process. Finding parts becomes a matter of minutes instead of hours.

Parts and Orders

The spare parts inventory system is connected to the work order system by allocations. An allocation is a reservation for a part required for a particular work order. Every time a work order is assigned to a technician and started, a document called a pick list is automatically printed (Fig. 4). This instructs the stockroom clerk to pull a set of parts and set them aside for a particular work order. Thus the technicians will always have available all of the parts necessary for the job when they begin. This saves the time involved in constantly going to and from the stockroom.

Moreover, the allocations provide for the accurate plan-




Fig. 2. Standard work orders are used to schedule preventive maintenance. On this screen, RC indicates a recurring (preventive maintenance) order.

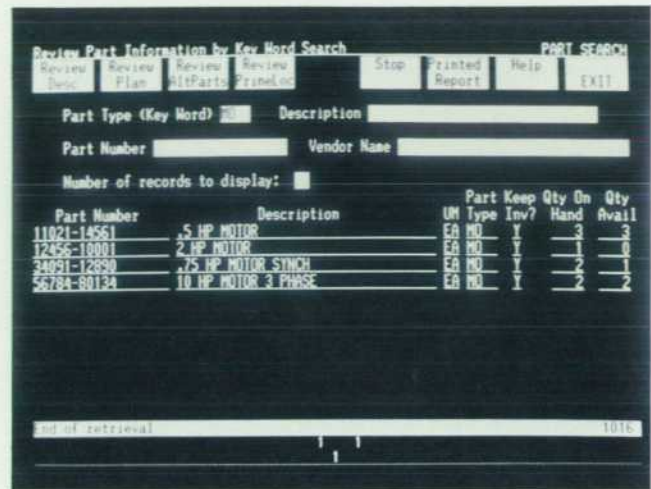


Fig. 3. This screen can be used to find out whether a particular type of part is in stock. The listing here was produced in response to the part type MO (motor) entered in the appropriate field.

ning of work orders. When allocations are initially added, they are considered soft, and merely serve as documentation. When the planners decide to commit to a work order, they also commit its allocations. However, the system will only allow allocations to be committed if there is uncommitted inventory for that part (the system tracks the level of uncommitted inventory internally). The system can be set up so that any attempt to commit past the level of uncommitted inventory for a part automatically generates a back order request. When a work order is started, its allocations are released, causing the pick list to be printed.

This reservation system helps maintenance planners avoid the overallocation of parts. Overallocation usually results in a work order holdup, in which technicians are assigned and work starts without the necessary parts being available. Reducing overallocation results in faster work and less wasted time. Reservations also provide a feedback loop between work orders and parts. The schedule for work orders can be determined by the availability of parts, and the demand for parts is determined by the scheduling of future work orders.

Organization of the System

HP Maintenance Management consists of three applications software modules:

- Work Order Control
- Parts Catalog
- Inventory Control and Purchase Order Tracking.

Fig. 5 shows these modules and their various functions.

Work Order Control

The Work Order Control module helps plan and schedule maintenance work including routine repair, preventive maintenance, emergency orders, standing orders, projects, and shutdowns. A maintenance engineer, operations foreman, or mechanic identifies work that must be done and enters a work order into the system. Work orders can be identified by type, such as a lubrication work order or a normal repair order. A priority can be assigned to indicate

the urgency of the work required. In addition, work orders can be identified by type of job performed to differentiate between discrete, standing, and recurring preventive maintenance orders. Work orders and tasks can be reviewed on-line by various selection criteria (Fig. 6).

The planner begins a work order by estimating the tasks, spare parts required, and costs. A new order can be entered by copying a previous order or work plan. The crafts assigned to the work order tasks are entered into the system. The next step is to authorize the work and if necessary, print a copy of the order. Safety procedures or special instructions for each task are also printed.

The Work Order Control module has the flexibility to allow personnel to enter work order information after work has been done. This ensures that the system does not get in the way of responding to emergencies or short jobs. Task information for quick jobs can be dispatched and printed at remote printers.

Management reports such as the total estimated labor, material, and additional costs for all open work orders can be displayed graphically using HP Decision Support Graphics (DSG/3000). Fig. 7 shows an example.

If the Inventory Control module is used, parts can be allocated to the work order. The selection process is done on-line by reviewing the spare parts bill of material for the piece of equipment and selecting the parts required for the repair.

In daily or weekly planning meetings, planners and supervisors agree on which orders to start. Those orders are approved and parts are released. With the Inventory Control module, the pick list for the spare parts is printed when the order is released (Fig. 4). The inventory control clerk uses this pick list as authorization to pull spare parts




Fig. 5. Block diagram of the HP Maintenance Management software package for the HP 3000 Computer. The principal modules are Work Order Control, Parts Catalog, and Inventory Control and Purchase Order Tracking.

from the stockroom.

While an order is in process, repair information is entered, such as the root cause of the failure and the actions taken (Fig. 8). This information is useful for later analysis. Once the work order is completed, the maintenance foreman or scheduler verifies that the work is done and closes the work order.

Closed work orders are maintained in a history file. Information such as the root cause of the failure and actual labor and material costs are also kept in the history file. Historical information can be reviewed on-line. For example, all pre-

PICKING LIST

WAREHOUSE 01 PICK LIST ID 33

Order Number 1784 Task Id 3 Extra Usage Id

MATERIAL REQUISITION PARTS

LOCATION	PART NUMBER	REQUISITN NUMBER	QUANTITY TO PULL	UM	QUANTITY PULLED	DESCRIPTION	DELIVERY POINT	CRITICAL PART
ROW B	6914-5928	1	60 EA		WASHERS	ELECSHOP	N
ROW B	6915-6126	1	30 EA		NUTS, 8-32	ELECSHOP	N

BACKORDERED PARTS

WHSE	LOCATION	PART NUMBER	REQUISITN NUMBER	BACKORDER QUANTITY	UM	PARTIAL FILL QTY	DESCRIPTION	DELIVERY POINT	CRITICAL PART
01	ROW B	6913-4728	1	30	EA	0	1.5" SCREWS, 8-32	ELECSHOP	N

OPEN STOCK PARTS

There are no OPEN STOCK parts for this order / task.

Fig. 4. When a work order is released, a pick list is printed, showing the material required to complete the work.

Review Work Orders by Selected Criteria							REV. NO.	TIME
Change Task	Update Craft	Change NO	Delete or Task	Stop	Printed Report	Help	EXIT	
Status/Delay	Order Mbr	Type Class	Supervisor	AssignTo	Sched Start	Range		
						(mddy)		
Status	Delay	SchedStart/Complete	Supervisor	AssignTo	Type	Class		
PLAN	NONE	02/01/85/02/14/85	BLINDH	BUNTON	RC			
	56821							
PLAN	NONE	02/11/85/02/13/85	BLINDH	JWALIN	RC			
	47089							
PLAN	NONE	02/25/85/03/01/85	LUCINA	LEUNG	RC			
	44769							
There are no more orders to display								
							60%	
							1 2	
							1	

Fig. 6. Work orders and tasks can be reviewed on-line by selected criteria. Here all of the work orders listed have a status of PLAN.

vious work for a particular piece of equipment can be reviewed. When appropriate, historical information is archived to tape.

Parts Catalog

One of the challenges in the maintenance department is to maintain accurate information about equipment and the spare parts used to repair the equipment. In many cases, this information is not kept up to date or the files are difficult to find. HP Maintenance Management helps gain control over this vital aspect of the maintenance department.

Using HP Maintenance Management, maintenance engineers identify and document the equipment that they maintain and the spare parts used for the equipment. The bill of material or structure relationships for the equipment and spare parts are numbered and entered into the system. The bill of material makes it possible to review all parts for a piece of equipment, and to review all the equipment

where a particular part is used. When the Parts Catalog is used with the Inventory Control module, inventory control information, including primary stocking location and minimum order quantities, is also entered.

Generic searches help identify parts or alternate parts when one is unsure of the specific part number. For instance, a search of parts whose description includes "3-inch valve" can be made.

Inventory Control and Purchase Order Tracking

Managing the inventory of spare parts is a necessity for the maintenance department. The first step in controlling inventory is to maintain accurate inventory counts and to identify where each spare part is located. Accurate inventory balances must be available to the maintenance planner so the proper spare parts can be planned for each order.

The Inventory Control and Purchase Order Tracking module is used to track purchase orders and maintain inventory balances of all spare parts. A terminal is used for on-line purchase order receiving. Inventory quantities and left-to-receive quantities are immediately updated when purchase orders are received.

If appropriate, the purchase order is closed. A receiving document is printed to accompany the parts or the purchase order for accounting control. If back orders exist for the part being received, a document can be printed indicating the quantities on back order and the departments and orders requiring the parts. Back orders are filled in date or priority order. Material can be issued from the stockroom for a work order, a back order, or extra use requirements. A pick list, printed by releasing a work order on-line, is used to identify which parts should be issued. All parts required for a particular work order are issued by entering just the work order number and an identifying number from the pick list. Issuing can also be done by exception when the need occurs. A partial issue can automatically create a back order.

A purchase order to be issued to a vendor is entered at

LMP Corporation
Maintenance Department - number of tasks for each status




Fig. 7. Management reports can be displayed graphically and plotted using HP Decision Support Graphics (DSG/3000).

Enter Repair Information for a Task				ENTER REPAIR
Change Task	Change NO	One Step	Workford Review	Rev Fail Task
				Help
				EXIT
Order Number	Task ID	Equipment ID	Location	
02871		CHEM HEATER # 1	J11 / K78 / /	
Task Description			Status/Delay	
REPLACE CORRODED IMPELLER			COMP / NONE	
Technician	Equip Status	Repair Hours	(hhh:mm)	
JOHN SMITH	WIP	0:30	(hhh:mm)	
		Equip Down	0:00	(hhh:mm)
Comments: NEEDED A NEW IMPELLER				
Root Cause	Component	Cond	Action	
6 12				
1				

Fig. 8. When the work is complete, the maintenance technician uses this screen to record the cause of the failure. Failure data can be reviewed at a later date, sorted by cause or other criteria.

a terminal in the purchasing department and includes necessary tracking information such as quantities and due dates as well as descriptive information. Blanket purchase orders can be identified. The purchase commitment report shows weekly cash requirements to meet the purchasing plan.

Each part can be stored in multiple stocking locations and in different warehouses. When multiple locations are used, receipts and issues are tracked for each inventory location. There is an audit trail of each transaction that causes a change to inventory quantities. This audit trail can be reviewed on-line.

Many work orders in the maintenance department are planned for shutdowns or major projects. The spare part requirements for these orders are reserved to ensure that the parts will be available when work is started.

The Shortage/Preshortage report identifies parts that might run out, based on the parts scheduled for current work orders. An order point technique is used to launch orders when the available inventory level drops below the order point for the part.

Customizable Technology

HP Maintenance Management was developed using the customizable technology tools used in HP Materials Management/3000¹⁻³ and HP Production Management/3000. These tools use VPlus/3000 and Image/3000 to support a data-independent software environment. A central dictionary called APPDIC is used to store information about user application data bases. Software systems such as HP Maintenance Management then use the information in this dictionary (an Image data base) to manipulate the data. In this way, the customer has complete control over the appearance of the data. To change lengths of fields and other data-oriented features, the user need only change APPDIC and not any code. This makes customizable software products not only easier to modify by the customer, but also easier to update and correct by HP. A system administrator can customize parts of APPDIC while users are using the system.

The Application Customizer² is one of the most attractive features of the application. Information about the layout of data bases, VPlus/3000 screens, and transactions is stored in APPDIC. This information is entered and updated with a screen interface. Data bases can be generated and modified with the Customizer. Formspec, the forms design utility, can be accessed through a function key. The user can define names to be associated with the logical device numbers of terminals on which to run the application. The names of many items can be changed without affecting the operation of the application.

The format is one of the important aspects of customizable applications. A format can describe a data set layout, a VPlus screen, or a data file layout. The format contains fields that can represent Image data items, VPlus fields, or fields on a report or file record. The fields can have not only the same data types as Image but also additional types (e.g., 24-hour time, dates, display numeric, etc.). These extra types enable additional data editing without having to specify it in VPlus processing specifications. Each format or field has a unique number associated with it. Although

users can change the format and field names, they cannot change the numbers associated with those names. Fields and formats can also be added by the user.

The formats for data sets can also have Image information other than the fields. APPDIC identifies which fields (data items) are sort and/or search items. It also tells which data set formats are masters, automatic masters, or detail data sets. While users can add automatic masters for data items, they cannot delete those put in by the application programmer.

Users can define additional characteristics for screens. The user can indicate to which screens to transfer control from the current screen, based on the eight function keys. More than one screen can use the same transaction; this enables users to keep screens simple, since every screen does not have to contain every field on a data set. The application programmer specifies which fields are critical for processing a transaction, so that these fields cannot be deleted from the screen.

While users cannot change the code being executed, they can request some additional processing. One way they can do this is with processing specifications. Each screen can have a list of operations to be performed on fields in the formats. Users can also specify where the results of the operations will go. These operations are performed before updates are made to the data base.

If the user finds that more complex operations need to be performed, the advanced customization (user exit) option may be purchased. User exits are customer-written COBOL II procedures, which are loaded using the MPE LOADPROC intrinsic and called during critical portions of code execution. These procedures are passed buffers for the appropriate data sets and the screen. The user can update the buffers used by the application and/or open other data bases and update them. Users can indicate an error by setting the value of a parameter. This lets the transaction know that processing should stop. User messages can also be displayed on the screen by setting the appropriate parameter.

Given that the user can customize the product so much, the application programmer must use special intrinsics to manipulate data, since crucial information about formats and fields is not known until run time. Most of the information in APPDIC is put into extra data segment (EDS) files so that accessing the information is quicker. The application programmer then uses Application Customizer (AC) intrinsics to get this information from the EDS files or to perform operations on the data. There are AC intrinsics that move, compare, add, multiply, divide, subtract, and so on. These intrinsics use the unique field or format numbers instead of the names, since the numbers cannot be changed by customization. Since users can add fields to screens and data sets, the AC'MOVE'CORR intrinsic was developed to move corresponding fields from one buffer to another.

Software Productivity Tools

While the AC intrinsics make customizable applications manageable, the parameter lists and code to load information from the EDS files into memory can make coding more difficult. Consequently, the HP Administrative Productivity Operation (APO) developed a new tool specification

The Need for Plant Maintenance

In facilities without a structured preventive maintenance program, emergencies account for 50%-90% of the maintenance work and unorganized lubrication and service work account for the rest. According to Joseph W. French, a senior maintenance consultant in the Applied Technology Division of Du Pont, the annual maintenance bill for U.S. business is in excess of \$200 billion. This figure is increasing at the rate of 12% per year. Furthermore, as manufacturing plants become more sophisticated, the always low supply of skilled technicians for equipment maintenance has become critical. These technicians can no longer be effectively used by hopping from one crisis to the next. Company executives are being forced to look for a better management system for plant maintenance. The mode of operations has to become planned and predictable.

The cost of production is also going up. High-technology equipment costs money, and production managers are finding that they cannot afford to keep it idle. Yet these high-tech machines are more likely than ever to require regular maintenance to perform at peak efficiency. A computerized maintenance system is a must to keep the complex technology in tune.

Another driving force for better maintenance is the increasing popularity of flexible manufacturing centers and the just-in-time manufacturing technique in certain industries. Here the problem is that production is very sensitive to unplanned interruptions. A high degree of coordination between production and maintenance is needed to keep these types of factories successful. In a just-in-time system, the failure of one machine can shut down an entire assembly line.

In the past, managers of maintenance departments have not had the tools to handle the enormity and diversity of even standard operations. For many years this was a tolerable situation; the inefficiency of maintenance was tolerated because maintenance was an overhead operation that wasn't seen to be integral

to the success of production. The current reemphasis on manufacturing, especially high-quality automated manufacturing, has changed this attitude.

Some industries have always known the importance of plant maintenance. The customer letter that started HP Maintenance Management was from a petrochemical processor. Maintenance has always been a vital part of the petrochemical industry, and indeed of most of the process industries, because the breakdown of even one key part can disable a many-miles-square plant. In 1979, the top ten U.S. petrochemical companies alone spent \$3.4 billion on maintenance, 2% of the total expenditure in the U.S.A.

More recently, the capital-intensive industries, which use large numbers of very expensive machines, have begun to make increased commitments to getting full and efficient use out of their machinery. These companies, typified by the automotive and aerospace industries, are trying to save production expense by making better use of maintenance.

High-volume electronics manufacturers use the just-in-time manufacturing technique to improve product quality and reduce inventory costs. The success of this method of manufacturing depends directly on minimum machine downtime. While the machinery itself might not be expensive, the cost of having it and the workers idle can be astronomical because of the high volumes involved. Moreover, improperly maintained equipment introduces poor quality that the production technique is designed to eliminate.

There is clearly a need for an organized method of doing plant maintenance. The key to HP Maintenance Management is its focus on the central problems of plant maintenance: control or requests for service, control of spare parts, control of skills use, and planning of maintenance operations.

system called Toolspec, along with customizer tool (CT) intrinsics, and a Toolspec "compiler" (actually a code generator) to make the use of AC intrinsics easier. The Toolspec compiler generates code for a transaction that will use AC intrinsics to load the formats from the EDS files and get space for the data buffers. The compiler gets its information from TOOLDB, a data base that contains information on which formats, fields, and work files are being used.

With its screen interface, the Toolspec program allows the entering of buffers and fields being used. Application programmers can specify what names they want to use to refer to data buffers for particular formats for a given transaction. Application programmers can also specify variables, which are names for given transactions for fields on particular buffers. Work files with sort criteria can also be specified. The names of the buffers and variables can be used in CT intrinsics, which call the AC intrinsics. Although most field operations on fields with AC intrinsics require the format number, field number, and buffer address, the CT intrinsics only require the specification of the variable. For example, AC'ADD, which adds two fields and puts the result in a third field, requires 12 parameters. The corresponding CT intrinsic, CT'ADD, requires only three parameters. Error messages are also printed by the CT intrinsics.

The new intrinsics also provide improved functionality

to the user. HP Maintenance Management is the first customizable manufacturing system to provide maintain transactions, which can add, change, delete, and review data all in one step (see Fig. 1). In addition, the CT intrinsics allow the use of review-by-criteria transactions, which provide a very flexible data base inquiry system to the user (see Fig. 6). Using these transactions, the user can specify a limited set of fields to be matched to the records on the data base. When the transaction is executed, the values specified for these fields are compared to records in the data base, and records that match all of the specified fields are retrieved and displayed.

The entire product team participated in writing an external specification (ES) based mostly on customer input. Unlike previous ES efforts, which had concentrated mainly on the technical aspects of the software package, the HP Maintenance Management ES attempted to describe how the product would appear to the users. The various life cycles or continuing processes of plant maintenance were described in this document using the jargon of maintenance personnel instead of computer engineers. A simulator was developed that mimicked the functions of the new product. This simulator became part of the ES, enabling the lab team to show the outside world how the product would look when it was completed. This simulator was used to demonstrate HP Maintenance Management nearly two years

before the final product was released.

In this atmosphere there was a lot of risk, but also a chance to make new contributions. There were never preconceived notions about what we should do or could do. A typical result was a comment by one of our first customers, who praised the ES because it had obviously been written by maintenance engineers rather than computer people!

We learned that success is possible in a market where no one has set a standard. The key ingredient is responsiveness. Without widely accepted standards it becomes imperative to maintain constant contact between the customer and the product team. The standards then are written in the process of matching product to problem. To do this, the development process must be short enough so that prototypes of the product can be made available to selected customers throughout the design process.

Perhaps another way of stating these lessons is this: the "next bench" effect⁴ works for software. By using our customers as the "next bench" we were able to meet a need that we might not have seen within the company.

Development of HP Maintenance Management

HP Maintenance Management started with customers. In fact, a customer wrote a letter asking if HP had a system to computerize plant maintenance. It was evident from preliminary investigation that a market for such systems existed, and that HP could make a contribution.

Unlike most of the manufacturing areas for which HP had developed software in the past, no external standards existed in the plant maintenance area to serve as a basis for a software product. Instead, customers provided information from the very beginning. HP engineers went to process companies and to HP field people to gather data about plant maintenance. Because customers wanted the product as soon as possible, a decision was first made to design the product by making only slight modifications to an existing product, HP Materials Management/3000.¹ Using existing software tools, HP Materials Management/3000 could be tailored to meet the needs of the maintenance department instead of the production department. In this way, the new software product could be developed quickly and still remain compatible with the other HP Manufacturer's Productivity Network products.

However, slight modifications to HP Materials Management/3000 were not enough to make the right solution for plant maintenance. The engineers determined that the problem of making firm schedules for manufacturing parts was much different from the problem of balancing work orders, parts, and people in a situation where schedules are always changing. HP Materials Management/3000 was too rigid in some places, and too flexible in others. For instance, the major goal of HP Materials Management was reducing inventory to a minimum level. As a result, parts without future demand gradually would disappear from inventory. This was unacceptable to HP Maintenance Management since the future demand for a spare part is often not known.

The stage for innovation was set. Although Materials Management/3000 itself was not used, a good portion of

the Materials Management/3000 code was modified for use in the new system. The leverage gained by this move probably saved us over a year in development time.

Acknowledgments

So much work obviously could not have been done by just a few people! In fact so many people have had a hand in the success of HP Maintenance Management that it would be impossible to list them all here. Instead, we would like to focus on the people who really made a substantial difference: Royal Linden, our former project manager, now a section manager in the Financial Systems Operation of the Administrative Productivity Division, Steve Baker, our former product manager, now a product marketing manager for the Manufacturing Productivity Division, the HP Maintenance Management team, which has over the years included (besides the authors) Marc Barman, Lois Watson, Karen Bunya, Tom Holden, Nancy Lipstein, Bob Passell, and Toshi Hirasawa. We would also like to thank Neal Streit, the product manager for HP Maintenance Management, Steve Dentel and Barbara Chang, the product support engineers, Steve Law, Doreen Petersen, Daphne Quintana, Ken Jercha, and Roz De La Torre, who set very high standards of quality while producing the documentation and training for the product, Walter Utz and Mark Bussman, former section managers, who provided the much needed support and faith to keep the project alive, Bill Sudlow, our current section manager, for doing the same today, Bob Greenfield and Norma Sengstock, who supported us from the field, and Brenda Leung and project manager Susan Luciano, the newest additions to HP Maintenance Management's continuing success.

