


JULY 1974

HEWLETT-PACKARD JOURNAL



Powerful Data Base Management Systems for Small Computers

For the first time, sophisticated data management is available to the user of computers as small as the HP 2100 and HP 3000.

by Richard E. McIntire

COMPUTER INFORMATION SYSTEMS harness the processing power of the computer to collect and organize data and make it easily accessible to the businessman, scientist, or other user who needs it. A computer information system may be thought of as a set of applications programs using common data bases through a data management system, as shown in Fig. 1. The data management system is composed of two separate software entities: the operating system and the data base management system.

In the past it has been generally accepted that computerized information systems and, similarly, data base management systems, required large-scale computers. But as computer performance has expanded, the conventional concept of a data base management system has changed. Increasingly evident is a shift away from large central processors designed to serve all needs and all users. The demand now is not necessarily a bigger or faster computer, but a less costly, more flexible system tailored to serve definite data base requirements and process the total expected workload at a lower cost.

Small-Computer Data Management

Hewlett-Packard's new data base management system, IMAGE, and its companion data base inquiry facility, QUERY, for the first time make sophisticated information management available to the user of small, low-cost computer systems based on the HP 2100 and HP 3000 Computers. The design objective for these new software systems was to make common data bases and data management services accessible to users ranging from those who know nothing about the system and want to use it without learning complex programming languages to those who understand the system well and want to manipulate its inner workings to their advantage.

IMAGE is designed for the internal user, that is, the

information systems manager, the data base manager, the systems analyst, the computer system specialist, and the programmer—the user who is interested in the computer, its associated input/output devices, computer programs, and the influence of these components on data entry, organization, and retrieval. QUERY is designed for the external user, the non-



Cover: The images on the head represent information – credit card records, inventory records, student and course data, and so on – information that must be organized and made readily available to those who need it (and protected from those

who don't). IMAGE, a sophisticated new data base management system for HP 2100 and HP 3000 Computers, has information management capabilities formerly available only in large computer systems.

In this Issue:

Powerful Data Base Management System for Small Computers, by Richard E. McIntire	page 2
Quality Frequency Counters Designed for Minimum Cost, by Lewis W. Masters and Warren J. O'Buch	page 11
A Versatile Bipolar Power Supply/Amplifier for Lab and Systems Use, by Santo Pecchio	page 15
An Automatic Exposure Control for a Lab-Bench X-Ray Camera, by John L. Brewster	page 20




Fig. 1. General concept of a computer information system.

specialist, whose needs are to retrieve, analyze, and report information to support his function and decision processes.

IMAGE consists of a set of programs that create and maintain complex data structures known as data bases, and a set of library procedures that enable users to access, modify, and report on the data content of the data bases. There are two versions: IMAGE/3000 for HP 3000 Computers and IMAGE/2000 for systems based on HP 2100 Computers. IMAGE/2000, a subset of IMAGE/3000, has many but not all of the features of IMAGE/3000. IMAGE/2000 operates under the control of the disc operating system for HP 2100 Computers (DOS-III) and IMAGE/3000 is executed under the control of MPE/3000, the operating system for HP 3000 Computers.

IMAGE/3000 Features

IMAGE provides powerful software tools that help the data base manager define and create a data base tailored to his requirements. It has a network data structure that allows cross-referenced access to collections of data down to the smallest unit.

The data base manager defines data sets and their interrelationships just once. Thereafter, applications programmers can search, retrieve, or update their data bases from host-language programs without concern for the details of accessing the data base. Host languages can be FORTRAN, COBOL, or SPL (the HP 3000 Systems Programming Language).

IMAGE provides facilities for combining files from many applications into one data base, so sets of data

that occur in two or more old files are consolidated and need be stored only once.

A flexible security scheme lets the data base designer control access to any subset of the data base, down to the smallest unit. The user must include a password or level word identifying his access. Reading and updating are treated as two separate operations, and the system checks each data transfer request to determine the user's access before allowing the read or update operation.

IMAGE Subsystems

IMAGE consists of three basic elements (Fig. 2): a data base definition subsystem (DBDS), a data base management subsystem (DBMS), and a data base utility subsystem (DBUS).^{*} The data base manager uses DBDS to define the data base and DBUS to create and maintain the data base. The applications programmer, the principal user of IMAGE, writes procedural programs using a host language and the data base management language, DBML, which operates on the data base using DBMS.

DBDS is independent of the applications programs. This system allows the data base manager to define all aspects of data base organization. His data base definition is called a *schema*. Using the data base definition language, the data base manager defines data items, security levels, and relationships and mappings between data sets.

^{*} DBDS corresponds to the CODASYL data base task group's data base description statements and schema processor. DBMS corresponds to CODASYL's data base manipulation language and library routines. DBUS corresponds to various CODASYL utility routines.





Fig. 2. IMAGE subsystems. IMAGE/2000, for HP 2100-Series Computers, is a subset of IMAGE/3000, which runs on an HP 3000 Computer System.

DBMS provides the means for applications programmers to access an IMAGE data base. DBMS is a set of stored library routines invoked by CALL statements in host-language applications programs. DBMS serves as the interface between the data base and the applications programs, and either can change without affecting the other. DBML is a non-procedural language. It relies on the host language to provide a framework and the capabilities required to manipulate the data. Thus, the application programmer uses the full data processing power of his host language and leaves data base structure and access activity to the DBMS. DBML is not an inquiry language and does not provide for selection criteria in the form of Boolean expressions; however, QUERY does allow such selection criteria.

The third IMAGE subsystem, DBUS, consists of a set of support routines that run as stand-alone programs. DBUS permits building, dumping, and restoring data bases and assists in their restructuring.

Preparing a Data Base

Four steps are required to prepare an IMAGE/3000 data base (see Fig. 3). The first step is for the data base manager to use the data base definition language to define the data base structure. Fig. 4 is an example of a typical data base definition, or schema. Second, a DBDS program is employed to process the schema, the result being a disc-resident table describing the data base. This table is called a