References

1. N.C. Federman and R.M. Steiner, "An Interactive Material Planning and Control System for Manufacturing Companies," *Hewlett-Packard Journal*, Vol. 32, no. 4, April 1981.
2. L.E. Winston, "A Novel Approach to Computer Application System Design and Implementation," *ibid.*
3. B.D. Kurtz, "Automating Application System Operation and Control," *ibid.*
4. W.R. Hewlett, *Inventions of Opportunity*, Hewlett-Packard Company, part no. 92233B, 1983, p. vii.

Development of a High-Performance, Half-Inch Tape Drive

The design of a low-cost, high-density tape drive for backup of large amounts of on-line computer system memory requires a sophisticated combination of technologies and careful project planning. This new drive's greatly improved reliability reduces maintenance costs and downtime.

by Hoyle L. Curtis and Richard T. Turley

SATISFYING EXISTING AND FUTURE customers with products that are price and performance competitive remains a constant goal in the computer equipment industry. Adding to the challenge is the desire to maintain compatibility with earlier generations of devices while making the continuing performance improvements and cost reductions that customers have come to expect from electronic product vendors. This is particularly apparent in the technology of tape drives used as computer peripherals.

Because of its versatility, a half-inch tape drive is often selected to provide backup of computer system memory, allowing recovery from system failures or other instances of data loss. It is historically associated with the computer industry to such an extent that often just the picture of a reel-to-reel tape drive is used to convey the image of a mainframe computer.

This identity is well-founded because of the different roles that half-inch tape has played in the growth of the computer industry. It was first used in 1952 by IBM in their Model 726 tape drive, which had a storage density of 100 bits per inch (bpi). Improvements have come on a regular basis with storage densities of 1600 characters per inch introduced in 1965, and 6250 characters per inch in 1973. Half-inch tape can be used for on-line memory, off-line storage, backup memory, data interchange, transaction logging, and software distribution. Although its use as on-line memory has all but disappeared because of its slower access speed, the other uses continue to flourish. Because of its early use in most computer systems, it is a universally accepted medium for data interchange between computers from different manufacturers. This has been assured by international standards such as those administered by the American National Standards Institute (ANSI).^{1,2}

Hewlett-Packard's newest half-inch tape drive, the HP 7978A (Fig. 1), is an industry standard nine-track, reel-to-reel drive that provides improved functionality and reliability at a significantly lower cost. This tape drive is designed to provide HP computer systems with a reliable backup device for large amounts of on-line disc memory by placing up to 140 megabytes on a standard 10½-inch reel. It also provides for the interchange of data with other HP or non-HP computer systems by supporting both 6250 GCR (group coded recording) and 1600 PE (phase encoding)

ANSI formats.

Compared with the earlier HP 7976A Magnetic Tape Subsystem which it replaces, the HP 7978A makes significant contributions in several areas. The HP 7978A is half the cost and five times more reliable. Additional improvements include better data reliability (full parity checking on internal data paths), half the size, less than half the power consumption at under 400 watts, lower acoustic noise (less than 60 dBA), and improved serviceability (30 minutes




Fig 1. Designed to back up computer systems with large amounts of on-line disc storage, HP's newest half-inch tape drive, the HP 7978A, features high performance and increased reliability for a low cost. This streaming tape drive uses both 6250 GCR and 1600 PE storage formats.

MTTR). These attributes were gained by combining a market-driven product definition, a good technical implementation, and proper project execution. While these three areas may seem obvious, their correct sequencing and interaction are the heart of both the product and the project.

Product Definition

When HP's Greeley Division surveyed the market for contributions that could be made in half-inch tape drives, they found cost of ownership (purchase price plus maintenance costs) to be the highest on the list of customer concerns. Increased reliability was also found to be an area where HP could make a contribution. The study noted that the complexity of existing tape drives was primarily responsible for their high cost and low reliability. This complexity was required to provide the performance needed to achieve a high data rate. However, raw performance means little if a system cannot use it, so extensive modeling of entire computer systems was conducted.

Prices in the tape drive industry had not followed the cost reduction trends experienced by the rest of the computer industry, making the cost of a tape drive an increasingly larger percentage of the system price. An astounding increase in magnetic disc storage on systems meant that backup time had become an increasing burden on system uptime. This trend created a demand for higher capacity and faster transfer rates, which could only be satisfied by a drive capable of using the 6250 GCR format. Because of the large tape libraries at most computer centers and the need to interchange data, it was apparent that the older 1600 PE format should also be usable in a new tape drive.

The task then became the definition of a tape drive that could satisfy all of the customer needs and that would fit into the HP computer family. While this may seem straightforward, the integration of a new peripheral into an extensive family of mainframes and a large diverse software base proved to be no mean feat.

A number of studies were done to determine whether current standards (either industry or HP) and/or leveraged designs could be used in the HP 7978A definition, since

this would reduce engineering effort and make development easier in general. The most obvious of these were the previously mentioned ANSI data format standards for 6250 GCR and 1600 PE.^{1,2} Others included the HP-IB (IEEE 488) interface and a command set leveraged from earlier tape and disc drives. A powerful microprocessor (the 68000) was chosen because development tools were already available for it in-house, left over from the design of the HP 9000 Model 226 Computer.

The choice of 6250 GCR as the recording format had significant implications. This is currently the most advanced and complex recording technology used for tape. It requires real-time encoding of data into a nine-track format with three bytes being added to every seven bytes of data. Reading GCR data requires real-time decoding and error correction of up to two bits per byte. A density of 9042 magnetic flux reversals per inch must be placed on the tape. To optimize performance, a tape speed of 75 inches per second (ips) is used for both reading and writing. A rewind speed of 250 ips is used. The resulting drive provides a data rate of about 470,000 bytes per second.

Technical Implementation

An architecture had to be defined that would allow HP's technology, the market definition, and areas of leverage to be combined into a viable product. See Fig. 2 for a high-level diagram of the HP 7978A architecture. This architecture permitted the establishment of a detailed, low-level product definition based on modularity with minimum connectivity. Each module was then divided into sub-modules that reached a level of complexity that could be assigned to individuals for implementation. A further benefit of the early architecture definition was allowing tests to be designed for functional modules, both for internal diagnostics and for manufacturing test.

After the overall architecture of the HP 7978A was established, it was possible to review the technologies that could be best used to implement the product. The following technologies became crucial to the final performance and reliability of the HP 7978A:

(continued on page 14)




Fig. 2. Basic architecture of the HP 7978A Magnetic Tape Subsystem.

LSI Simplifies Tape Drive Electronic Design

LSI technologies of various forms are used in the HP 7978A Magnetic Tape Subsystem. The decision to use LSI was made early, allowing certain advantages during the development of the product. There are a number of good reasons to use custom and semicustom LSI circuits in addition to standard off-the-shelf LSI components. Among these are reductions in cost and size. Although not of direct importance in a large tape drive, certain advantages come from size reduction. The reduction in parts count directly affects large system problems such as noise, propagation delay, and timing. In addition, the use of LSI circuits reduces the need for complex designs, thereby reducing the problems associated with these designs.

When LSI components are used, the power supply requirements are greatly reduced, because the number of parts required is vastly reduced. An example of this is the HP 7978A's read formatter electronics. Were it not for the use of two LSI parts, hundreds of off-the-shelf parts would have been required (see article on page 19).

Architectural freedom is a prime motivation for using custom and semicustom LSI components. Having the freedom to decide the discipline best suited for design and being able to implement functions not found in standard off-the-shelf parts allows the designer to go straight to the solution with fewer restrictions.

Design Philosophy

The very nature of LSI technology lends itself well to the "Do it right the first time!" strategy used within quality circles. In keeping with this strategy, the approach for the HP 7978A was to design for LSI from the start. This affected the architecture and design strategy a great deal.

A primary consideration was tool development. HP's Greeley Division's strategy was aimed at developing tools for LSI design and making those tools available to every designer. In addition, the support of development projects such as the HP 7978A was of primary concern. Each custom LSI chip developed for the HP 7978A's electronics was done by two designers and one of the designers was provided by the HP 7978A design team. An efficient mix of LSI design skills, tools, and product knowledge resulted from this approach.

The Mead/Conway approach to LSI design was chosen.¹ This incorporated the use of lambda-based design having simplified electrical and topological rules. Additionally, having two "tall-thin" designers on each custom part helped guarantee success. This concept simply means that the designers are knowledgeable in every aspect of the chip design from the highest level of architecture down to the layout of individual cells.

Software Simulation

The need to check the design to verify correctness led to the use of software simulation of LSI on the HP 7978A. This was chosen over emulation techniques for several reasons. First, it was very easy to try new ideas and make modifications in software. In addition, the designer could choose the level at which the design was simulated, thereby simulating only the most critical timing situations at the transistor level and keeping the remainder of the design at a much higher level.

Closing the loop on a design requires two steps. Initially, the functional design must be completed. This requires the designers to define the function of each subsystem thoroughly, beginning at the topmost level. From the top, a functional decomposition

must be done to yield concise blocks that can be fit together into a floor plan. These blocks should be fairly autonomous. Software simulation of the architecture gives the designers confidence in the design and verifies algorithm correctness.

Once the functional design is complete, the second step, verification, requires a transistor-level layout simulation. The simulator we used was extracted directly from the layout of the chip. Verification requires that the chip be simulated using the same stimulus as the functional simulator and the results compared with the simulator's results. Verification of the layout in this manner guarantees correctness, with the exception of critical timing situations.

Cells with layouts that are timing critical were checked using HPSpice, a simulation program used to model the analog behavior of a circuit. It is useful in modeling digital circuits where timing information is important. An HPSpice deck was extracted from the cell layout data and used by HPSpice. Once critical timing matters were resolved, the layout simulator was used to guarantee correctness.

This method of software simulation requires that the designer thoroughly investigate the environment of the chip to create realistic stimuli. The point of vulnerability is the designer's imagination.

Following this procedure resulted in fully functional LSI parts. Testing of these components in the HP 7978A yielded zero layout mistakes and a very small number of design errors.

Architecture

The HP 7978A's formatter function was broken down into two blocks, the data path and the control for the data path. The register-to-register construction of the data path fits well into the structured MOS design methodology. In addition, extensive modulo 2 arithmetic with large numbers of lookup tables and registers can be readily implemented using custom LSI. Once designed, the channel was stepped and repeated to form the parallel 9-bit environment inherent in the HP 7978A.

Controlling the data path using a series of state machines and counters fits well into PLA-based logic. The technology chosen for the HP 7978A was that of the universal synchronous machine (USM) developed at HP's Greeley Division. Its ease of implementation and other benefits made the USM fit nicely into the "designed for LSI" strategy of the HP 7978A project.

A PLA-based synchronous machine allows simplicity of programming in the sum-of-products form. Within the USM, the PLA area is segmented to allow its product lines to be broken and independent or semi-independent PLAs to be constructed. Storage at the PLA outputs in the form of flip-flops and feedback paths facilitates state-machine design. The flip-flops can be of several types to allow even more versatility.

Where great flexibility was needed, or in the case where the design did not have a good functional description, off-the-shelf LSI components such as microprocessors were used. The main controller in the HP 7978A makes use of a Motorola 68000 and the servo controller is an Intel 8051.

Reference

1. C. Mead and L. Conway, *Introduction to VLSI Systems*, Addison-Wesley, 1980.

Jimmy L. Shafer
Development Engineer
Greeley Division

- LSI circuits
- Intelligent systems
- Electronics/mechanics optimizations
- Modeling
- Magnetics
- Analog/digital electronics
- Precision mechanics.

LSI became a critical factor in implementing a high-reliability product by replacing approximately 1200 TTL parts with eleven LSI circuits and 100 TTL support parts. Microprocessor control within the architecture allowed great flexibility and improved stand-alone capabilities of the product such as self-test and diagnostics. (See box, page 15.)

Electronic and mechanical optimizations, such as using random-access memory buffers to match host system performance rather than high-speed and high-acceleration mechanisms to do start/stop operations, allowed more reliability to be obtained and cost reductions to occur. Modeling was an essential part of all design decisions. Models allowed us to predict system performance for various tape drive parameters accurately. A tape drive specification was not selected randomly, but rather based on its effect on overall system performance. A good example is the chosen tape read/write speed of 75 inches per second. A higher speed was physically possible, but would not have provided improved performance in operation with a computer. However, a rewind speed of 250 inches per second could be implemented since no system interaction occurs and a minimal rewind time is a performance goal.

The magnetic characteristics of the tape and the read/write head were viewed both as a magnetics problem (waveforms and bit shift) and as a communication theory problem (transmission of data through various transforming media). At all times, data reliability was the most important design criterion.

Analog and digital electronics played roles in each sub-

system. The read amplifier was an especially sensitive area with large signal amplification needed and minimal noise injection required for low error rates. The data buffer area had to be designed both to accept and to transfer data asynchronously, which led to a semicustom LSI controller design.

Precision mechanics is the heart of the tape handling system. The tape is sensitive and delicate and must be guided precisely. Many component tolerances are held to a few ten-thousandths of an inch to ensure a parallel tape path from reel to reel. At 9042 flux reversals per inch, the data can be easily distorted or destroyed unless exact speed, acceleration, and tension control are maintained.

Working in the above technologies required a major commitment to design tools. These tools came from various sources (HP proprietary, commercial CAD vendor, HP 64000 Logic Development System, etc.).

The architecture permitted design trade-offs that minimized real effects and used HP's strengths to a maximum. This was particularly noticeable in the decisions on hardware versus software implementations. Major revisions, such as the addition of "immediate response" to improve system performance (see page 30), could be made without a total redesign. The architecture was firmed up by specifications that defined each design area. These specifications were considered to be contracts between the design engineers. If any deviation was proposed or required, a formal change in the appropriate specification was made. With boundaries clearly defined, designs and subsystem testing could be carried on concurrently. By using this technique, designs with significantly different characteristics (LSI, software, mechanics, etc.) could be combined into a single product and efficiently managed. The specifications and other related documentation were put together in a book that was updated regularly and republished for anyone who needed data on the product or its design project.




Fig. 3. Organizational growth model as discussed by Greiner.³

Program Management

Since this was the largest project that HP's Greeley Division had ever undertaken, the plan for the project implementation received early attention. It was also very apparent that more than just an R&D team had to be assembled if the HP 7978A project was to achieve the improvements that other functional departments felt to be essential. These covered the range from accurate cost estimates to manufacturability goals. Meeting HP's corporate objective of significantly improving product reliability, the HP 7978A design goal of lowering the failure rate of a half-inch tape drive by a factor of five required a divisional commitment.

Thus, the project became a program very early. This program was centered around a business team, which drew members from all departments in the division and the R&D team. The milestones that determined when the project was ready to move to the next phase of development were established by these two teams. These goals formed a formal development plan, which was published and used as the final test to judge if a development phase was complete. This *a priori* definition avoided many arguments about what constituted a phase and limited discussion to considerations of whether or not a fixed set of goals had been

System Integration

The design of the HP 7978A tape drive was a challenging task involving large project groups and new technologies. Yet the completion of the product effort was only the first step in making the product successful. Without host system support, a new tape drive cannot be used for its intended purpose. For all of its capability, this new drive would remain unable to perform any useful function without complete integration of the product into computer systems. Initially this might seem as simple as writing a quick software routine to transfer data from system disc drives to the HP 7978A tape drive. However, an HP 3000 computer system involves a very complex set of interacting subsystems which all need to be considered when a new product is supported. The HP 3000 file system, tape drivers, diagnostics, and intrinsic system functions had to be modified to take full advantage of the advanced features available in the HP 7978A. In addition, complete environmental, EMC (electromagnetic compatibility), safety, reliability, and performance testing was required on all supported system configurations. Thus the addition of an HP 7978A drive to an existing HP 3000 system represented a complete product design in itself.

There are a few keys to successful system integration. These keys are:

- Designing the right product
- Communicating well
- Anticipating needs
- Using in-house expertise.

Designing the right product is a simple rule often overlooked by product designers. Attributes of the right product include low cost, high reliability, and high performance. However, just as important are the attributes that make it easier to integrate the product. Consistent architecture and I/O requirements make the design of drivers easier. Advanced internal diagnostics must complement system-level diagnostics. Complete and consistent documentation of the product contributes to easier integration.

Timely and accurate communication is the key to eliminating confusion and providing rapid feedback on problems. The HP 3000 Computer is produced by HP's Computer System Division in Cupertino, California. The HP 7978A is a product of HP's Greeley Division in Colorado. Early in the program all of the key players were identified with each department represented. Counterparts in each division were encouraged to interact directly with one another. One manager in each division formed the primary contact and was charged with full responsibility for coordinating the efforts of that division. Face-to-face meetings were held with all participants from each division and each department.

The frequency of meetings increased as deadlines approached. This encouraged personal responsibility for assigned tasks. Since the HP 3000 division had many programs proceeding in parallel, this was an important way to focus everyone's efforts on the HP 7978A periodically.

With two divisions committing large resources to the program, it was imperative that all the required support be available when it was needed. From the HP 7978A team's point of view, we needed to provide everything required to make the other division's effort efficient and successful. We used PERT charts to schedule all development and testing. This gave us a clear idea of when prototype units would be required and in what quantity. Since prototype HP 7978A drives were at a premium, it fell to the coordinating manager to understand the HP 3000 division's tests well enough to provide efficient scheduling. One of the most difficult tasks proved to be the tracking of prototype hardware upgrades. These upgrades usually consisted of ROM changes, but often included hardware changes. It was critical that these changes be coordinated and tracked so that errors reported in the other division's testing could be appropriately evaluated. Eventually we found that all prototypes had to be upgraded at the same time so that there was never a question of product "generation." A particularly successful aspect of anticipating needs grew out of our own need for consistent and complete documentation. Since the HP 7978A project was large, it was imperative that even early documentation be controlled. Valuable time can be lost when one engineer proceeds using dated or erroneous information. Since our documentation was standardized from the start, it made it much easier to provide everyone with useful information before the information was needed.

Using in-house expertise to augment host division efforts was very useful. Our division already enjoyed a proven reputation for design and manufacture of reliable, high-quality hardware products. Our coworkers from the HP 3000 division were software experts drawn from the HP 3000 R&D labs. Thus our experience made us the more logical group to work with the other division's environmental test and reliability engineering departments on hardware testing. To this end, we performed most of the required system hardware testing, leaving the HP 3000 R&D team to concentrate on driver development and software performance testing.

Richard T. Turley
Project Manager
Greeley Division

achieved. Thus, the entire organization was always aware of the goals for each development phase.

The goal of the business team was to manage the HP 7978A product development as a program. Other departments could assign members well in advance of critical interaction periods, permitting them to contribute to the conceptual definition of the product. Their continuing role in the program ensured that the product was optimized in all areas of concern. Each member of the business team was responsible for publishing an interaction and support plan that stated the exact relationship of that member's department to the program.

The R&D team was initially an architectural definition team with key members from all relevant engineering disciplines. After the architecture was solidified, subteams

began detailed planning for the product. An early realization was that the use of LSI technology would be an extremely important part of the HP 7978A. A separate project group was formed to ensure that appropriate processes and tools were available (see box, page 13). Thus, from its very inception, the development of the HP 7978A was a multi-group project.

The early knowledge of this program's complexity and potential staffing growth led the management team to evaluate the structure of the program and their own roles in it on a continuous basis. The management of complex and dynamic projects can be viewed as similar to organizational growth as detailed in reference 3. This model, shown in Fig. 3, implies that a number of crisis stages of management will be reached as an organization grows. This pattern

was true for the HP 7978A program. As the size of the organization and the complexity of the interactions increased, several management crises arose. In the early stages the problem was the division of responsibility and the span of control, so that the project managers felt that each of them was a full partner in the program. Later, it was realized that management specialists were required in certain areas such as software or system integration. The later stages of the program required that a program manager make decisions across department and division boundaries.

Many management lessons were learned as the program evolved. While none was unique or new, they still were learned by the traditional methods of finding a problem and working it out. The management team's approach to ensuring success was an outstanding experience, a classic example of management by objectives. The objectives for the product were agreed upon by the division's functional staff and the implementation was delegated to the responsible managers. The program was placed under a program manager with each functional area providing resources to ensure that the program was successful. Reviews were held for major checkpoints or when serious issues surfaced.

The critical management tools for this program turned out to be technical competence and flexibility, egoless and noncompetitive management teams, willingness to experiment with the organization, immediate response to critical issues, and ownership of all problems. This tool set allowed

a broad and changing group of managers to contribute to the program effectively.

Acknowledgments

The HP 7978A program had a rather large project team and it would be very difficult to single out each of the many individuals who contributed during the long development cycle in the wide range of technical and management issues. To those not mentioned here, thank you for your contribution to the success of the HP 7978A.

The authors would like specifically to recognize the project managers on the HP 7978A, who were able to be significant technical contributors while accepting the responsibility for managing all aspects of the project. These were Alan Richards, Gordon Thayer, Clair Nelson, John Meredith, Tom Bendon, and Don Stavely.

Jim Jonez provided invaluable market information and was an integral contributor in making major product decisions.

Special thanks must also go to Curt Brown, our section manager, and Rex James, our R&D lab manager, for their inspiration and program focus.

References

1. ANSI X3.54-1976 GCR.
2. ANSI X3.39-1973 PE.
3. L.E. Greiner, "Evolution and Revolution as Organizations Grow," *Harvard Business Review*, July-August 1972.

Write and Read Recovery Systems for a Half-Inch Tape Drive

by Wayne T. Gregory

THE ANALOG ELECTRONICS SYSTEMS connected to the head assembly in the HP 7978A Magnetic Tape Subsystem are responsible for writing and reading data on the tape in an error-free manner. These basic functions are separated into the write electronics, data protection, and read electronics.

Analog Write Electronics

The purpose of the analog write system is threefold. First and foremost, the circuitry must ensure that data is never accidentally altered or destroyed. Second, the circuit is responsible for recording digital data in the form of magnetic flux transitions on half-inch-wide magnetic tape. This second task must be accomplished such that the recorded amplitude and flux transition spacing are within the limits specified in the ANSI specification for data interchange using half-inch magnetic tape. Last, this system must be able to erase all data formats that have previously been

specified for half-inch tape by ANSI. These formats are 800 NRZI,¹ 1600 PE,² and 6250 GCR.³

Fig. 1 shows a block diagram of the analog write electronics. Control signals from the master controller set up the current drivers for operation in either PE (phase encoding) or GCR (group coded recording) mode and control operation of the erase current driver. The other control signals to the write system come from the write enable sensor on the tape drive hub and from the system power supply. These signals are used by the write enable electronics. The write system returns to the master controller the write enable status and a signal that indicates that the magnetic recording head is plugged in.

The circuit shown in Fig. 2 translates the digital data signals from the write formatter to currents flowing in the recording head. Since the magnetic flux generated by the recording head is approximately proportional to these currents, the digital data becomes regions of alternating flux

polarity stored on the magnetic media.

The circuit is a cascode differential current driver which is used to switch current from either of the two windings of the center-tapped recording head. Connected in the emitter circuit are two magnitude-controlling current sources. One current source is set for a constant value and the other is pulsed each time a data transition occurs. The common-base transistors are used as an externally controlled write-protect circuit.

The HP 7978A uses a dc erase system which, when energized, sets up a constant magnetic flux across the half-inch width of the recording tape. Whenever data is recorded, the erase head is energized and the magnetic tape is passed across it before reaching the data recording head. The erase circuit is a simple switched current source with write protection features.

Write compensation is used during the recording of data to improve the quality of the recorded flux transitions. This compensation improves two undesirable attributes of the recorded data: bit shift caused by a low-frequency-response head/media system and baseline shift caused by a system using a dc erase with a nonsaturating recording format.

A simple way to understand the main cause of bit shift is to think of the output of the recording system as a summation of the responses caused by each flux transition. Each transition produces an output pulse from the recovery head with a pulse width that is a significant portion of the nominal bit spacing (see Fig. 3). Furthermore, because of the shape of the recording head flux distribution and the media motion, the pulse is not symmetrical. As can be seen in Fig. 3, when two adjacent pulses are summed, the locations of their peaks are pushed apart. Since the digital information is stored in the location of each peak (i.e., each digital transition), compensation that decreases bit shift yields a low recovery error rate and simplifies the recovery problem. The HP 7978A's write compensation scheme reduces bit shift by effectively narrowing the isolated pulse response so that the summation of two adjacent pulses does not push their peaks apart. This is accomplished by including a small opposing current transition following each data transition. This accelerates the decay of the waveforms shown in Fig. 3, effectively narrowing the pulses. This technique, which is commonly used in half-inch magnetic tape recording, has been described in detail

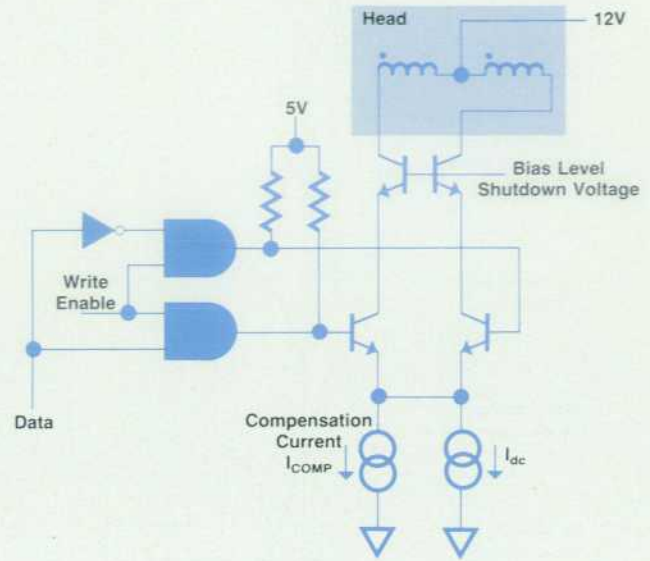


Fig. 2. Simple schematic of current driver for magnetic recording head.

in articles and texts on magnetic recording.^{4,5} The same circuitry also aids data recovery by improving the baseline shift caused by the dc erase current. This improvement is made by driving the write driver with more current when it is recording a flux transition. This current is of the opposite polarity to the current generating the erase state flux.

Data Protection

The recording and erase circuitry for a mass storage device must provide protection against accidental destruction of data. This is true for any storage device and even more true if the device is a backup device such as the HP 7978A. The ANSI standards for half-inch data formats provide for a plastic ring which fits into the back of each tape reel. If this ring is removed, the tape is write protected. When a tape reel is loaded on the HP 7978A, an optical sensor measures whether this write enable ring is installed. This information is reported to the master controller and on a direct line to the analog write circuitry. If the write enable ring is not installed on the tape reel, then the record and erase circuits are disabled by hardware. The analog write

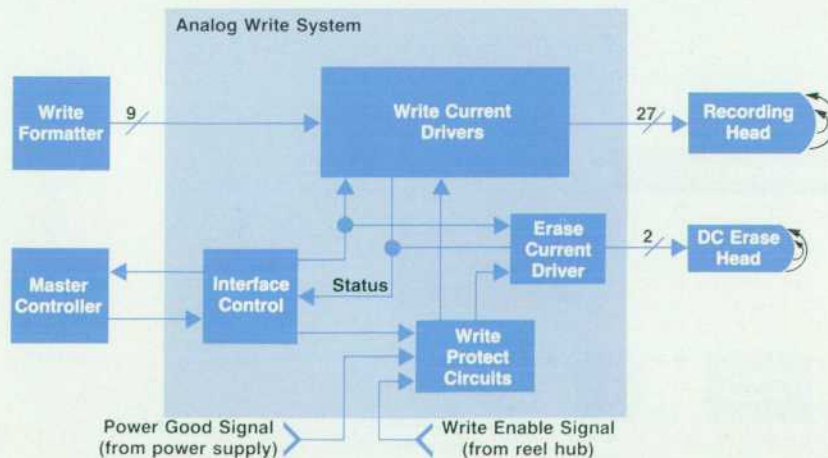


Fig. 1. Block diagram of the analog write system for the HP 7978A Magnetic Tape Subsystem.




Fig. 3. Bit shift can occur when two adjacent signal pulses of opposite polarity are recorded. The locations of their peaks tend to be pushed apart on the tape.

circuit also reports to the master controller the state of this write enable signal. Therefore, the write-protect condition is signaled to the master controller by two different pieces of hardware and the write/erase functions are disabled by circuitry in the write/erase system.

Another condition that must be prevented is an accidental destruction of recorded data during a power-fail condition. The HP 7978A provides this protection by monitoring the output of a signal generated from the primary switching voltage on the system power supply board. When this signal falls below a specified level, the write/erase system is held in a special power-fail write-protect mode.

The final data protection scheme is a self-check used by the master controller to verify that the erase and write enable control signals are functioning properly. This allows the master controller to detect faulty write or erase hardware.

Working together, these write protect systems ensure that valuable backup data is never lost through a system malfunction or during system power failure.

Analog Read Electronics

The purpose of the analog read recovery system is to

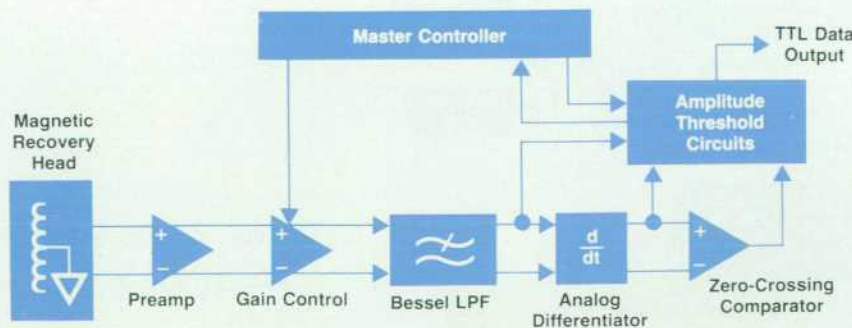


Fig. 4. Block diagram of analog read recovery system.

convert the flux transitions recorded on the tape to digital transitions. This conversion must be accomplished without altering the transition spacing by more than 1% relative to the nominal transition spacing of the data rate.

As shown in Fig. 4, the read recovery system is a standard recovery scheme which uses a preamplifier, low-pass filter, active differentiator, and zero-crossing detector to convert the magnetic flux transitions to logic-level signals. These circuits, working together with an amplitude threshold detector, provide a recovery system that is capable of detecting dropouts and data gaps, automatic adjustment of system gain level, and recovery with a 33-dB signal-to-noise ratio, which adds less than 1% phase error.

The system preamp, which for noise reasons is mounted close to the magnetic recovery head, provides a gain of 40 dB with a 3-dB bandwidth of about 10 MHz. After it is amplified by the preamp, the data signal is amplified further in an adjustable gain control stage which is capable of providing a gain from 0 to 9 V/V. The gain control stage drives a passive fourth-order, linear-phase low-pass filter which has a 1-dB bandwidth of about 400 kHz. After this point, the signal is processed by the differentiator stage. The differentiator converts the amplitude peaks to zero crossings by passing the signal through a bandpass filter function that has a peak amplitude response around 500 kHz ($1.5 \times$ the maximum data frequency) and then rolls off at a rate of -60 dB/decade. The signal is then input to the zero-crossing voltage comparator and converted to a TTL-level signal. The TTL signal is processed by the amplitude threshold circuits to determine if a logic transition is valid data and then sent to the clock recovery phase-locked loop to be processed further.

References

1. Phase encoding, ANSI X3.39-1973.
2. Group coded recording, ANSI X3.54-1976.
3. Nonreturn to zero, inverted, ANSI X3.22-1973.
4. F. Jorgensen, *Magnetic Recording*, TAB Books, 1980, pp. 262 and 270.
5. L.E. Ambrico, "Pulse Crowding Compensation for Magnetic Recording," *Computer Design*, Vol. 11, no. 6, June 1972, pp. 55-60.

Digital Formatting and Control Electronics for Half-Inch Tape Data Storage

by Jimmy L. Shafer

READ RECOVERY and read/write formatting of the nine-track data in the HP 7978A Magnetic Tape Subsystem are fairly complex tasks. Made up of four major assemblies, the digital formatting and control electronics handles read formatting, write formatting, data deskewing, clock recovery, and on-board diagnostics. The system is designed to support both GCR (group coded recording) and PE (phase encoding) formats in a transparent manner. To minimize board space and reduce cost, the functions are heavily integrated.

Read Formatter

The read formatter in the HP 7978A must read and decode data, correct errors, and detect problems in the retrieved and corrected data. The read formatter design (Fig. 1) centers around two LSI parts. The Read IC is a custom LSI chip that contains the data path portion of the design. A semicustom LSI chip is used to implement the read formatter controller for the data path. Together, the two integrated circuits perform the task of formatting the read data, which includes error correction and detection.

The Read IC is largely composed of a register-to-register data path that is controlled by the read formatter controller. The Read IC makes use of synchronous logic design and is required to run at a slower clock rate than would otherwise be possible with TTL design. The result is that the Read IC contains pipelining to allow parallel calculations while the data propagates along the path. In this way, an MOS design structure is maintained. The cost of pipelining in LSI is relatively cheap, thereby justifying the trade-offs.

GCR Format Architecture. The overall block diagram for the Read IC appears in Fig. 2. This block diagram includes the circuitry used for all three data formats. Some portions of this circuitry are used less extensively for particular formats. The data path consists of a 5-to-4-bit converter, mark detection circuitry, error detection circuitry, error correction circuitry, track-in-error accumulator, and error correction verification. The chip architecture is based

primarily on the requirements of the GCR data format. ANSI Standard X3.54-1976 gives the details of this format,¹ but the following brief overview should be enough to understand the read formatter.

The GCR format is designed to add sufficient data redundancy so that errors in up to two tracks of recorded data can be corrected. To make this possible, the data is divided into seven-character blocks. An eighth character must be added to complete a data group. This character is supplied by the error correction character (ECC.)¹ The ECC is formed by using the preceding seven data characters as input to a polynomial generator. Each data group is divided into two four-character subgroups (see Table I).

Each subgroup of four characters (including the parity bits) is now expanded to a storage subgroup of five charac-

Table I
Data Group to Storage Group Conversion

Physical Tape Tracks	Data Group		Storage Group	
	Data Subgroup "A"	Data Subgroup "B"	Storage Subgroup "A"	Storage Subgroup "B"
1	DDDD	DDDE	XXXXX	XXXXX
2	DDDD	DDDE	XXXXX	XXXXX
3	DDDD	DDDE	XXXXX	XXXXX
4	PPPP	PPPP	XXXXX	XXXXX
5	DDDD	DDDE	XXXXX	XXXXX
6	DDDD	DDDE	XXXXX	XXXXX
7	DDDD	DDDE	XXXXX	XXXXX
8	DDDD	DDDE	XXXXX	XXXXX
9	DDDD	DDDE	XXXXX	XXXXX
Group Positions:	1 2 3 4	5 6 7 8	1 2 3 4 5	6 7 8 9 10

NOTE: Data is recorded in 9-bit characters (across the tape) in blocks 10-bits long. This 90-bit group is called a "storage group." Before conversion to the recorded code values, there are eight linear bits made up of seven data bits and one check bit. This group of 72 bits is called a "data group."

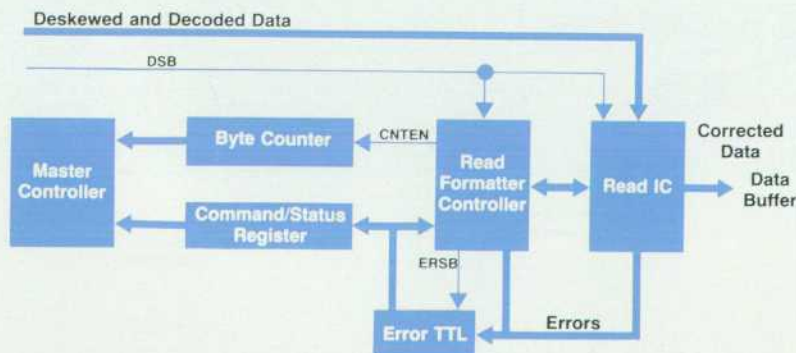


Fig. 1. Block diagram of read formatter subsystem for HP 7978A Magnetic Tape Subsystem.

ters. This is done by taking four bits, one from the same bit position of each character, and translating them to a 5-bit code as defined by the ANSI standard. (See Table II for conversion definition.) When reading, these 5-bit storage subgroups must be decoded back into the 4-bit data subgroups. This is handled by the 5-to-4-bit converter within the Read IC.

Data reliability has already been enhanced three ways. First, the extra bits added to the data groups increase the Hamming distance* between data bytes, making it less likely to read one byte as another. These bits are actually added to ensure that no more than two zeros occur in a row, making self-clocking possible with NRZI (nonreturn

*The number of symbol places in which two words of the same length differ from one another. Also called the Hamming metric.

Table II
Record Code Values

Data Values (Group Positions: 1 to 4/5 to 8)	Record Storage Values (Group Positions: 1 to 5/6 to 10)
0000	11001
0001	11011
0010	10010
0011	10011
0100	11101
0101	10101
0110	10110
0111	10111
1000	11010
1001	01001
1010	01010
1011	01011
1100	11110
1101	01101
1110	01110
1111	01111

to zero, inverted) encoding. Second, a parity bit is included to identify characters with single-bit errors. Finally, the ECC contains redundant information about the data in the group, making it possible to determine which bits, if any, are in error. The read electronics also provides track-in-error signals for any track with low output or excessive phase delay, thus increasing the correction capability.

As an extra margin of safety, the GCR format is designed to ensure that any data bit changed as the result of an error correction is one that really should have been changed. The format uses two overall cyclic redundancy checks (CRC and Auxiliary CRC) as verifications for data integrity. The checks are meant to indicate the occurrence of errors, but they have no capability to correct them. Two CRCs are used so that one will identify any errors that the other is incapable of detecting. CRCs (like the ECC) are generated by modulo 2 polynomial arithmetic. This arithmetic is performed by the read formatter and the results are compared to the CRC and ACRC supplied with the data, thereby verifying both good and corrected data.

The CRC is used in its own group, known as the CRC group. The CRC group also contains a character called the residual character. This character is two counts (modulo 32 and modulo 7) of the number of characters in the full data block. The modulo 7 count indicates how many valid data characters will appear in another group called the residual group, and the modulo 32 count is used to determine whether an odd or even number of characters exist in the block.

The residual group (see Fig. 3) contains other information of interest. It contains any extra data characters when there are not seven characters left to make a group. Thus the residual group will contain at most six valid data characters for the entire record. Pad characters are added as required to ensure that six characters in the residual group are filled. The seventh character is the ACRC. This group is treated as any other. An ECC is generated and the 8-character group

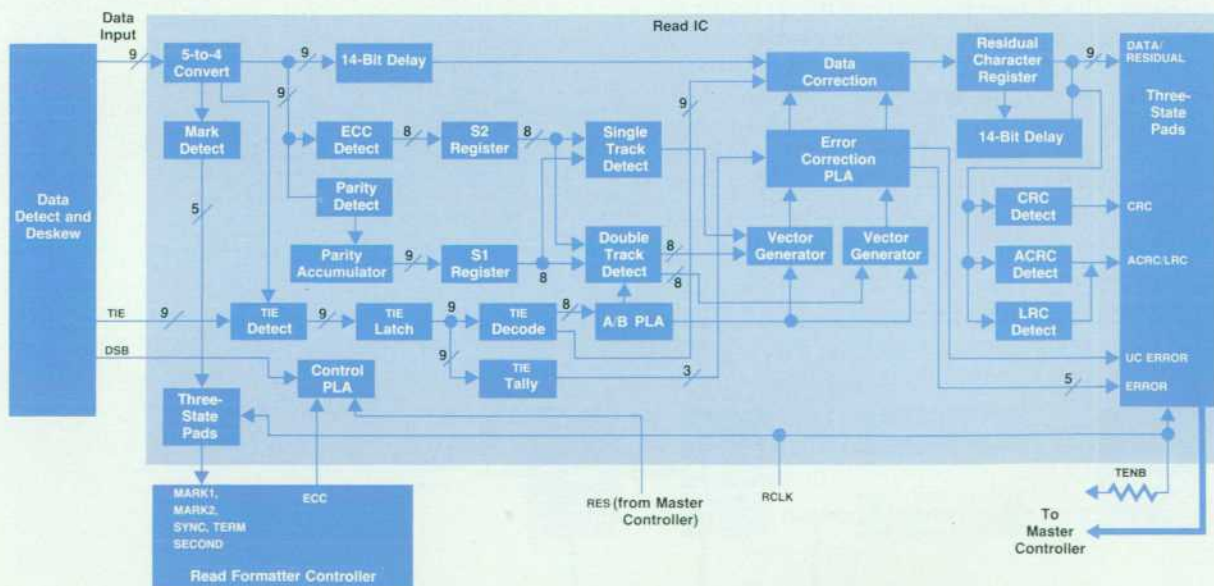


Fig. 2. Block diagram of custom LSI Read IC used in the read formatter.

is translated to a 10-character storage group.

The CRC and ACRC are arranged in the architecture so that the CRC generator can include the ACRC as an input to its polynomial. The CRC also makes use of whatever pad characters are inserted into the data blocks while the ACRC does not. However, both the ACRC and the CRC are grouped into 7-byte data blocks. Each of these blocks has an ECC generated and appended at the end. The blocks are then translated into two 5-character storage subgroups as illustrated by Tables I and II.

The blocks described so far have explained the generation of data groups, the CRC group, and the residual group. All that remains in any data record are the control marks. Control marks include the preamble, the postamble, the resync fields, and the end mark. Referring to Fig. 3, the preamble is used to announce the arrival of a data block under the read head. It identifies the beginning of the record with the TERM and SECOND marks. It provides 14 sync subgroups to allow the read electronics to sync with the data. A MARK1 indicates the arrival of data. Next, a set of 156 data groups is written on the tape. This set is followed by a resync burst to allow the read electronics to relocate its position. A resync burst consists of a MARK2, two SYNCs and a MARK1. This sequence of 156 data groups followed by a resync burst continues until all data (except for any residual amount of data less than seven characters) is written to the tape. Now an end mark is written. Following the residual group and the CRC group, a postamble is written. The postamble is a mirror copy of the preamble except that the last character ensures that the tape will be left in the unfluxed state. The postamble signals that the end of the block has occurred and can be used as a preamble if the tape is read backwards.

Additional control marks can be written outside of a data

block. These marks serve to identify the start of the tape and indicate the recording format being used. In the GCR format, there is an automatic read amplification burst to help calibrate the read electronics to the signal levels on the tape. Finally, a tape mark is available for the host computer to use to indicate file bounds. Interblock gaps appear between data blocks or tape marks.

This discussion leads to the use of the mark detect function indicated on the Read IC block diagram (Fig. 2). This block contains nine decoder PLAs (programmable logic arrays). The Read IC decodes the control marks with each clock and the read formatting controller decides which ones are valid. These marks become internally generated stimuli that are used in the read formatting controller for decision making and timing of control functions.

PE Format Architecture. The Read IC can read data in the phase encoding format at 1600 characters per inch.² Referring to the data path block diagram (Fig. 2), the PE data path uses the GCR data path with certain areas used less extensively. Transition information from the read head is encoded to include a phase and data window for each byte, according to ANSI Standard X3.39-1973. The PE format has no imbedded control marks or any redundancy codes. The data block consists of a preamble, the data, and a postamble. The only additional marks are a format identification at the beginning of the tape and tape marks included at the host computer's option.

PE guarantees that a transition occurs in every data window. Each phase window (level restoring) may have a clock transition, depending upon the data to be written and the previous state of the tape.

Read Interfacing. As shown in Fig. 1, the interface for the read formatter is made up of three functions. The first is the interface to the master controller. The read formatter




Fig. 3. GCR tape format. Legend: D=data character, H=pad or data character, X=residual character, E=ECC, C=CRC, N=ACRC, L=last character, and B=CRC or pad character.

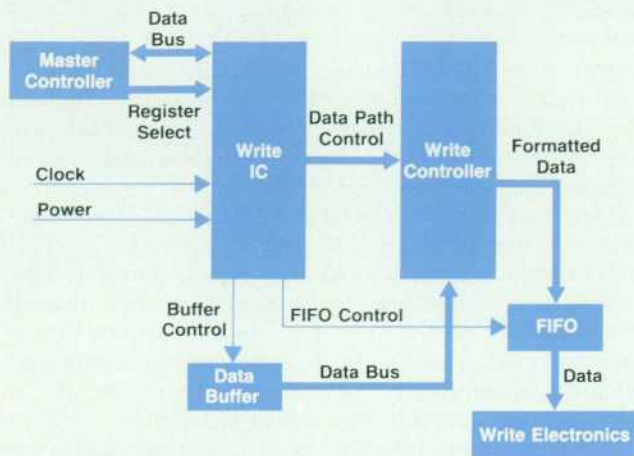


Fig. 4. Block diagram of write formatter subsystem.

accepts information from the master controller that indicates the density of recording and the mode of operation. A status word indicating the success of the read operation is accessible via this interface to the master controller. Finally, the master controller can access the byte count to determine the number of bytes passed through the read formatter.

The data buffer interface allows the read formatter to pass the corrected and formatted bytes to the data buffer. A full handshake occurs between the two functions to ensure data integrity. The data is in the form of an 8-bit byte with a parity bit.

The last function is the data detect and deskew interface. This function deskews the data before sending it to the read formatter. In addition to the nine data tracks, nine track-in-error pointers are generated by the deskew function and provided to the read formatter through this interface.

Write Formatter

The format recorded on the tape by the HP 7978A is selected by the host. For GCR, the write formatter must encode the data into GCR storage groups, which include the error correction character ECC. Also, control groups must be inserted into the data stream properly to meet the specifications according to ANSI. PE data also requires control groups, but not the error correction information. To allow read verification, the cyclic redundancy checks must be calculated and inserted into the data stream. In addition to these, other bursts required by the supported formats must be generated by the write formatter. These

include tape marks, density IDs, and automatic read amplification bursts.

Two LSI parts make up the heart of the write formatter design (Fig. 4). The Write IC is a custom CMOS LSI chip, which makes up the data path of the write formatter. The other LSI chip, a semicustom design, forms the write controller. Some TTL circuitry, including a FIFO (first-in, first-out) buffer, makes up the remainder of the system.

Using the register-to-register pipelining scheme previously mentioned for the Read IC, the Write IC forms the data path portion of the design. Like the read formatter, it is synchronous and controlled externally by its controller, the write controller. The Write IC is implemented using HP's NMOS-CP process and replaces approximately 90 SSI/MSI/LSI parts, thereby reducing cost and enhancing reliability. As shown in the block diagram (Fig. 5), the circuit consists of a parity checker, a 7-byte-deep shift register, polynomial generators for CRC, ACRC, and ECC characters, a residual character generator, a 4-to-5-bit encoder, a control mark generator, a phase/flux encoder, and multiplexers for controlling the flow of data across the chip.

Write Interfacing. Referring to Fig. 4, the interface for the write formatter is made up of three functions. The master controller indicates to the write formatter the mode of operation and the density of recording across a command interface. Also, the master controller issues to the write formatter the type of block that needs to be formatted. These include data, tape mark, or some other control burst. The status of the write operation can be read by the master controller across the status interface.

Next is the data buffer interface. Eight-bit bytes of data are sent across this interface with a parity bit. An end-of-data (EOD) bit accompanies the last byte. A full handshake takes place between these two subsystems to guarantee data integrity.

Finally, the write formatter interfaces to the write electronics. This is accomplished using a FIFO buffer. The purpose of the FIFO buffer is to create a synchronization boundary between two asynchronous systems. Data leaves the write formatter system in nine-track form.

Data Detect and Deskew

Before the read formatter can attempt to read the data, the data bits must be realigned across the nine tracks. Misalignment of the bits is caused by mechanical variances (such as head skew) and electrical propagation delays. The job of realigning the nine-bit bytes is done by the data detect and deskew (DDD) circuitry. In addition to realigning the data in the nine tracks, this circuitry must also identify




Fig. 5. Block diagram of custom LSI Write IC used in the write formatter.

the type of burst under the head and report this information to the master controller. This allows the master controller to manage the read/write operations appropriately.

In Fig. 6, the block detection portion of the architecture can be seen. Track activity recognition is fed to the master block detect to allow a tally to be done for block detection purposes. Block type information is relayed to the master controller. In addition, gap information is sent to the read formatter and to the servo controller.

The job of recognizing track activity is performed within each of the slave deskew subassemblies. The slave deskew subassemblies are implemented using semicustom LSI chips called USMs (universal synchronous machines; see box, page 13). There are two tracks of slave deskew in each 40-pin USM, requiring that five of these chips be used.

The nine slave deskew subassemblies send their track activity information to a single USM called the master block detect. Here, a tally of this information is used to determine the type of block currently under the head. These include data blocks, automatic read amplification bursts (in GCR), tape marks, and density IDs. Upon encountering a new block type under the head, the DDD circuitry interrupts the master controller and makes this information available to it. In addition, the master block detect USM acts as an interface to the master controller. A bidirectional read/write port allows handling of commands and status. Information from the master controller is relayed to the remainder of the DDD circuitry by the master block-detect USM.

In addition to detecting track activity, the slave deskew USMs must also detect control marks for deskewing the data. This requires that each track identify sync marks within the data stream and report this information to the master deskew USM. Each slave deskew USM has control over its individual FIFO buffer. These 64×4 -bit buffers allow deskewing across the nine data tracks and create a synchronization boundary between the clock recovery sub-system and the DDD circuitry.

To ensure that the maximum allowable skew of 25 bits is not exceeded, the master deskew USM checks the skew and can force the release of the FIFO buffers as necessary. A majority vote is taken by the master deskew USM to determine when enough of the tracks have been deskewed to recover the data. In all cases, the master deskew USM

will wait as long as possible to allow all nine tracks to be deskewed.

Error pointers are generated within the deskew portion of the DDD circuitry. These track-in-error pointers accompanying each track are an accumulation of the lock-detect information provided by the clock recovery subassembly and the internally generated error information of the DDD circuitry. Tracks that cannot be deskewed and "dead tracks" fall in the second category.

As shown in Fig. 6, the DDD subassembly interfaces to the master controller, the clock recovery circuit, and the read formatter. Data into this subassembly is from the clock recovery circuitry in the form of nine tracks of data accompanied by individually recovered clocks. An indication of data integrity from the clock recovery circuitry in the form of a lock-detect signal also accompanies each data track.

Deskewed data leaves the subassembly and is sent to the read formatter with error pointers indicating bad tracks. The data is strobed to the read formatter in 9-bit bytes and a full handshake is done between the two systems.

The master controller has access to information about what type of block is currently under the head, as well as information about the success of the deskew operation. Blocks are verified by the data detect and deskew circuitry, both when they are written and when they are read.

Clock Recovery

The clock recovery board contains nine channels, each made up of a phase-locked loop and a discriminator (Fig. 7). The phase-locked loop reconstructs a data clock that is used to clock the data into the DDD circuitry. Detection of phase lock and loop control are performed by the discriminator.

The phase-locked loop for each channel is made up of the phase detector, filter, bias network, and VCO as shown in the block diagram. Each loop feeds back the data clock signal DCLK to the phase detector where it is compared with the incoming data and correction is applied to the VCO. The amount of correction is determined by the phase detector. The filter integrates the correction signal and filters out high-frequency phenomena such as bit shift, asymmetry, and noise. The resulting error voltage is filtered and applied to the VCO.

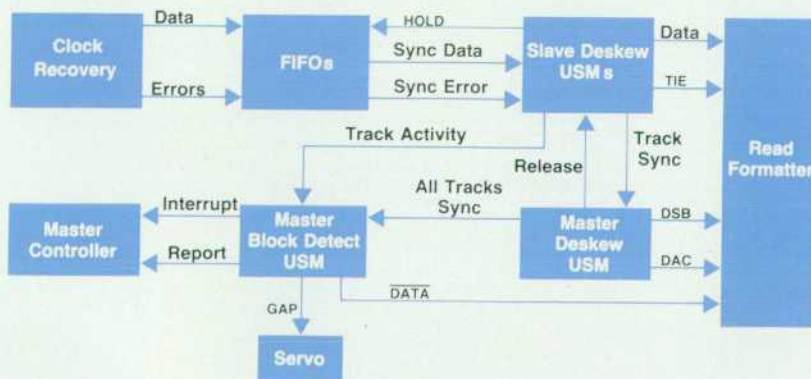


Fig. 6. Block diagram of the data detect and deskew circuitry.




Fig. 7. Typical channel of the clock recovery subsystem with common control circuits.

The discriminator is a state machine that creates two signals. The lock detect signal LD is sent to the data detect and deskew section of the formatter board. This signal indicates when the probability is high that the loop has acquired the data and is locked.

CLAMP is the second signal created by the discriminator. It is used to select between data entering the phase-locked loop or the write clock. The write clock is the data rate clock created on the write subsystem for writing data to tape. The ability to switch between the two is used to force the loop to maintain lock to a frequency that is very close to the actual data frequency. If CLAMP is true, forcing the loop to lock to the write clock, and data is encountered, the discriminator sets CLAMP low, allowing the loop to acquire real data. Whenever an extensive dropout occurs (128 windows without data) the discriminator sets CLAMP high, forcing acquisition to the write clock.

Control signals from the master controller determine the environment in which the phase-locked loop is expected to function. When the density of recording is changed from GCR to PE, the loop switches its divide-by-N circuitry by a factor of three to allow recovery at PE densities.

Diagnostics

Diagnostics have been built into the read electronics to allow self-test capability. Two modes of this self-test can be selected by the master controller. One mode loops data from the data buffer through the write formatter, data detect and deskew circuitry, and read formatter back to the data buffer. This mode checks the formatter subsystem. Errors can be injected by allowing the master controller to select various tracks to be "killed." In this way the error correction capability of the read formatter can be checked.

Second, the data can be looped in a similar fashion with the addition of the clock recovery subassembly. This checks the clock recovery circuit's ability to acquire the data and maintain lock. In addition, the inherent delays associated with the electronics force the data detect and deskew subsystem to perform synchronization and deskewing.

References

1. "American National Standard Recorded Magnetic Tape for Information Interchange (6250 CPI, Group-Coded Recording)," X3.54-1976, American National Standards Institute, June 10, 1976.
2. "American National Standard Recorded Magnetic Tape for Information Interchange (1600 CPI, Phase Encoded)," X3.39-1973, American National Standards Institute, 1973.

Streaming Tape Drive Hardware Design

by John W. Dong, David J. Van Maren, and Robert D. Emmerich

AS A STREAMING TAPE DRIVE, the HP 7978A Magnetic Tape Subsystem requires much simpler mechanical and electrical hardware than a start/stop drive with a capstan motor. The main items that are simplified are the tape path hardware and the servo electronics. Fig. 1 shows a block diagram of the HP 7978A. A start/stop capstan drive would require another motor as well as an extra tape buffering device. Capstan motors and tape buffering devices are fairly expensive. Eliminating these items reduces the factory cost substantially.

The HP 7978A uses two controllers to run the drive, a master controller and a servo controller. Based on a 68000 microprocessor, the master controller controls the data path electronics, the HP-IB (IEEE 488) interface, the front-panel display and keyboard, and the servo controller. The servo controller, which uses an 8051 microprocessor, controls the tape movement servo system and corresponding mechanical hardware.

Performing the same function on both drives, the data path electronics for a start/stop drive are very similar to the electronics for a streaming drive. The data path electronics are much more complex in a drive that uses the group coded recording (GCR) method of formatting data. In previous tape drives, this method of formatting data was implemented in TTL circuitry. The HP 7978A tape drive uses LSI components to reduce the data path electronics parts count from thousands to hundreds, resulting in a great increase in reliability and a large reduction in cost (see box on page 13).

Tape Path

The HP 7978A tape path has several major components: the main casting, the head plate assembly, the buffer assembly, the speed sensor, the supply hub and motor, and the take-up reel and motor. The latter five components are mounted on the main casting as shown in Fig. 2.

These components serve two major functions. The first is to guide the magnetic tape past the head in a predictable manner, compatible with the standard written tape formats used by the HP 7978A. These include the usual tracking and constant-speed constraints. Nine-track tape drives also have the additional constraint of writing the nine parallel bits of a single byte within a given distance along the tape. The distance is specified by the established tape format for compatibility. This is the skew requirement, which for the 1600 PE (phase encoding) format is 625 microinches over the width of the half-inch tape.

The HP 7978A tape path is designed to minimize skew caused by the variations in the positions of various components. These effects were modeled by treating the tape as a cantilever beam with four supports. These supports are the buffer and speed rollers and the two fixed guides on the head plate assembly. The model worked well in predicting the effects of component variances on skew, and was used in determining the final positions and allowable tolerances of the various tape path components.

The second major function of the tape path components is to move the tape gently and without damage from one reel to the other. This is facilitated by minimizing the

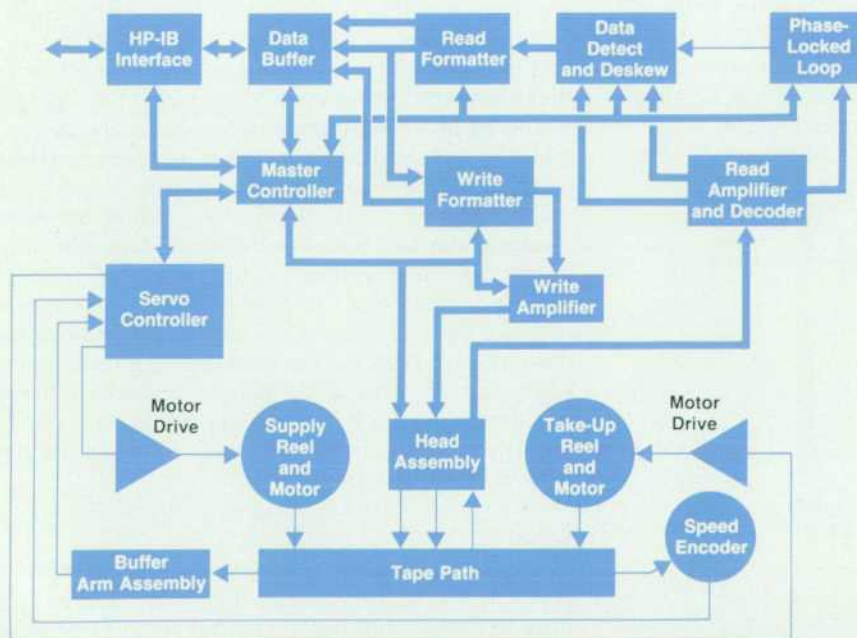


Fig. 1. Block diagram of HP 7978A Magnetic Tape Subsystem. This half-inch tape drive is based on a streaming tape design as opposed to the more common start/stop tape drive designs.

number of components touching the tape. Only seven items contact the tape between reels. Only three of them contact the oxide side of the tape, which is the side where the data is written.

The rollers are set at less than ten minutes of arc from perpendicular to minimize edge stresses on the tape. High edge stresses are one of the major causes of tape damage and subsequent loss of data. The spring washers on the fixed guides, which help guide the tape, exert a very low edge pressure on the tape, gently persuading the tape to move up against a known reference edge to minimize off-track errors. The tape cleaner has five ceramic blades, each flat on top. The resulting ten edges, two per blade, remove loose oxide and other debris as the tape passes over them. These blades exert a low contact pressure on the tape compared to other tape cleaner designs using sapphire or tungsten carbide blades, which exert a very high contact pressure on the tape.

These low edge and contact pressures are easy on the tape, but they are also easy on the components themselves. The components of the tape path are designed to achieve an 8000-hour tape running life. With the exception of the main casting, these components are modular. Since they are precisely set at the factory, none of the components require any adjustments when they are replaced by service personnel in the field. In comparison, most other half-inch tape drives require the use of specially written tapes and/or special height adjustment tools to replace some of their major tape path components.

One of the major components that requires no adjustment is the head plate assembly. The two fixed guides, the tape cleaner, and the magnetic head are mounted together on this assembly. To obtain the modular, no-adjustment feature, the magnetic write gaps must be aligned to one half minute of arc. This is done by mounting the head on an adjustable tongue integral to the head plate as shown in Fig. 3. A differential screw with a pitch of 288 threads per inch provides the fine mechanical adjustment necessary. The head is mounted on the tongue so that the write and read gaps are centered above the tongue's pivot point. This is done so that the off-track error caused by the skew adjustment is negligible. The full travel of the skew adjustment causes an off-track error of only 60 microinches.




Fig. 2. Main casting assembly and location of tape-path components for the HP 7978A.

The magnetic head is ceramic coated to meet the 8000-hour design life. The fixed guide washers are made of alumina ceramic, as is the tape cleaner. This very hard material resists the abrasive action of the tape very well and enables these parts to meet the design life.

The 8000-hour design life also applies to the buffer assembly. The tape is very abrasive on its edges where it is slit to size. Therefore, the roller flanges are chrome plated to reduce wear. In tape drives, one of the major causes of bearing failure is lubricant washout caused by the liberal use of tape-drive cleaning solutions. The rollers on the buffer assembly use shielded bearings and plastic caps to prevent cleaning solutions from washing out their lubrication.

The buffer assembly is an important part of the tape servo system. It maintains the tape tension and helps buffer the servo from any major tape movement to which the tape servo can't respond. Therefore, the servo requirements played a major role in the design of the buffer assembly. The amount of tape buffering was increased several times and the effective mass of the tension arm was reduced to enable the tape servo to function well.

The buffer assembly uses a differential optical sensor, which uses a slotted mask over two phototransistors. The optical sensor gives absolute position feedback to the tape servo telling the location of the buffer assembly arm.

The tape servo also influenced the speed sensor design. The speed roller diameter was optimized to minimize the effective inertia on the tape. This also helps reduce the chance of the tape's slipping on the roller during high tape accelerations. The roller is grooved to prevent tape slippage by allowing air to escape from between the tape and the roller. The sensor itself is an HP one-thousand-line optical encoder. The roller is made of hard-anodized aluminum and is designed to last 8000 hours.

During development of the HP 7978A, a very interesting tape vibration problem occurred. It was a high-pitched squeaking sound, which occurred while the tape was running. The sound could be turned off by pushing the lower fixed-guide spring washer away from the tape. The tape and the washer were acting as a spring mass system. If the spring-mass-system resonant frequency was close enough to the longitudinal resonant frequency of the tape, a squeak would result. The longitudinal resonant frequency of the tape is determined by the length of tape between the rollers on the buffer arm and the speed sensor.

Once this was understood, the problem was solved by changing the two resonant frequencies so that they were farther apart and wouldn't couple into each other. Changing the longitudinal resonant frequency by changing the distance between the rollers was undesirable in terms of design time, skew, and the need to minimize the solution's impact on the other components. The mass of the washer was therefore increased, reducing the spring mass resonant frequency substantially below the longitudinal resonant frequency of the tape span.

Servo

The HP 7978A's servo system is responsible for all tape handling and motion control across the read/write heads. A major design criterion was to provide this motion control at a constant tape velocity with minimum speed or tension

perturbation error. Since the servo system is a major electrical-to-mechanical interface, another focal point in the development effort was to provide high reliability and reduced service costs. HP's earlier half-inch tape drives (the HP 7970 and HP 7676) had start/stop servos. This means that they stopped between accessing records on the tape. All the accelerating and decelerating had to take place within the gaps between records on the tape. Since gaps can get as small as 0.28 inch (0.71 cm), the tape had to be accelerated very quickly. The reel motors could not provide enough acceleration, so a capstan motor was used. It had to be isolated from the large inertias of the reels to accelerate so quickly. This isolation was provided by large buffer arms or vacuum columns. They allowed the tape over the head to be accelerated quickly by the capstan motor, while the reels followed a slower acceleration profile.

The HP 7978A, on the other hand, is a streaming tape drive. That is, it is capable only of much lower accelerations and therefore tries to avoid stopping. The reduced acceleration requirements allow the servo to be greatly simplified. In the HP 7978A, the reel motors provide the necessary acceleration, so an inexpensive position encoder has virtually replaced the capstan motor and its drive electronics. One of the buffer arms has been eliminated also. This simplification contributes greatly to the reliability and reduced costs of the HP 7978A. Since starting and stopping are minimized, the acceleration duty cycle is greatly reduced. This reduces stress and strain on the tape, mechanics, and electronics, which contributes to improved reliability and reduced service costs.

Servo Design

The tape handling system is made up of two servo systems: the velocity servo and the tension servo. Regardless of the specific tape motion operation, the velocity servo maintains the required steady-state tape speed to within a $\pm 1.5\%$ tolerance and the tension servo simultaneously maintains an absolute tape tension of 10 oz ± 2.6 oz. Servo system isolation is stressed in the design by having the velocity servo control only the take-up reel, and the tension servo control only the supply reel. Hence, the two servos are physically connected only by the tape.

A good way to understand how these two servo systems work together is to view the velocity servo as the master servo and the tension servo as its slave. For example, the master velocity servo is fed a velocity profile input to follow, which will accelerate the tape up to a specific speed and then maintain this speed thereafter. While the master

servo is following its new velocity profile input, the tension servo, which knows nothing of tape speed profiles, sees a perturbation error in its own control system—a tension error. The slave tension servo then attempts to reduce this tension error, and in so doing, it effectively forces the tape speed of the supply reel, which it controls, to match exactly the tape speed of the take-up reel, which the master servo controls.

The velocity servo is a Type I control system which uses integral-plus-proportional control action. It is primarily a digitally controlled servo system implemented using a digital position encoder, an 8051 microprocessor, a digital-to-analog converter (DAC), a motor drive amplifier, and the take-up reel motor (see Fig. 4). The position encoder increments or decrements a counter in response to tape motion. The counter is read by the processor at precise, fixed time intervals to determine the speed. The processor performs the necessary control calculations and outputs a voltage to the motor drive board via the DAC. This voltage is amplified and applied to the appropriate speed-controlling motor, the take-up motor, which then moves the tape onto the take-up reel past the speed encoder, thus completing the speed control loop.

The tension servo is also a Type I control system which uses series lead compensation as well as feedforward control action. It is primarily an analog control system consisting of a buffer arm assembly, a motor drive amplifier, and the supply reel motor (see Fig. 4). The intelligence for this servo resides in the buffer assembly, which senses the actual buffer arm position using a differential optical sensor and provides the appropriate servo lead compensation. The output of the buffer assembly is summed with the D-to-A converter output of the velocity servo to achieve feedforward control action. This result is sent to the motor drive board where it is amplified and applied to the buffer-arm control motor, the supply reel motor. This motor moves the buffer arm via the tape on the supply reel, thus completing the tension control loop.

The tension servo is virtually transparent and yet entirely subservient to the velocity servo. It is transparent to the velocity servo in that it presents a relatively large mechanical impedance to the velocity servo system, which produces little perturbation in the velocity servo's frequency transfer function response. It is subservient to the velocity servo in that the controlling input to the tension servo is the velocity servo's controlled output, the speed of the tape.

The HP 7978A servo system was designed using localized feedback compensation techniques. This can be looked upon as a decentralized-intelligence system design. The velocity servo knows virtually nothing of the tension servo and sees it only as a relatively high and virtually nonexistent mechanical load impedance. The tension servo knows nothing of the velocity servo, but merely tries to maintain its buffer arm near its set location, regardless of the disturbances existing outside its own control loop. While the tension servo is maintaining this buffer arm position, it is the specially designed spring mechanism in the buffer arm that actually controls the tape tension. The reel motors are driven directly by voltage amplifiers, which make direct use of the motors' own individual back-EMF properties. This yields a useful velocity feedback compensating drive




Fig. 3. Magnetic recording head plate assembly. Skew is adjusted by turning the very-fine-pitch differential screw.

scheme.

The HP 7978A servo system maintains its tension and velocity within specification over a great range of different conditions. Tape reel sizes and pack radii can vary over a large range, but since neither detection of reel sizes nor the knowledge of tape pack radii on the reels is needed, the design's simplicity is maintained.

Starting and Stopping

The simplicity afforded by making the HP 7978A a streaming tape drive instead of a start/stop drive is not without price. Such simple actions as starting and stopping the tape lose their simplicity and become more of a technical challenge.

Stopping and then starting the tape within the distance of an interrecord gap is impossible in the HP 7978A because of its slower accelerations. It moves 4.3 inches of tape over the head just to bring the tape to a stop. Even though it must leave the gap to stop and start, it must be able to return to the gap at full speed to resume reading or writing where it left off. The HP 7978A performs this "repositioning" using the velocity encoder and its counter as an odometer. This same odometer setup is used when a record is retried. In that case, the HP 7978A needs to return the tape to the previous gap so that it can retry reading or writing.

One difficulty that emerges is caused by tape stretching. The tape tension at the buffer arm is fairly constant over its entire range of travel. But when the tape is moving, its tension at the velocity encoder changes because of the friction of the tape against the head assembly. When the tape is moving forward, the frictional drag *increases* the tension at the speed encoder. However, when the tape is moving in reverse, the frictional drag *decreases* the tension at the encoder. So the tape stretches more and becomes longer when traveling in reverse than when traveling forward. This introduces a cumulative error in the encoder-based odometer. Thus, when the odometer says that the correct gap is over the head, it is actually a small distance in front of it. This makes the preceding block move slightly closer to the head each time it is retried. To ensure that the HP 7978A is in the correct gap, the HP 7978A detects whether the preceding block is over the head when the odometer says that the gap should be. If it is, the odometer is updated to match the new position of the gap. This zeros out the accumulated errors from previous retries.

Gentle Tape Handling

The reliability of the data is another important concern in the HP 7978A design. The servo is designed to minimize the risk of damaging the data on the tape. This is helped by a very smooth loading process. When the tape tension is being established, the error signal in the speed control loop is forced to zero, so that when the loop is activated, the tape will not be jerked. The buffer-arm control loop is also activated at the point where it will jerk the tape the least. In addition, the loading process is somewhat tolerant of incorrectly threaded tapes. Small mistakes in loading the tape around the tape path will rarely result in any tape damage.

In the event something causes the servo to lose control (e.g., a hardware failure or loss of power), the servo is shut down using passive electrical components that sink energy from the motors in such a way as to protect the tape from bunching up or stretching. The user is also protected from injury in the event of a hardware failure. If the door should be opened while the motors are still turning (made possible only by a hardware failure), a safety shutdown mechanism implemented redundantly in both software and hardware forces the motors to stop.

Diagnostics and Confidence Tests

Key contributors to the reduced service costs of the HP 7978A are the diagnostics and confidence tests. They provide very good checking of the mechanism and electronics. The digital portion of the servo electronics is able to diagnose its own problems very well. This includes the 8051 processor and its system of chips, the state machines and I/O electronics, and the DAC and comparators.

Some of the servo's hardware, however, can only be tested with a user's help. Extensive confidence tests were therefore implemented. The BOT (beginning-of-tape) and EOT (end-of-tape) sensors, which are part of the servo hardware set, can be tested by selecting a confidence test that displays their status at the front panel. As a reflective surface is waved in front of them, the display shows whether or not they see the surface. Another test allows the speed encoder to be checked by displaying a count that

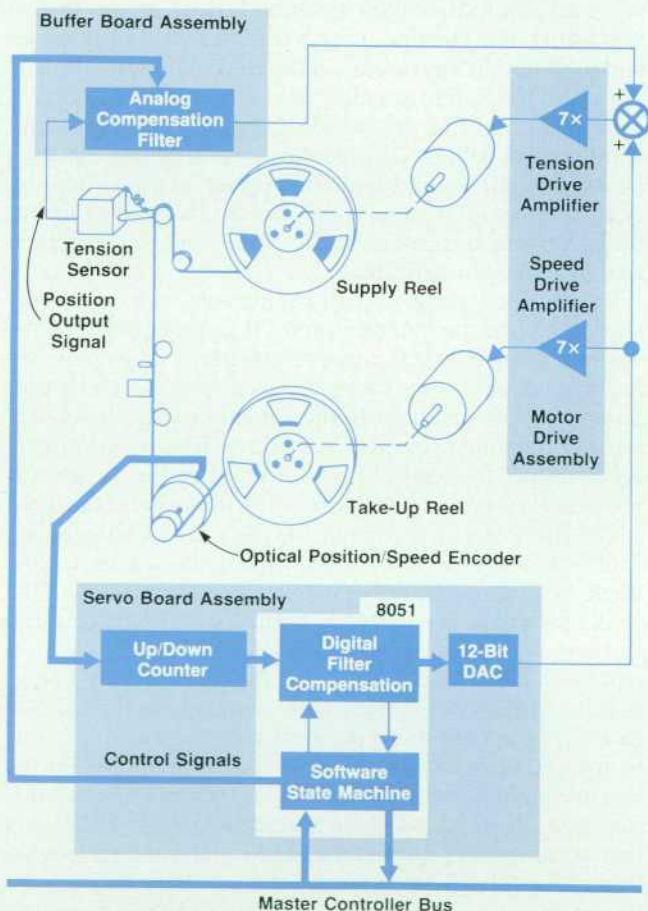


Fig. 4. Diagram of HP 7978A servo system.

increases as the encoder is rotated clockwise and decreases as it is rotated counterclockwise. The user can also test the buffer arm by moving it through its range. The front panel displays whether it is over or under tension, when it crosses through zero, and whether or not its optical sensor is operating properly. A failure is indicated by the numbers in the display not going through the proper sequence. The motor driver also has a confidence test. As the user steps through the different parts of the test, the motors are moved in various ways, exercising the basic capabilities of the motor drive amplifiers. With this test, a failure is indicated if the motors do not move properly.

Acknowledgments

Among those who played a major part in the hardware design of the HP 7978A were Kevin Wilson and Bob Gilbert, master controller design, Alfred Natrasevski, HP-IB interface design, Peter Way and Bob Emmerich, servo design, and Martha Chavez-Simmons, motor drive board design.

Gordon Thayer was responsible for a major part of the GCR VLSI design, along with Mike Tremblay, Bob Sobol, Jeff Osborne, and Rick Turley. Gene Briles designed a large part of the read circuitry. The phase-locked loop was designed and debugged by Jeff Kato and Tom Holmquist.

The tape buffer arm assembly and the speed sensor assembly were designed by Dave Sims and Dave Jarrett, respectively. Dave Jarrett also designed the hubs and the door. The industrial design was done by Jim Dow. Doug Domel and Dave Jones did the product design. Dave Jones also designed the aluminum base plate casting. Mel Crane was a great help in solving many of the mechanical problems that arose during the development of the HP 7978A.

Thanks to John Meredith, Tom Bendon, Alan Richards, Gordon Thayer, and Rick Turley for managing the later stages of the program in getting the product from the lab into production.

Many thanks also to Rex James for helping manage the program through to the end.

Firmware for a Streaming Tape Drive

by David W. Ruska, Virgil K. Russon, Bradfred W. Culp, Alan J. Richards, and John A. Ruf

THE MASTER CONTROLLER FIRMWARE of the HP 7978A Magnetic Tape Subsystem is responsible for the highest level of control. This firmware handles all interactions with the host system and the operator and oversees the execution of both the operational and the diagnostic functions of this half-inch tape drive.

The host system sends commands to the tape drive over the Hewlett-Packard Interface Bus (HP-IB, IEEE 488). The master controller firmware must receive and execute these commands and then respond with the appropriate results. In the execution of commands, the firmware will set up the tape drive's hardware subsystem and then detect either successful completion or an error condition. When an error occurs, the master controller firmware will process recovery procedures, retrying the operation if possible.

The operator interacts with the HP 7978A through the front panel. The master controller firmware receives the requests from the front-panel buttons, executes them, and updates the front-panel displays.

An important responsibility of the firmware is validating the fitness of the tape drive's hardware. If a failure is detected, the firmware will aid service personnel in the diagnosis of the problem.

Architecture

The architecture of the master controller firmware was designed using the principle of functional decomposition. The firmware contains three main subsystems: the channel program, the device program, and the diagnostic program. Each of these programs performs a specific function within

the tape drive. These programs use a command/report transactional communications scheme when interacting with each other. This architecture exhibits strong functional binding within the subsystems and a minimal amount of coupling between subsystems.

To facilitate maximum tape drive use, the master controller firmware supports queued operations. This is made possible by concurrent execution of the channel and device programs along with internal queues for commands and reports (Fig. 1). Through the use of read aheads when reading and "immediate responses" when writing, the tape can be kept streaming without the need of a complex queuing protocol with the host.

Channel Program

The channel program is responsible for processing all user interactions with the drive, that is, interface requests over the HP-IB from the host and front-panel requests from the operator. These requests generate command sequences which are sent to the device program for execution.

Streaming tape mechanisms require either a constant supply of data from the host or a constant acceptance of data by the host to maintain streaming. Otherwise, the mechanism must perform a tape reposition cycle whenever the data stream is interrupted.

There are two methods for smoothing data transfers with the host to help maintain streaming. Data can be buffered either within the host or within the tape drive. Buffering data in the host requires memory resources and the ability to queue requests to the I/O driver. Buffering data within

the tape drive frees the host from any special requirements to support a streaming drive. This buffering is also more advantageous, because it allows the drive to monitor the amount of used and unused buffer space. While writing, this capability lets the drive perform write holdoffs as necessary to extend streaming. When reading, the drive's data buffer allows the device to perform read aheads, anticipating the following host read requests.

Along with the data buffer, the channel program maintains command and report queues within the tape drive to monitor the placement of data within the buffer. These queues are also used to allow file marks and motion commands to be queued to the device by the channel program. Buffering and command queuing are transparent to the host system.

To allow data records to be queued in the buffer during writes, a new feature was added to the interface protocol. This feature is called "immediate response." Data sent from the host is stored in the drive's buffer and a command is placed in the command queue. An immediate report is sent to the host indicating that the data has been accepted. The host is then free to gather and send more data to the drive.

Data sent to the host system that was read off the tape is also stored in the drive's data buffer. Upon completing a requested read command from the host, the HP 7978A continues reading additional data records into the data buffer. This read-ahead operation continues until the data buffer becomes filled, the end of the tape is reached, or an error condition is encountered. In parallel with this read-ahead operation, the HP 7978A can continue to accept read commands from the host and fulfill them immediately with data records already stored in the buffer. Only when the host lags the tape drive to the point where the data buffer becomes full will the drive need to perform a reposition cycle.

The HP-IB interface protocol is a modification of the protocol used with the earlier HP 7976A tape drive. The HP 7976A allows command queuing of two read or two write operations. This allows an overlap of transfers between the host and the HP 7976A's buffer, and between the buffer and the tape. This queuing has to be performed explicitly by the host. To use this feature usually implies rewriting applications programs. The HP 7978A does not support queuing in its protocol. Instead, immediate response and read-ahead functions are used to make the streaming tape drive perform like a start/stop drive.

Device Program

The functionality of a streaming tape transport mecha-

nism can be maximized if the transport is allowed to run independent of outside timing constraints, and if it can implement intelligent algorithms for control under all situations. The device program provides this type of control.

All control of the tape mechanism and read/write electronics during normal operation is performed by the device program. As requests are received from the host or operator interface, they are placed into the queues by the channel program for the device program to act upon. The device program executes each of the commands, formats and returns a report to the channel program, and then goes on to the next command. The device program processes each command completely, including all exceptions.

A wide variety of recording errors and exceptions can occur during the processing of a host transaction. The ability to perform the proper recovery from errors automatically without burdening the host with unnecessary information and responsibility is built into the intelligent device program. Recording errors caused by such things as tape imperfections or debris in the tape path are recovered by the execution modules performing automatic retries.

Write retries involve repositioning the tape back before the record that was recorded with errors, erasing that portion of tape, and then skipping the next section of tape to attempt the write again. The data for the record is maintained intact in the data buffer until the record has been successfully written so that the host does not need to retransmit the information.

When a record is not read successfully on the first attempt, read retries are performed by repositioning the tape back before the record and reading it again with the read signal amplified. Retries for reads and writes occur until the transaction is successful or a maximum limit occurs, resulting in a hard error.

Different host and streaming tape data rates require that there be a decoupling of the two rates if high performance is to be maintained. In the HP 7978A, the decoupling is provided by queuing, data buffering, and a device program that runs independently. The channel program runs in sync with the host, handles the bursting needs of the interface, and queues them for the device program, which in turn runs in sync with the needs of the tape mechanism. The device program's only interaction with the channel program is at the beginning and end of each block of data, when it reports on one request and accepts a new one.

By running the device program in sync with the tape drive, the high-performance features of the HP 7978A can be implemented. Read aheads are performed by the device program's independently reading records in anticipation




Fig. 1. The firmware of the HP 7978A Magnetic Tape Subsystem contains three main modules: channel program, device program, and diagnostic program (not shown).

of the host's needs. The only interaction required is for the channel program to indicate when conditions are right for the host to desire read records. During writes, it is the implementation of immediate response that gives the drive its high performance. However, when the channel program is operating in immediate response mode, a report has already been sent to the host, leaving the device program free to write buffered data in the most convenient way. The device program holds off writing the data until a sufficient amount of data has been placed in the buffer or until a moderate timeout period occurs. This write holdoff groups write operations together and avoids an appearance of thrashing the tape back and forth.

Diagnostics

Diagnostics in the HP 7978A are split between the read/write functional electronics and the servo electronics. A 68000 microprocessor controls the read/write circuitry while a parallel 8051 processor controls the servo system. The 8051 accepts motion/diagnostic commands from the 68000 and returns status.

The 68000-based diagnostics verify the hardware operations of the HP-IB, data buffering and manipulation, and the read/write circuitry. The 8051-based diagnostics verify the servo hardware: buffer tension arm, speed encoder, motors, and servo control board.

The goal throughout the HP 7978A's diagnostics is to locate a failure to the board level. Once a failure is located, no further attempt is made to uncover any other failures. By testing from a core of independent hardware outward to the increasingly dependent hardware, the number of failures recorded without specific causes is minimized.

Approximately 75% of the drive is tested at power-on. The remaining 25% of the drive can be tested using the diagnostic tape tests or through normal operation. Any drive failure that is discovered is encoded into the host command report status and the internal diagnostic error log.

The HP 7978A has 62 diagnostic functions. Sixteen of these functions are performed during power-on self-test. The diagnostic functions can be categorized as follows: 12

utility/environment setting functions, 16 servo system tests, 15 board-level tests for the read/write functional electronics, 10 tape exercising functions, and 9 special functions of the HP 7978A. The utility/environment functions allow front-panel access to internal drive information and set the diagnostic test environments.

The servo system and read/write functional tests are designed to exercise a board or major subassembly. Examples of these tests include injecting signals into a circuit and verifying the outputs, or performing data loopbacks along the functional data paths.

The tape exercising functions allow reading and writing data from and to tape.

The nine special functions are: host-to-drive data loopback, remote download diagnostics, local and remote firmware update, image dumps of RAM and EEPROM, programmatically adjusting the read gain values, monitoring the amount of tape traveling over the read/write head, and an excessive soft-error rate function.

Execution of the HP 7978A's diagnostics is quite simple. A diagnostic test can be requested from the drive's front panel or through an HP-IB command from a host computer. Most diagnostics can be executed from either interface. However, some diagnostic tests require that special conditions be met before execution and some simply are not available from both interfaces. For example, the power-on self-test can be accessed remotely through a host or locally by an operator. Although the power-on self-test can be accessed remotely, the tape drive must be off-line for the test to be executed. If the drive is on-line, an error message is returned indicating the reason for nonexecution.

Acknowledgments

The authors acknowledge the efforts of other members of the firmware design team. The device program was written in part by Bruce Thompson, Dave Ellis, and Dave Burton. Diagnostics were written in part by Kevin Wilson and Dave Van Maren (servo diagnostics). Thanks also to Dave Jarrett, Doug Domel, and Alfred Natrasevski, who helped with the firmware QA effort.

Authors

March 1985

4 Maintenance Management

Irving Bunton, Jr.



Irving Bunton is a native of Chicago, Illinois and a graduate of the University of California at Berkeley (BSCS 1983). He joined HP's Manufacturing Productivity Division in 1983 and has contributed to the development of HP Maintenance Management. He is a resident of Oakland, California and enjoys music, especially jazz and blues.

Joseph L. Malin



A native of Washington, D.C., Joe Malin earned a BS degree in chemistry from Haverford College in 1978. He also did graduate work at the University of Pennsylvania. At HP since 1982, he worked on HP Maintenance Management and now develops manufacturing management software for the HP 3000 system. Joe's professional interests include manufacturing systems design and software engineering techniques. He also teaches Pascal programming at a local college. Now a resident of Cupertino, California, he enjoys skiing, drawing cartoons, and bicycling.

11 Half-Inch Tape Drive

Hoyle L. Curtis



Hoyle Curtis was born in Petersburg, Texas and attended Texas Tech University (BSME 1971 and MSME 1972). After joining HP in 1973 he contributed to the development of the HP 9871A Printer, was an R&D project manager for several printer projects and

for the HP 9895A Disc Drive, and was both an R&D project manager and section manager for the HP 7978A. Now a resident of Fort Collins, Colorado, he likes aquatic sports including sailing, windsurfing, and scuba diving. He participated in the Hobie 18 National Championships for sailing.

Richard T. Turley



With HP since 1977, Rick Turley managed the HP 7978A Tape Drive project and also contributed to the electronic design for the unit. Earlier, he was both a development and a production engineer for the HP 9876A Printer. Rick was born in Cranford, New Jersey

and graduated from Case Western Reserve University in 1977 with undergraduate degrees in both electrical engineering and mathematics. He continued his education at Colorado State University, from which he received an MSCS degree in 1981. He also attended Stanford University as an HP fellow in 1984 and received an MSEE degree. He lives in Fort Collins, Colorado with his wife and young son and enjoys hiking and cross-country skiing.

16 Write and Read Recovery Systems

Wayne T. Gregory



Tom Gregory joined HP's Civil Engineering Division in 1979 and was the responsible production engineer for the HP 3805, 3808, and 3810 distance and angle meters. After transferring to the Greeley Division, he contributed to the analog electronic design of the HP 7978A. Tom was born in Oceanside, California and now lives in Loveland, Colorado with his wife. In his spare time he enjoys backpacking, chess, and skiing.

19 Digital Formatting

Jimmy L. Shafer



Jim Shafer was born in Shreveport, Louisiana, and attended Louisiana Technical University, from which he received a BSEE degree in 1981. He joined HP in the same year and was responsible for several aspects of the electronic design of the HP 7978A. He also had pro-

duction engineering responsibilities after the product was introduced. Most recently he has worked on VLSI design for hard disc drives. Jim lives with his wife and new baby in Greeley, Colorado. His interests include skiing, softball, hiking, hunting, fishing, and camping.

25 Tape Drive Hardware

Robert D. Emmerich



Currently lead engineer for the Greeley Division technical analysis lab, Robert "Diamond Bob" Emmerich has been with HP since 1976. His contributions include work on the HP 250 Computer, NMOS II process control and IC testing, quality testing for the HP 9895A

Disc Drive, and evaluation of wear on flexible disc media. A native of Colorado, Bob was born in Denver, was educated at the University of Colorado (BSEE 1976), and is a resident of Loveland. He says he spends his spare time "reconstructing electronic hulks."

John W. Dong



John Dong was born in Phoenix, Arizona and educated at the Massachusetts Institute of Technology, from which he received the degrees of BSME (1975) and MSME (1976). With HP since 1976, he was the product designer for the HP 9876A Printer, a production

engineer for the HP 9825 and 9815 Computers, and a mechanical designer for the HP 7978A. John's work on the head plate assembly for the 7978A has resulted in a patent application. He is a registered professional engineer in the state of Colorado and a resident of Fort Collins. When he is not working on the basement of his house, he enjoys skiing and hiking.

David J. Van Maren



Dave Van Maren has been with HP since 1980, the same year he was awarded a BSEE degree from the University of Wyoming. He first worked in the Vancouver Division on the HP 293X Printers and then was a production engineer in the Greeley Division. He

was responsible for the servo controller firmware for the HP 7978A. Dave is a native of Casper, Wyoming and presently lives with his wife and two sons in Fort Collins, Colorado. He is a member of Right to Life and both he and his wife are instructors in natural family planning methods.

29 Tape Drive Firmware

David W. Ruska



A native of Detroit, Michigan, Dave Ruska attended Michigan Technological University, from which he received a BSEE degree in 1982. After coming to HP in the same year, he was responsible for channel program and HP-IB interface specifications for the HP

7978A and 7974A Tape Drives. Dave is a resident of Greeley, Colorado and has many outside interests. He raises Arabian horses, plays the piano, and is currently designing a digital music synthesizer. He also makes stained glass windows and enjoys photography and bicycling.

Bradfred W. Culp



Born in Ogden, Utah, Brad Culp received a BSEE degree from Virginia Polytechnic Institute in 1978 and joined HP in the same year. He resumed his studies in 1984 through the HP resident fellowship program and recently earned an MSEE degree from the

University of Illinois. At HP, he has contributed to the development of an inkjet printer as well as the HP 7978A. Brad is a resident of Fort Collins, Colorado and enjoys hiking, bicycling, and skiing.

Virgil K. Russon



Virgil Russon was educated at Brigham Young University, from which he received BSME and MSME degrees in the same year (1979). At HP since 1981, he worked on the buffer assembly and firmware for the HP 7978A Tape Drive. Virgil was born in Lehi, Utah and

now lives in Greeley, Colorado with his wife and four children. He is active in his church and works on numerous home projects. He also enjoys woodworking and art.

John A. Ruf



A native of Milwaukee, Wisconsin, John Ruf served for four years in the U.S. Air Force before receiving a BSCS degree from the University of Wisconsin at La Crosse (1981). A short time later he joined HP's Greeley Division, where he has

worked as a software engineer on the 7978A Tape Drive. John lives with his wife and two sons in Fort Collins, Colorado. He enjoys tennis, jogging, and developing software design aids.

Alan J. Richards



Now R&D manager for Colorado Time Systems in Loveland, Colorado, Alan Richards came to HP in 1968. He was a project leader or manager on the HP 9825A Tape System, the HP 9874A Graphics Digitizer, and the HP 7978A. In addition, he was

a production engineering manager for the Greeley Division and HP's campus recruiting manager at Kansas State University. Alan lives in Loveland, Colorado, is married, and has two children. He likes hiking, backpacking, orienteering, and cross-country skiing. Last summer he realized a lifelong ambition by taking his family on an extended vacation to Alaska.

34 Winchester Backup

Sterling J. Mortensen



Sterling Mortensen received his BSEE degree from Brigham Young University in 1974 and came to HP in the same year. As a design engineer, production engineer and a project manager, he has contributed to the hardware and software design of the HP

250 Computer and the HP 9895 and HP 82901 Disc Drives. He was the software architect for the HP 9144A. He was born in Rexburg, Idaho, lives in Fort Collins, Colorado, is married and has twin sons as well as two other children. His interests include target shooting, photography, and church activities.

Donald A. DiTommaso



Currently manager of a new product development program at HP Laboratories in Bristol, England, Don DiTommaso was program manager for the HP 9144A. He also originated the idea for the project, managed design work on the 9144A drive mechanism, and contributed to controller design. Don is an alumnus of the University of Cincinnati (BSEE 1973) and the University of Colorado (MSEE 1975) and came to HP in 1975. He first worked on the design for HP 9845 Computers and later set up an R&D laboratory team in HP's Singapore Operation. He is also named as a coinventor on a patent on the design of the HP 9144A read after write technique.

John C. Becker



Now project manager for the HP 9144A, John Becker has been with HP since 1964. He also managed the development effort for the 9144A drive mechanism. Earlier, he managed a production engineering group and contributed to the development of the HP 9835

and 9825 Computers. He is named coinventor on a reed relay patent and is an author of papers published by the IEEE and the HP Journal. Born in Lincoln, Nebraska, John received his BSEE and MSEE degrees from Colorado State University in 1972 and 1974. He currently resides in Loveland, Colorado, is married, and has four children. His outside pursuits include tennis, jogging, sailing, church activities, and travel.

37 Tape/Disc Controller

Mark L. Gembarowski

Author's biography appears elsewhere in this section.

Craig L. Miller



A native of Sandy, Utah, Craig Miller attended Brigham Young University (BSEE and MSEE, both 1976) and joined HP in 1977. He was a hardware designer for the HP 9144A as well as for the HP 250 Computer and the HP 9895A Disc Drive. At present, he is working on the design of a mass storage controller. Craig is married, has five children, and lives in Fort Collins, Colorado. His interests include softball, volleyball, basketball, church activities, and travel. He also coaches for an elementary school soccer league.

39 Tape Data Integrity

K. Douglas Gennetten



Doug Gennetten came to HP in 1978 and has worked on the HP 9144A controller and mechanism, as well as on high-speed dynamic RAM test hardware and fault detection algorithms. He received his BSEE from Walla Walla College in 1978 and is continuing graduate

level study at Colorado State University and at National Technical University. Doug lives in Fort Collins, Colorado and is married. His outside interests include skiing, sailing, hiking, woodworking, visual arts, and exploring the back country by jeep.

44 Head/Tape Interface

Mark E. Wanger



Mark Wanger joined HP in 1980 with a background in satellite control systems. He worked first as a production engineer, and then became an R&D engineer and contributed to tape path control and hardware design for the HP 9144A. A native of Long Beach,

California, Mark studied at the University of California at Santa Barbara (BSME 1979), the Massachusetts Institute of Technology (MSME 1980), and Colorado State University (MSEE 1983). He lives in Fort Collins, is married, and enjoys backpacking, basketball, and sailing.

David J. Schmeling



Dave Schmeling has worked on the HP 9144A Tape Drive since coming to HP in 1979. He contributed to the mechanical design, environmental testing, and production of the product. Dave was born in Santa Ana, California and graduated from Stanford

University with an MSME degree in 1979. He lives in Fort Collins, Colorado and is newly married. His outside interests include Bible study, skiing, woodworking, and finishing the basement of his home.

Walter L. Auyer



A native of central New York State, Walt Auyer joined HP in 1971 and has worked as an R&D laboratory engineer in both the Civil Engineering Division and the Greeley Division. He received an MSME degree from Arizona State University in 1969. Before coming to HP, he was employed in the aerospace industry. Walt lives in Fort Collins, Colorado, is married, and has four children. He enjoys golf, bridge, and vacation travel with his family.

Charles H. McConica



A native of Ventura, California, Chuck McConica is an alumnus of the University of Denver (BSME 1974 and MBA 1976). He also studied at Stanford University, where he received an MSAM in 1975 and is presently a PhD candidate in mechanical engineering.

Since joining HP in 1979, Chuck has contributed to the development of inkjet head technologies and the manufacturing process for flexible disc heads and magnetic tape head mounting. His work has resulted in a patent on inkjet multiple-drop suppression. Chuck is a resident of Windsor, Colorado, is married, and has two children. His interests include family activities, backpacking, skiing, bicycle touring, and sports cars. He also serves as a soccer coach and helps local high school math and science teachers with lecture material.

47 Software Methodology

Mark L. Gembarowski



Mark Gembarowski joined HP in 1978 after receiving an electrical engineering degree from Michigan State University. In his first assignments, he contributed to the hardware design of the HP 9111A Graphics Tablet. Currently, he is manager for a half-

inch tape drive project. A native of Saginaw, Michigan, he now lives in Fort Collins, Colorado and enjoys softball, water and snow skiing, and photography. He also started a Fort Collins youth soccer program in 1979.

Low-Cost, Highly Reliable Tape Backup for Winchester Disc Drives

Designed for use on small to midrange computer systems, this new quarter-inch cartridge tape drive packs up to 67 megabytes onto a single cartridge.

by John C. Becker, Donald A. DiTommaso, and Sterling J. Mortensen

AS WINCHESTER DISC DRIVES have dropped in size and price, a matching backup peripheral has been needed that is also smaller in size and lower in price. The HP 9144A, Fig. 1, meets both of these needs. This peripheral can be used with HP disc drives from 5 to 132 Mbytes. It is supported on the HP 3000 family of computers, the HP 1000 A-Series, the HP 9000 Series 200 and 500 Systems, and the HP 150 Personal Computer. It is also available in one box with a Winchester disc drive as the HP 7946A, which also gives one-button full disc image store and restore. The HP 9144A can therefore meet the needs of a wide range of computer users.

The HP 9144A stores data on removable cartridges that contain 1/4-inch-wide magnetic recording tape. The cartridges come in two different capacities. The lower-capacity cartridge has 150 feet of tape with a data capacity of 16.7 Mbytes, and the higher-capacity cartridge has 600 feet of tape with a data capacity of 67 Mbytes. The HP 9144A has read after write capability, so that data can be rewritten automatically if there is an error during the write process. There is no time penalty for this read checking. It also has error correction capability, so that entire 256-byte frames can be replaced. The cartridges can be interchanged with the tape drive used in the HP 7908, 7911, 7912, and 7914 Tape/Disc Drives.

Although Winchester disc backup is its primary use, the HP 9144A also meets several other needs, such as data interchange and software distribution. High data interchange reliability from one HP 9144A to another was an important design criterion. This allows data exchange from one system to another. Software distribution can be a problem when the software takes many flexible discs, but with an HP 9144A, it can all be sent out on a single tape cartridge.

With the HP 9144A, the user can back up a 60-Mbyte disc in as little as 30 minutes. A 12K-byte buffer in the drive and optimized system software minimize waiting for data transfers and commands from the host. Operating in streaming mode, the drive provides up to 2-Mbyte/minute backup performance, depending on host software.

To protect the user's data and minimize drive failures, an electromechanical servo control feature is built into the tape drive. This increases both tape and drive motor life by providing gentle acceleration and deceleration as the drive stops and starts.

Internal diagnostics provide quick isolation of failure and ease of repair. This means that the user will experience minimum downtime and lower maintenance costs. A light

on the front panel warns if the drive is inoperative and rear-panel diagnostics indicate which replaceable unit has failed.

The HP 9144A fills the same space on a desk as an in-basket. It can also be packaged in a pod configuration or mounted in a standard 19-inch rack.

Disc-Like Format

User data stored on the HP 9144A's cartridge is formatted into blocks that contain 1024 bytes each. Each block contains a header that tells how many bytes are stored in that block, so that partial data blocks can be stored easily. The format allows any block to be accessed separately for reading or writing. Therefore, the tape can be written to and read from just like a disc. This is a distinct advantage because all the operating system commands that can be used on discs can also be used on the tape. The tape can have a directory just like the disc. Files can be stored onto the tape and then retrieved through the directory just as if they were on a disc. Files can be deleted and new ones added anywhere there is free space according to the directory. Alternatively, a file mark command exists so that the drive



Fig. 1. The HP 9144A 1/4-Inch Cartridge Tape Drive offers convenient, low-cost backup of 67 megabytes on a single cartridge. Read after write capability and exclusive-OR error correction provide improved data reliability.

can be used like 1/2-inch tape drives, where different files are found by searching for the appropriate tape mark. Different computer operating utilities will use the cartridge tape in one of these two modes depending on whether the utilities were originally written to work with discs or 1/2-inch tapes.

There are 16 tracks on the removable cartridge. The drive can write or read one track at a time. These tracks are written in serpentine fashion with the even tracks going from left to right and the odd tracks going from right to left. At the end of one track the head is just moved up to the next track and the motor starts moving the tape in the opposite direction.

The combination of serpentine format and read after write capability requires the magnetic head to have two pairs of gaps. One pair has the read gap on the left side of the write gap so that read after write can be done when the tape is moving from right to left. The other pair has the read gap on the right side of the write gap so that read after write can be done when the tape is moving from left to right.

Data Integrity

Providing superior data integrity was a primary goal for the HP 9144A. Achieving a low error rate on a cartridge tape product requires more fault tolerant design than hard or flexible magnetic media products. Tape, when moving, continuously changes shape as it conforms to the hubs, guides, and head. This can create a higher level of internally generated contamination. A robust design is required to tolerate both particle contamination and other causes of errors such as defects in the tape and externally generated contaminants.

The write and read processes are the two important areas of concern in providing a high level of data integrity. The HP 9144's write process provides assurance that data is written on the tape correctly and the read process incorporates extensive features to ensure data recovery.

Five levels of data protection are implemented in the HP 9144A (see article, page 39). Starting with the media, only certified data cartridges are specified for use on the HP 9144A. Defects that exist on the tape are bypassed by use of a sparing table, which is generated when the tape is certified. A defective portion of the tape will not be used; that block will be replaced by one of a group of extra blocks which have been reserved as spare blocks.

The next step to ensure that data is written correctly is the read after write process. The HP 9144A uses a four-gap wide-write, narrow-read magnetic head. All data written on the magnetic tape is read immediately after the recording process.

Data retrieval is the next level of protection; highly optimized data separation and read algorithms are incorporated in the HP 9144A.

Error detection and error correction provide another level of protection. A CRC (cyclic redundancy check) error detection process is used on every frame and key. Also, 50% data redundancy is recorded to provide error correction. A block of 1024 bytes is completely recoverable if as much as any two adjacent 256-byte frames are completely destroyed.

The last level of protection is media monitoring. The

media are monitored by recording the number of passes in the user's log on the tape. After a cartridge has exceeded a maximum number of either insertions or passes, the write-protect light located on the front panel flashes, warning the user to back up the data on that cartridge.

These protection levels provide a very high confidence of correct data recovery.

Reliability

A significant improvement in reliability was a major goal for the HP 9144A. The specific objective established to meet this goal was to improve the reliability by a factor of two over existing comparable products. This would be measured by both in-house and field failure rates.

The approach used to reach this goal was to take specific actions in design, testing, and manufacturing. The steps taken in the design were derating of components, using a higher level of circuit integration, and maintaining a low temperature rise within the enclosure. For example, all components are derated to less than 70% of their specified maximum power. A fan is used to limit the internal temperature rise to less than 5°C. The controller uses two custom VLSI chips and two gate arrays to minimize the component count (see article, page 37).

An extensive test program was implemented, including environmental, error rate, compatibility, and strife testing on 70 units. Strife (stress and life) testing was designed to test beyond specifications to find out what would fail, understand why, and implement improvements.

Process control was implemented in production. Control charts are used to monitor critical parameters such as head alignment. Production audit tests are used to verify error rate performance. Compatibility and strife testing are used as tools for continuously improving the quality of the product.

The results are very encouraging. Based upon the latest environmental and strife testing, strife yields have increased significantly. Building the HP 9144A in production has been extremely smooth, which is indicative of both a solid design and a solid production process.

Acknowledgments

It is difficult to thank all who contributed to this major project. Special thanks to the unnamed people throughout the division from manufacturing, quality assurance, marketing, and finance. Thanks go to lab management for their guidance on this project. Rex James and Rick Spangler as lab managers provided long-term support. Frank Carau and Don Stavely as section managers provided both product and technical support. As project managers, we express our thanks to the engineers who really made the HP 9144A possible. The electrical design team consisted of Doug Gennetten, Larry Paille, Joe Eccher, Tom Oliver, Craig Miller, and Lynn Watson. Walt Auyer, Dave Schmeling, Keith Walden, Dan Dauner, Chuck McConica, and Mark Wanger were members of the mechanical design team. The software team included Gordon Nuttall, Mark Gembarowski, Steve Henry, Frances Cowan, Bob Kaphammer, Perry Wells, and Ken Nielsen. Gene Briles, Mike Niquette, and Peter Way were members of the system integration team. The transfer team members were Kevin Bier, Leo Embry, Lonnie Ford, and Terry Tuttle.

A Design Methodology for Today's Customers

The HP 9144A Tape Drive was designed by Hewlett-Packard engineers to achieve a high level of customer satisfaction. To do this it was necessary to go beyond the standard approach of designing a drive and controller to meet specifications. The marketplace challenged us to go beyond the specifications to provide ample margins to accommodate real-world variances and to provide consistent quality and HP reliability.

Why was such an approach necessary? Let's look at an example. The DC600 Tape Cartridge is a high-technology product providing over 100 megabytes of unformatted storage capacity in a very convenient and low-cost package. Tape of comparable quality, such as 1/2-in tape, is used mostly in environmentally controlled computer rooms and is typically handled by technicians familiar with high-density computer tape. On the other hand, the 1/4-in tape cartridge reaches first-time computer users in a wide range of environments and applications and is most often treated like an audio cassette. Yet, these customers need high capacity and reliability and demand ease of use. To meet these requirements, the media and drive system of the HP 9144A had to achieve four times the track density of state-of-the-art 1/2-in GCR (group code recording) tape drives.

Clearly, we had to use an innovative design approach and invent the product from the ground up with the customer in mind. We had to design the mechanism, hardware, and software by forming interdisciplinary teams that were highly focused on a problem area and whose goals were to provide reliability and high performance with simple, elegant, and low-cost designs. These teams felt a high level of ownership of a design area and were able to design margin into the product to account for real-world variances and to identify all the key variables that would affect the long-term performance of the system.

In harsh environments, for example, the cartridge sometimes exhibits a sudden burst of high-frequency speed variation that typically causes tape drives considerable problems with data integrity. Our design philosophy challenged us to develop a third-order phase-locked loop that can track this high-frequency variation in tape speed. We integrated the loop with the firmware of our formatter to provide our customers with smooth streaming performance and data integrity.

Another reality of our marketplace is that customers want this streaming tape drive performance in spite of the fact that often the data is gathered and transferred to the tape drive in sequences of short bursts. These short data transfers can wear out the cartridge, cause high stresses on the drive motor, and typically result in poor throughput. To solve this problem, we chose a drive motor with a superior brush design, then surrounded it with a control system that very gently accelerates the tape, and finally, provided software enhancements to help the host system achieve full streaming mode performance even for short data transfers.

Our team approach was especially fruitful in combining software with mechanical design. We use software algorithms to

compensate for hysteresis in the head stepping mechanism and even to adjust the placement of data tracks to accommodate thermal bending of the cartridge baseplate.

We also took this methodology outside of HP to our key vendors, who became part of our team, designing margin into their components. In the area of the read/write head we worked together closely and produced a design that is totally insensitive to $\pm 25\%$ off-track alignment. We surrounded this with a mechanism design that minimizes tape wander in the cartridge and a read channel preamp that eases the design constraints on the head and yet achieves full compatibility with HP's current cartridge product line.

Now that we had a product that was designed from the customer's point of view, with margin to accommodate real-world variances, we challenged our team to submit it to a comprehensive test program that was planned to guarantee the quality, reliability, and consistency of every unit we produce. Each designer was responsible for creating a stand-alone test system that would verify the functional performance and margin of the design. Software, mechanism, and electronics designers teamed up to develop functional tests that prove the consistency of the design and demonstrate that all the key variables that affect performance have been identified and brought under statistical control in the *design* phase of the project. These functional tests are now in production; all assemblies receive complete functional tests after automatic component verification. As a result, the current HP 9144A error rate is so solid that we have statistically proven over an order of magnitude of margin on production units over the complete range of environments. All key variables are charted to be in control on a regular basis in production.

Acknowledgments

Such a design approach required the active participation of everyone on the project. Sterling Mortensen architected and managed our programs of structured software design, software QA, and team building. John Becker, who contributed to the testing effort that helped HP to win the Deming Award, designed a program of strife and environmental testing that, under his guidance, was highly effective in identifying system problems and eliminating potential hardware problems.

The final principle of our design methodology is to provide for continuous product improvement. Thus we have put into motion a program of LSI for the controller and cost reductions for the head and drive electronics. John managed a coordinated program between the lab and manufacturing that has resulted in a reduction of controller electronics from three boards to one. These improvements provide equivalent or higher performance and are fully interchangeable.

Donald A. DiTommaso
Project Manager
Greeley Division

Tape/Disc Controller Serves Integrated Peripherals

by Craig L. Miller and Mark L. Gembarowski

THE DESIGN GOAL for the controller of the HP 9144A Tape Drive was to provide a structure permitting one or two mass storage peripherals to operate cohesively for on-line and backup functions. The product's architecture design was a joint venture with HP's Disc Memory Division to provide a common controller for integrated disc and tape peripherals.

The HP 9144A controller provides high performance and flexibility. The controller architecture allows either a fixed disc, a 1/4-inch tape, or both to operate in a common environment. A related, combined disc/tape product, the HP 7946A, provides not only read and write capabilities to and from the host but also local data transfers from disc to tape or tape to disc.

The controller functions are segmented into two areas. The functions common to basic mass storage controllers are lumped together onto one controller called the host dependent controller (HDC). Functions that are specific to the tape or disc drive are contained on a device dependent controller (DDC). One HDC can support one or two DDCs (see Fig. 1).

HDC Characteristics

The HDC interfaces between the host computer and the DDC, which provides low-level control of the mass storage mechanism. The HDC provides high-level functions, including:

- Providing the host interface (HP-IB/IEEE 488) and handling the associated protocol
- Operating a microprocessor (6809) for basic control and status reporting
- Providing buffer RAM, enabling data to be read from the device and verified or corrected before being transferred to the host
- DMA capability.

Fig. 2 is a block diagram of the HDC.

The DMA controller is separated into two independent channels, input and output. Each channel has a FIFO buffer

which is four locations deep. Addresses for transfers are programmed into the FIFO buffer. Each FIFO buffer location is independent. Each channel can be monitored and new addresses can be programmed into the FIFO buffer during a transfer to allow high performance and uninterrupted transfers. Each channel can support 1-Mbyte/s concurrent transfers.

For example, suppose the host computer wants to perform a selective file backup of the fixed disc onto the 1/4-inch tape. The host computer sends commands to the device specifying a file to be transferred from the disc to the tape unit. While that file is being copied locally from the disc to tape, the host can send another command to copy a different file from the disc to tape. Continuing this sequence allows a user to back up the disc drive selectively. The fixed disc could possibly keep the tape streaming in this selective file backup since the buffer RAM (approximately 11K bytes) on the HDC provides data speed matching (see "Making a Start/Stop Tape Stream," later in this article). The tape drive transfers 1K bytes of data every 30 ms, allowing the fixed disc sufficient time to perform most seeks and to read additional data.

If the host computer initiates a backup operation from the disc to tape, the host is free to use the HP-IB for other transactions during the backup operation, since the backup is local to the peripheral. An HP-IB parallel poll will indicate when the backup is complete.

The HDC runs a small operating system serving three users: the host interface firmware, the firmware to control the tape, and the firmware to control the fixed disc in a dual-peripheral configuration. The 6809 microprocessor is time multiplexed between the users. Approximately 30% of every millisecond is devoted to each user, and there is a total of 10% overhead in switching between users.

The interface between the HDC and the DDCs has a fixed definition to provide for independent software and hardware designs. There are two buses with associated strobes or handshake lines. One bus is the command and

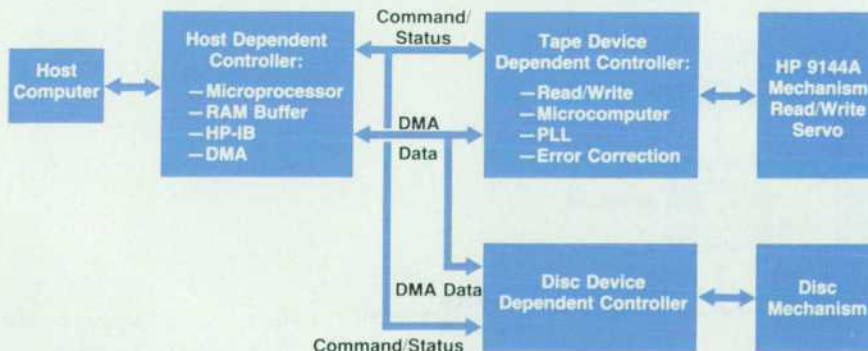


Fig. 1. The controller for the HP 9144A Tape Drive can interface one or two mass storage peripherals (tape or disc) to a host computer. The controller functions are partitioned into a host dependent controller and one or two device dependent controllers.

status bus for commands from the HDC to the DDC and DDC status monitoring by the HDC. The other bus is a DMA data bus for data transferred to and from the tape or disc. The separation of these two functions allows monitoring of status and issuing of commands to the tape or disc during data transfers.

DDC Characteristics

Functions that are specific to the tape drive are contained on a device dependent controller (DDC). These functions are reading and decoding data, writing and encoding data, checking and correcting errors, and basic device-specific control.

The DDC for the 1/4-inch tape interprets the specific requirements for the tape. There are two basic processes on the DDC: the write process and the read and data recovery process. A 6805 microcomputer oversees these two processes.

The write process, under control of the microcomputer, requests data from the HDC. The data is serialized and the appropriate identification, error detection, error correction, and formatting data are appended to the user data.

The read and data recovery process locks onto the serial data stream and decodes the format to identify the tape position. The read algorithms are tightly coupled to the data recovery circuit for controlling the data separator for optimal performance. The read process also strips away the overhead information and provides the user data to the HDC for later transmission to the host. Error correction is performed in real time by exclusive-OR hardware (see article, page 39).

The microcomputer controls both the read and the write processes. It also communicates with the HDC for commands and status. Physical tape positioning for seek, read, and write operations is controlled by the microcomputer.

The beginning and end of the tape are sensed electronically by the read process algorithms. The beginning-of-tape and end-of-tape positions are written in a special format

found nowhere else on the tape. The read algorithms sense these conditions and inform the servo controller to stop the tape motion. For fail-safe protection, the servo automatically stops the tape if the signal is lost. This prevents unspooling of the tape. Previous tape products used LED sensors to sense the beginning and end of the tape. These proved to be less reliable as the product became older or was subjected to heavy use.

The read algorithms associated with the data recovery circuit are very fault tolerant, enabling the read process to recover from many abnormal situations.

Four custom LSI circuits were developed for the HDC and the DDC. The first two are CMOS gate arrays, one having 2000 gates and integrating the DMA controller, real-time clock, and programmable time interrupt on the HDC, and the other having 440 gates and integrating the high-speed precompensation circuit, DMA handshaking, VCO control, and data recovery circuit for the DDC. These two integrations required the development and enhancement of a set of tools initiated by HP's Roseville Division. These tools and thorough simulation allowed the two gate arrays to work exactly as defined the first time.

The tools are based on HP's Testaid-IC logic simulator program. Timing and fault analysis can be performed using this program. The gate array specifications were added to the Testaid-IC library on the HP 1000 Computer System. The timing, loading, and functional descriptions of the library parts were loaded into a data base. Drawings of each library part were created and the design and connectivity were entered using HPDraw on the HP 3000. The drawing file was transferred to the HP 1000, where timing was added to the drawing data base. Functional, performance, and fault analysis were performed on the HP 1000 using Testaid-IC. The files were converted to the gate array format and transferred to the gate array vendor. The tools and analysis were comprehensive enough to provide full functional performance on the first design pass.

The other two LSI circuits are NMOS semicustom state

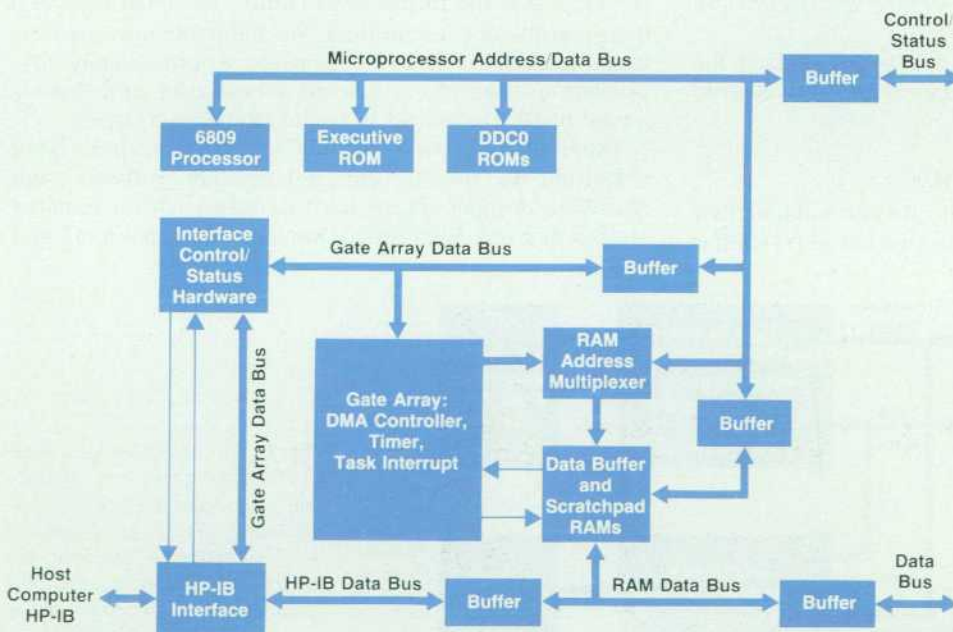


Fig. 2. Host dependent controller block diagram.

machines. Advanced tools developed by HP's Greeley Division for a universal state machine (a semicustom part) enabled fast and efficient integration of the read and write state machine algorithms that encode and decode the HP 9144A data format. These state machines also control the read and write functions.

The error rate was improved significantly by integrating the data recovery circuits in a gate array. The integration allows tight control of timing and closely matched delays in portions of the circuit.

Making a Start/Stop Tape Stream

One of the key performance improvements of the HP 9144A over the tape drives in the HP 7908, 7912, and 7914 is the addition of streaming mode. Streaming is a mode of operation in which the tape continues to write or read data without repositioning. The HP 7908/12/14 Tape Drives would stream only for the short amount of time that the host computer could buffer data. On the single-controller versions of these products, the tape would stream only for an image copy of the disc.

The HP 9144A has a significant performance advantage over its predecessors because streaming is not dependent on host buffer size, and streaming can be maintained even on selective file store and restore operations. The HP 9144A saves over one second per host transfer because it has streaming built-in.

How is this improvement realized? In normal start/stop operation, the drives (HP 9144A and HP 7908/12/14) reposition the tape upon completion of the execution phase of the transfer. (In execution phase, the data is sent to or from the host). Once the repositioning is completed, a report is sent to the host. Repositioning moves the tape back far enough so that if the host sends a command to read or write the next sequential block on the tape, the drive can

accelerate and read or write the target block at the proper speed of 60 inches per second.

We make the tape drive stream rather than operate in start/stop mode by eliminating the repositioning at the completion of reading or writing the data. The stream mode of writing is called "immediate reporting." This name evolved from the fact that the report byte is sent to the host computer immediately after the last byte of data has been transferred to the tape drive controller's RAM buffer, rather than after it is all written to the tape (start/stop mode). The host, because it has seen the report byte, can queue another write command to the tape drive (provided it is for a sequential address) and the tape drive will stream as long as the command is sent before the RAM buffer is empty. The only change that is necessary in the host driver is the addition of the set option command to enable the immediate reporting mode.

Streaming while reading data for the host presents a different problem. The tape drive is the source of execution data for the host. Normally, the controller repositions the tape as soon as all the data has been sent to the host, and then sends the report byte. In the streaming mode, which is called "read ahead," the report byte is sent to the host as soon as the data transfer is complete, and the drive continues to read data, buffering it in the RAM on the tape drive controller. If the host sends another read command to the tape drive for the next sequential address, the tape drive will not have to reposition to get data it already has in RAM, and hence it will stream. This streaming will continue as long as another sequential read command is sent before the RAM buffer is filled with read ahead data. Read ahead mode is enabled with exactly the same set options command and parameter byte that are used to enable immediate reporting.

Cartridge Tape Data Integrity Ensured at Five Levels

by K. Douglas Gennetten

THE NEW HP 9144A ¼-INCH CARTRIDGE Tape Drive provides a convenient, low-cost, high-reliability alternative to flexible discs and ½-inch tapes for backing up valuable data. The remarkable data reliability of the HP 9144A is ensured by the melding of several important data protection techniques. Extensive testing has proved that the typical HP 9144A's data error rate is more than ten times better than the specified rate of less than 1 error in 10¹⁰ bits. This performance is achieved through careful attention to design margin in many areas of the product.

Read After Write

The HP 9144A provides the first level of data integrity through precise flux-change placement on the media followed by immediate verification of written data.

Encoded data is first preconditioned to compensate for peak shift or spreading, which occurs when two flux changes are placed in close proximity. Each data edge is phase shifted by an amount determined by the surrounding data pattern. Two flux changes placed close together will tend to repel and shift apart in opposite directions. Precompensation of these two edges improves this condition by

writing the edges closer than nominal by a controlled amount. Upon relaxation, the edges return to their proper positions. The write circuit precisely records the flux changes into the media, removing any previous data.

The final written position of these flux changes is tightly controlled to ensure adequate performance margin in the data retrieval process. The read circuits are carefully designed to introduce no significant phase shift onto the read data. This is done using high-order phase-linear filters.

The read after write capability enables the drive to detect and automatically correct or spare out errors caused by media defects or foreign particles. This powerful advantage, typically found only in larger, more expensive tape storage systems, guarantees that the recorded data is initially defect-free.

Certified Cartridge

The second level of data integrity is provided by the HP DC600 High-Density Data Cartridge. This cartridge contains magnetic media formulated from high-coercivity materials which are applied using a special process to reduce error-producing dropouts. These characteristics permit reliable, high-density, minimal-defect recording.

Cartridges used on the HP 9144A are preformatted and certified at the factory. This provides additional protection and allows immediate use of new cartridges. No time-consuming certification or formatting is required.

The HP 9144A Tape Drive uses a block-oriented format written over 16 tracks in a serpentine block-serial fashion. For example, all the blocks on track 0 are written in the forward direction first; then the head steps to track 1 where all the blocks are written in the reverse direction. This formatting method allows random access as well as streaming operation. The drive uses preformatted tapes with storage capacities of 16 megabytes (150 ft) and 67 megabytes (600 ft). Fig. 1 shows the tape format.

Each 1024-byte block is written between interblock preformatted keys placed vertically across the entire width of the tape. The blocks contain six 256-byte frames: four for user data and two for error correction. Even tracks are writ-

ten from the beginning of the tape (BOT) to the end of the tape (EOT). Odd tracks are written from EOT to BOT.

Keys contain the physical block number, which is also reproduced in each frame header along with other information describing the block. All frames and keys end with a 16-bit cyclic redundancy check (CRC) character, which provides very reliable error detection, and a postamble.

The tape format contains special prerecorded BOT and EOT patterns which allow solid-state beginning-of-tape and end-of-tape detection without the use of conventional optical sensors, which are more costly and less reliable.

Data Separation and Clock Recovery

The third level of data integrity assurance takes the form of accurate data retrieval. Phase-linear signal filtering followed by optimized data separation provides data recovery margin well above that found in other high-performance tape systems. Rigorous design and optimization methods (described below) have provided the HP 9144A with enhanced read margin over a wide spectrum of cartridge conditions. The data retrieval process is exceptionally tolerant to large levels of phase error caused by peak shift and noise as well as extremely rapid and large speed variations, both being common conditions in high-density, belt-driven cartridge tapes.

The user data to be recorded on the magnetic tape is first intermeshed with a clock in a form of modified frequency modulation (MFM) encoding. MFM data encoding, when combined with optimized, low-noise data recovery, enables higher bit densities and improved speed tracking compared with other commonly used tape codes. MFM also has the advantage of being very easy to implement. For each 1 a flux transition is written at the center of the bit cell and for each 00 pair a flux transition is written between the bit cells (see Fig. 2). Before being recorded onto the tape, each encoded transition is first precompensated (phase shifted) to counter the expected peak shift that occurs when flux reversals are placed side by side at high densities.

Data retrieval requires reliable separation of the MFM-en-

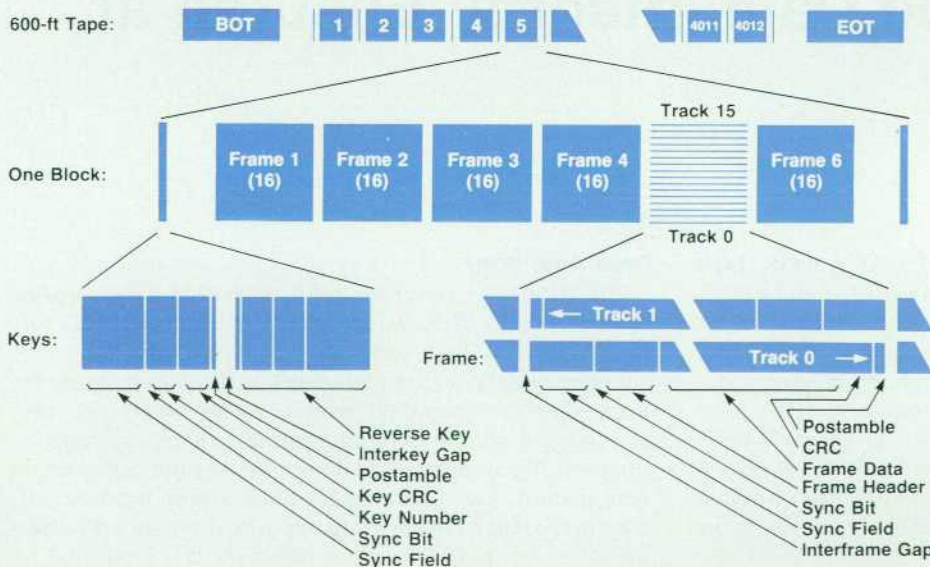


Fig. 1. DC600 tape format. All blocks on track 0 are written in the forward direction first. Then the head steps to track 1 and writes data in the reverse direction, and so on.




Fig. 2. MFM encoding example.

coded data and clock. This objective presents unusually difficult challenges when the data is stored on cartridge tape at very high flux densities. The encoded data read back from the tape is degraded by two major sources of distortion: tape speed variation and phase shifted data.

Tape speed variation in the DC600 cartridge is dominated by several large components at 30 Hz, 100 Hz, and 6000 Hz (using 60 inches per second tape speed), which are beyond the control bandwidth of the capstan servo system and, in fact, are not seen by the capstan servo encoder. Of these, the 6-kHz component is the most severe. When all three components are overlapped, they can add up to speed variations of over 15% of the normal read/write speed. Speed variations evoke frequency modulation of the data signal, which requires accurate demodulation during data separation.

Phase shift is caused by several sources. The most pronounced include imperfectly compensated peak shift, phase nonlinearity of the read head, and electrical noise. Phase shift evokes phase modulation of the data signal, which must be filtered during data separation. The HP 9144A read amplifier and filter electronics have a high degree of phase linearity and do not significantly add to the total phase shift.

The tracking of speed variations and the filtering of phase shifted data are done in the data separator by a phase-and-frequency-sensitive closed-loop servo system termed the phase-controlled loop, or PCL (Fig. 3). Higher and faster speed variations require quicker response by the PCL while higher phase shift magnitudes are best handled by a slower response. The resulting conflict requires a careful compromise and is handled particularly well by the PCL, which is an optimized third-order system (second-order systems are most common). The higher-order loop, when accurately optimized, is better able to make such compromises by minimizing the average and instantaneous real-time error (RTE), thus providing the necessary margin for reliable data recovery during the extreme circumstances commonly

found in inexpensive tape cartridges.

Data Separator

The HP 9144A data separator (Fig. 3) is composed of the phase-controlled loop,* an MFM decoder, gap and lock detectors, and a synchronous state machine. The state machine monitors gap, lock, and data status of the separator and controls the PCL modes.

Frequency-lock mode. This mode is invoked during sync field frequency acquisition, which occurs at the beginning of every data area, and during gaps, when the loop is frequency-locked to a crystal-based clock. The frequency-lock mode inserts a frequency divider ($\div 2$) in the PCL feedback path. This divider causes the VCO to acquire lock at twice the incoming data rate; this is necessary for providing the two-windows-per-bit needed for MFM. The frequency-lock mode also disables the pulse gate function of the phase/frequency detector, thus allowing the frequency of VAR (VCO/2) edges to be compared with the frequency of the REF (incoming data) edges.

Phase-control mode. This mode is used during all data tracking and decoding. It is invoked after the frequency-lock mode has achieved a stable locked condition in the loop, but before the data actually begins. The phase-control mode disables the VCO/2 divider, making all VCO edges available to the phase/frequency detector for phase comparison. The pulse gate function is enabled, blocking all extra VCO edges and removing the loop's frequency-discrimination capability. The phase detector attempts to match each REF edge with one and only one VAR edge. From one to three extra VAR edges occur between each REF edge. The phase/frequency detector matches each REF edge with the one VAR edge that is within one half window before or after the REF edge. Once a valid match is made, the phase/frequency detector sends a pulse to the loop filter

*The more common term "phase-locked loop" is avoided here because the PCL is never actually phase-locked as in a frequency synthesizer or FM receiver; it is frequency-locked but exhibits large, rapidly changing phase error.




Fig. 3. Data separator block diagram.

that is proportional to the phase error and of the appropriate sign. The loop filter then processes this signal and sends the filtered response to the VCO, which begins to nullify the detected phase error.

The lock detector is nothing more than the data decoder combined with a lock-detect algorithm in the state machine. When locked and in a sync field, the data decoder will have a stable output. Any cycle slip or other lock failure will show up as an unstable decoder output and is easily recognized.

The gap detector is a resettable counter that increments one count for each VCO clock and is reset by each data edge. When the count is at or above five, the circuit alerts the state machine that a gap or dropout has been detected. With the format used, it is important that gap detection be rapid and reliable. This circuit meets that need by using a tracking time base rather than an external crystal or one-shot timer. A gap is detected based on the count of bit times, rather than a fixed time period.

Propagation delays in a design such as this can easily accumulate and produce asymmetry in the form of large steady-state phase error (i.e., reduced RTE margin). This design incorporates delay equalization, which balances the data path from the data source to the phase/frequency detector with the data path from the source to the MFM decoder. The result is improved operating margin and enhanced data integrity.

PCL Optimization

Optimization of the PCL's tracking dynamics is of paramount importance. Two common optimization procedures for data recovery systems are 1) trial and error adjustment of the loop filter parameters while observing the error rate performance and tracking ability under typical data conditions, and 2) reduced-window accelerated testing in combination with the first procedure. Reduced-window testing artificially reduces the RTE margin by an adjustable amount during tracking of normal data. Another common testing method is subjecting the system to artificially generated data patterns from a pattern generator modulated by a variable-frequency source. This last method is of little value, since it simulates speed variation without any superimposed phase modulation. Under such conditions a loop will perform much better than in actual use where pattern-dependent peak shift and random noise induce phase modulation of the data, which greatly reduces the frequency demodulation (speed tracking) ability of the system. These optimization procedures do not provide best-compromise performance, which is necessary for ensuring the best overall error rate. These inadequate methods all incorrectly assume a direct link between available tracking and decoding window margin (RTE margin) and error rate. A system optimized to provide the best compromise provides the best possible performance without going to a higher-order response. It can therefore be termed a best-fit optimized system. The HP 9144A PCL is this type of system.

A best-fit optimized PCL will provide an absolute minimum error rate when subjected to the entire range of expected conditions. This loop, however, will not exhibit the maximum possible average RTE margin and will not have the maximum possible instantaneous RTE margin

during most typical conditions. This is not undesirable; although the margin is reduced during noncritical conditions, it is still reliably maintained well above necessary levels. The reserve capability needed for worst-case conditions (relatively infrequent surges in speed and/or phase shift) causes minor overshoot during less strenuous typical conditions, resulting in reduced average and instantaneous RTE margin. Higher-order loops are more effective in allowing such trade-offs, resulting in improved error rates. A loop optimized to typical conditions using one of the common procedures noted above will not perform as well in any data recovery environment (especially cartridge tape drives) exposed to transient noise sources, head and media imperfections, and speed variations.

Best-fit PCL optimization requires an accurate evaluation of the data modulation during readback. The most critical element of this process is the careful capture and analysis of transient worst-case conditions during all operating environments. This information is used in the first phase of the optimization by applying theoretical analysis and de-




Fig. 4. This diagram illustrates the three possible phase-controlled loop (PCL) operating points. The two sources of data modulation are phase modulation (abscissa), and frequency modulation (ordinate). Increased levels of phase modulation (noise and peak shift) reduce the amount of tolerable frequency modulation (speed variation). This is shown by the PCL operation-limits line. Any modulation falling to the right of the line or above it causes tracking failure. Line 1 represents a typical set point: relatively large margin (distance from modulation area to limits line) exists for typical operating conditions, and failure occurs during transient conditions of large speed variation. In general, increasing the PCL's bandwidth will rotate the line clockwise (line 2). This increases the margin available for transients (but not enough in this case) while reducing the margin for typical conditions. By using a third-order loop (rather than second) and optimizing both bandwidth and phase margin (damping) one can achieve limits-line translation with little or no rotation, thereby improving overall margin (line 3). The margin provided for typical modulation is set larger than that provided for rare events of large modulation. This is essentially a "proration" of margin, which provides the best overall error rate.

sign techniques to arrive at important parameter values such as bandwidth and stability factors (see Fig. 4). These transient conditions are then simulated by using a special signal source which overlays controlled phase and frequency modulation onto a predetermined data pattern. These simulated conditions are used in the second phase of optimization by observing the instantaneous RTE margin during minor parameter adjustments.

During the HP 9144A development, a PCL test system was designed to aid the best-fit optimization process. This test system provides two separate functions: 1) modulated data simulation—combining variable data patterns with controlled simultaneous frequency and phase modulation—and 2) data capture and analysis—high-resolution data gathering and automated analysis of tracking performance and RTE margin. The modulated data generator allows simulated phase modulation from random noise (asynchronous) and from pattern-dependent peak shift (synchronous) to be superimposed onto frequency modulation, thereby providing a realistic, controllable, comprehensive test data source.

Other noteworthy contributions to data integrity are included in the HP 9144A's data separation functions. No adjustments of the circuitry are possible or necessary during the production process and low-tolerance components can be used. Both factors contribute to the low cost and ease of manufacture of the product. All of the timing-critical signal paths are carefully balanced and integrated into a single LSI chip designed for the HP 9144A, providing significantly improved margin and reliability and lower cost.

The phase/frequency detector designed into the HP 9144A's PCL is an improved design that eliminates synchronization boundary problems commonly found in classical sequential phase detectors. When switching from one signal (clock) to the other (data), the phase of the new data source is unknown and can often cause a glitch in the phase synchronization process, which may cause a transient error condition. To avoid this, the asynchronous state machine was carefully analyzed to identify and resolve any possible state transitions that could cause race conditions or hung states.

Error Correction

While the DC600 cartridge is capable of storing data without significant degradation for several years, data loss caused by dust particles, stray magnetic fields, and extremely long storage intervals is possible. The HP 9144A protects the user from such losses through the fourth level of data integrity assurance: post-retrieval error correction. Large portions of the recorded data—up to 512 contiguous




Fig. 5. Exclusive-OR frame generation during data recording.




Fig. 6. Exclusive-OR error correction of defective media. Data can be recovered when up to two complete adjacent frames are in error. Correction is done on the fly; there is no need to reposition the tape and retry. Retry algorithms are used only as a last attempt to recover data.

bytes—can be destroyed without inducing an unrecoverable data error. This feat is made possible by the use of redundant data embedded within the user's data. The prevailing cause of transient errors in a cartridge tape drive is foreign particles, which lift the tape off the head surface, causing loss of a large number of contiguous bits. The HP 9144A's exclusive-OR error correction is especially suited to correcting such errors.

Error correction also serves to extend motor and cartridge reliability by minimizing the need for tape repositioning. Without error correction, each error-causing dust particle or media defect would initiate a reposition and retry operation over the data area, reducing motor and cartridge life. Starts and stops (two are required for each repositioning) are the primary cause of motor wear.

User data is written onto the tape in the form of 1024-byte blocks. Each block is segmented into six 256-byte frames. The first four frames contain the user data while the last two frames contain redundant data (see Fig. 5). The redundant frames are generated concurrently with the write process by executing a bit-serial exclusive-OR operation on the previous frames: frame 5 is frame 1 combined with frame 3, and frame 6 is frame 2 combined with frame 4.

Defective data areas of up to two consecutive frames (512 bytes) can be successfully corrected by virtue of the previously recorded redundancy. Fig. 6 illustrates the process for correcting a defect that overlaps into frames 2 and 3. Frame 2 is reconstructed by performing an exclusive-OR operation on frames 4 and 6. Similarly, frame 3 is reconstructed from 1 and 5.

In general, defects are correctable that affect up to one odd-numbered frame and up to one even-numbered frame. For example, a block is recoverable if defective in frames 1 and 6 if the remaining frames are error-free.

Media Monitor

The fifth and final level of protection against data loss is furnished by the automated media monitor. A built-in use log automatically alerts the user of worn media by flashing a light when it is time to replace a used tape. This allows replacement of the media to avoid losing data, yet enables full use of the expected lifetime of the cartridge.

Controlling the Head/Tape Interface

by Walter L. Auyer, Charles H. McConica, David J. Schmeling, and Mark E. Wanger

ONE OF THE MOST CRITICAL AREAS in the HP 9144A, as in most magnetic storage devices, is the head-to-media (tape) interface. A considerable amount of engineering effort was spent both in examining and in the subsequent control of this interface to ensure data integrity as well as unit-to-unit interchangeability.

For optimum performance, the tape must be in intimate contact with the read/write head and accurately aligned to it. The most important alignment parameters to control are those that influence azimuth angle, tape flying height, and off-track error. The azimuth angle (head-gap-to-tape-path orthogonality) affects phase distortion and signal amplitude; both are important in data recovery. The tape flying height (distance between the tape and the head) is a function of tape tension, tape wrap angle, head penetration, and head geometry; this also affects phase distortion and signal amplitude. Off-track error (head gap position relative to the written track) can introduce erroneous data sensing (from adjacent tracks or other residual data), resulting in unwanted noise, which causes phase distortion. In practice, off-track error is more a function of the step motor positioning system than the head alignment process.

Control of the head/tape interface involves four major areas:

- Head mounting
- The head actuation system
- Head design
- Cartridge referencing.

Head Mounting

In the head mounting process the read/write head is aligned and mounted in the drive. The six degrees of freedom (three angles and three displacements) that affect signal amplitude and phase distortion must each be addressed




Fig. 1. HP 9144A read/write head, showing the six degrees of head mounting freedom and the read and write gaps for the two directions of tape travel.

(see Fig. 1). There are two schools of thought regarding head mounting precision. One favors creating a precision interchangeable mechanical assembly, including the head. The other favors mounting the head relative to the drive volume, reducing sensitivity to manufacturing tolerances, but sacrificing interchangeability. The HP 9144A design team, after carefully considering both options, chose the second because 1) the head has a lifetime ceramic wear coating which virtually eliminates the need for field replacement, and 2) the mechanical parts require less precision and therefore cost less. Since the mechanical parts do not require a high degree of precision, the design team was




Fig. 2. Preliminary head mount fixture.

able to concentrate its effort on the head mounting tools and procedures to make the process fast and precise.

The head mounting process involves two steps: preliminary head mount and azimuth adjustment. The head is first aligned and clamped to a precision fixture, which simulates a data cartridge (see Fig. 2). The fixture is then inserted into the drive, which clamps and references it to the drive mechanism. This establishes five of the six degrees of freedom (all but azimuth). The head is then tacked with adhesive to a bracket, which in turn is held to the drive by a small, clothespin-like clamping fixture. The preliminary head mount fixture is then removed and the drive is ready for azimuth adjustment. The small clamping fixture holds the head and bracket assembly tightly against the head carrier, yet allows rotation to adjust the azimuth angle.

Azimuth adjustment for any type of magnetic storage device can be tedious and costly, calling for a high degree of operator skill and a large investment in equipment. The HP 9144A design team made a significant contribution in this area by designing a process that is fast, inexpensive, and requires no special operator skills. The azimuth adjustment process takes advantage of the head design and the media format. The head has two pairs of gaps (total of 4) that allow bidirectional read/write operations (see Fig. 1). The object of the azimuth adjustment is to set the in-line head gaps perpendicular to the direction of tape travel.

The data cartridges are preformatted with keys written across the width of the tape and spaced 1.7 inches apart, which functionally indicate the physical address of the data. By reading the keys with one vertical pair of in-line gaps (one read gap and one write gap) the time delay between the arrival of a key at each gap can be measured. An azimuth adjustment test station was designed using two read/write boards and a servo control board from the HP 9144A and an HP 9000 Model 216 Computer to control the process. A tape is inserted into the drive and run back and forth.




Fig. 3. Azimuth angle calculation.




Fig. 4. Azimuth adjustment fixture.

As the keys cross the head, the time delay between the in-line gaps is measured and stored. A simple algorithm is used to calculate the azimuth angle from the time delay, the tape velocity, and the in-line gap separation distance (see Fig. 3).

Initially, a rotation fixture is clamped on the head (see Fig. 4). As described above, the data from a test cartridge is used to compute the existing azimuth angle. Computer-controlled corrections direct the head actuation system to step the head up or down an appropriate number of steps, thus rotating the head relative to the tape path. The final azimuth correction is set by factoring in the statistical distribution of all the azimuth readings for both pairs of in-line gaps for both directions of tape travel. Finally, a correction is made to compensate for the azimuth of the test cartridge. When the head azimuth angle is adjusted to the required specifications, the head is tacked in place with a fast-curing adhesive and all the fixtures are removed. The head is now precisely mounted relative to the drive.

Next, the head is checked for magnetic and electrical parameters. If it meets the specifications, the head and bracket are bonded to the head carrier using a high-strength adhesive.

The results of this head mounting process have been impressive. The whole mounting procedure, which took one hour during the early project stages, can be completed by production personnel in only six minutes. The azimuth angle specification is ± 6 minutes, but the process is yielding a majority of heads mounted in a ± 2 -minute range. No part of the process requires operator judgment in setting or aligning the head. Azimuth angle, zenith angle, head penetration, and yaw are measured on each drive and recorded automatically in a data base. Control charts are used to monitor the process and maintain its consistency.

Head Actuation System

The function of the head actuation system is to place the read and write gaps accurately and repeatably on the tracks on the tape and perform such functions as edge finding and unlocking the cartridge. The off-track error is directly related to the accuracy that can be maintained in the working range of the head actuation system. A four-phase dc step motor and a lead screw result in a travel of 0.02 mm per step. The working range is 8 mm (400 steps), although the overall range is greater to allow for unlocking the cartridge. The positional accuracy over the working range is within ± 1 step ($\pm 0.25\%$). Hysteresis in the system is compensated by using an antibacklash spring in conjunction with a routine in the servo electronics that causes the head to approach all targets from the same direction.

Head Design

The mechanical design of the head addresses two aspects of the head/tape interface, one being off-track alignment and the other being head/tape contact or tribology. Both issues are important for achieving a reliable signal from the media.

Tracks are placed on the media with an open-loop positioning system. After an initial referencing to the tape

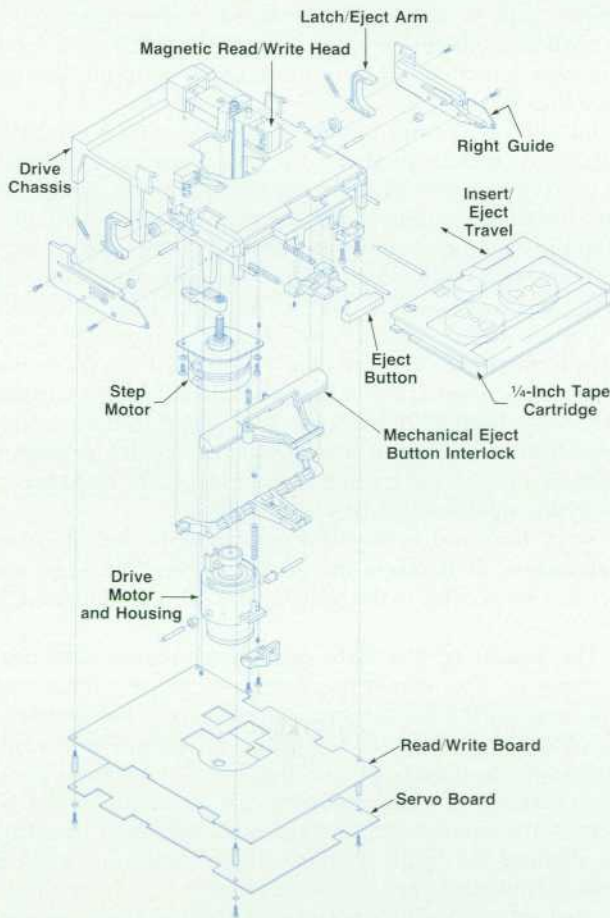


Fig. 5. Exploded view of the HP 9144A.




Fig. 6. DC600 Tape Cartridge.

edge, any error in positioning the head or vertical tape movement within the cartridge (tape wander) leads to track misalignment and possible loss of data. The design of the head gap geometries is aimed at providing a margin for positioning error. Current head technology precludes having a trim erase geometry* with read after write capability, so a wide-write, narrow-read design is specified. The 0.013-in write gaps are 0.004 in wider than the 0.009-in read gaps. The result is that the read gap positioning can vary ± 0.002 in without any signal degradation. The width of the write gaps is limited by the track density requirements and subsequent data overlap.

The design goal for the head in relation to the head/tape interface was to supply an extremely reliable data signal over the range of cartridge variations while not requiring an extremely accurate head mounting process. Head contour design must consider variations in tape tension and wrap angle, which directly affect tape flying height. The gap configuration compounds the problem since there are four gaps, which must be held in intimate contact with the tape, which is traveling at 90 inches per second. To provide some perspective, the Van der Waal equation for signal loss from spacing (flying height) is $55 \text{ d}/\lambda \text{ dB}$. At 10,000 flux reversals per inch, $\lambda = 100 \mu\text{in}$ and a flying height of $5.5 \mu\text{in}$ (about the surface roughness of the media) results in a 3-dB loss. The final design uses a double-bump configuration with a consistent curvature on each bump. The curvature allows for variations in wrap angle while still maintaining good head/tape spacing.

We chose to use a ceramic wear layer to enhance head life. The trade-off was difficulty in working with materials of different hardness in the machining and profiling process. To get the necessary profile required working closely with the head vendor on the design and the manufacturing process.

The result of careful design and extensive testing is a head that gives a very reliable signal that is insensitive to variations in cartridges, tape tension, and tape wrap angle.

Cartridge Referencing

Equally important in controlling the head/tape interface is the relative position of the tape cartridge (and thus the

*In a trim erase geometry, two erase heads, one on each side of the data track, trim the track to its proper size after the data is written by the write head. This geometry is common in flexible disc drives.

tape itself) to the HP 9144A tape drive mechanism. Esthetic considerations (insertion and ejection forces and travel, "feel," appearance, etc.) of the mechanism also were challenging aspects of the design.

The drive mechanism, primarily a mix of motors, plastic hardware and die castings (see Fig. 5), accepts and captures an inserted cartridge, positions it relative to the read/write head, stabilizes it during dynamic drive functions (tape movement relative to the head and vice versa), and upon request, unlocks the cartridge for the user to eject.

ANSI Specification X3.55-1977 defines the parameters for the ¼-inch tape and 0.67 × 4.00 × 6.00-in tape cartridge used in the HP 9144A. The defined cartridge reference datums (cartridge position planes) are the forward (or leading) edge and the top surface of the 0.10-in-thick aluminum baseplate (see Fig. 6). Internally, the tape guide pins are positioned relative to these datums, thus establishing the tape path relative to the cartridge base plate.

The user inserts the cartridge until it firmly snaps into place. Insertion resistance forces (friction, motor capstan reaction, latch/eject arm reaction) build to approximately

3 lb. To smooth and minimize this resistance, the guides are made of Teflon™-filled plastic (not glass-filled—glass causes friction) and incorporate guide rollers for aft cartridge stabilization. During the last millimeter of insertion travel, the dual-purpose latch/eject arms are released. This results in 3.5 lb upward and forward forces applied to the cartridge, pulling the cartridge reference datums audibly, firmly, positively, and repeatably against their corresponding mechanism references. A stable head/tape interface is thus established.

Another feature of the mechanism is a mechanical eject button interlock, which prohibits the user from inadvertently ejecting the cartridge while the head is in an active, or up, position. This interlock decouples the eject button from the head/tape interface, thus preventing accidental bumping of the eject button from disturbing any read/write operations that may be in progress. Cartridge ejection must be preceded by a request for unload (a pushbutton on the front panel) which initiates a down-load sequence, subsequent head stepping down, and release of the mechanical eject button interlock.

Software Methodology Preserves Consistency and Creativity

by Mark L. Gembarowski

THE PROJECT TEAM that developed the HP 9144A Tape Drive also developed a software methodology that is now the way in which all programming is done at HP's Greeley Division. The methodology provides designers with a consistent, common approach to software design without limiting their creativity and coding style. This software methodology addresses two major questions:

- What can we do about schedule?
- How "good" is the code?

The software methodology does not generate code until the final step of the design. Yet, we can maintain or even beat schedule because the code generated is very easy to debug, extremely flexible (features can be easily added), and extendable. Many of the subroutines written for the HP 9144A are being used on new products under development because of the modularity of the code that the design methodology produces.

The software design methodology embodies a number of design ideas seen in the current literature.¹⁻³ They include structure charts (hierarchy charts), structured programming, top-down design, structured walkthroughs, and structured analysis. We took the best attributes of these methods and combined them into a methodology that gives the designer a series of design steps that offer additional insight into the design at each subsequent phase. The

methodology generates design documentation as a product of each design step, and provides management with a series of design checkpoints for measuring progress.

Four Phases

The specific steps of the software design methodology are the definition phase, the hierarchy phase, the input-process-output phase, and the coding phase.

Definition Phase. The output of this phase is a document that details what the code must do, what hardware environment it resides in and must interface to, and a list of reference documents. This design step may include system architecture design, some structured analysis, and some testing of existing products one is trying to improve.

Hierarchy Phase. Once you know what you need to do (definition phase), you need to know how to do it. A hierarchy is a graphical representation of how the code accomplishes all of the functions listed in the definition. It shows individual modules (subroutines) and the interface between the modules. An initial hierarchy is created and then refined to improve the functionality of individual modules while reducing the interdependence between modules. This design step uses the methods of hierarchical decomposition found in much of the current software design literature. Fig. 1 shows a sample hierarchy chart.

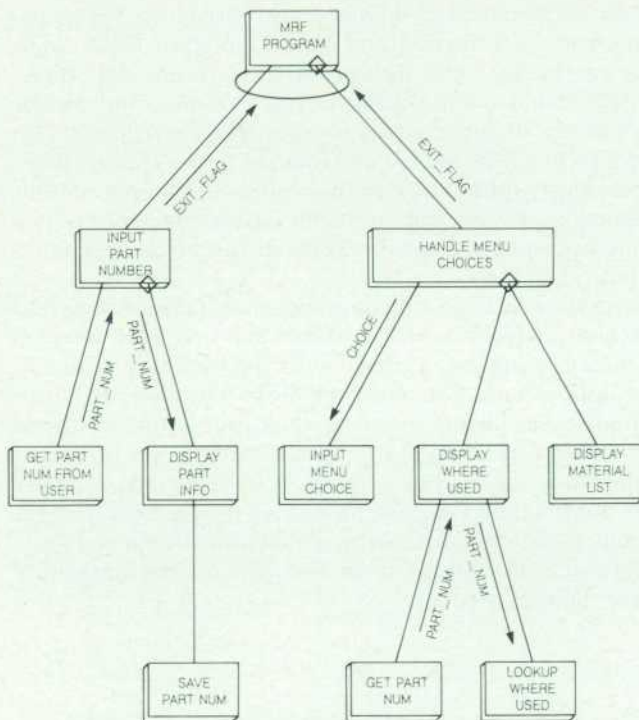


Fig. 1. A portion of a hierarchy chart. Such charts tell how the code accomplishes the functions required. Each box with a name in it is a code module (subroutine). The arrows on the lines connecting the modules indicate parameters passed between modules. The elongated circle below the top box indicates that the subordinate calls are repeated. The diamond symbols indicate conditional calls or a selection from a group of choices.

Input-Process-Output (IPO) Phase. Now that we have seen how we will do each function, we again develop what we will do, but this time the focus is on the individual modules rather than on the entire task to be accomplished. For each module in the hierarchy, we write a document (typically one-half page in length—it serves as an ideal header for the code) that shows the inputs to the module and their type (Boolean, real, integer, etc.), the process within the module (using inputs, subordinate module calls, and how the outputs are created), and the outputs and their type. For complicated algorithms, we expand the IPO to include a structured English algorithm to help the code writer understand exactly what the module must do. We use constructs such as IF-THEN-ELSE, REPEAT-UNTIL, WHILE-DO, and CASE and indentation to show the structure of the complicated algorithm. Once the input-process-output phase is completed, the code for any module can be written by any member of the team.

Coding Phase. Notice that no code has been generated before the design is finished. There are no constraints regarding what language the designers write in. They may be limited by design tools, but not by the methodology. When the coding begins, typically certain functions will need to be done first, most likely to help show functionality for a checkpoint meeting. Since the design is finished, portions of the coding can be set aside and done later, or new project members can do it when they join the team.

References

1. G.J. Myers, *Reliable Software through Composite Design*, Van Nostrand Reinhold, 1979.
2. G.J. Myers, *Composite/Structured Design*, Van Nostrand Reinhold, 1978.
3. E. Yourdon and L.L. Constantine, *Structured Design*, Prentice-Hall, 1979.

HEWLETT-PACKARD JOURNAL

March 1985 Volume 36 • Number 3
 Technical Information from the Laboratories of
 Hewlett-Packard Company

Hewlett-Packard Company, 3000 Hanover Street
 Palo Alto, California 94304 U.S.A.

Hewlett-Packard Central Mailing Department
 Van Heuven Goedhartlaan 121

1181 KK Amstelveen, The Netherlands

Yokogawa-Hewlett-Packard Ltd., Suginami-Ku Tokyo 168 Japan

Hewlett-Packard (Canada) Ltd.

6877 Goreway Drive, Mississauga, Ontario L4V 1M8 Canada

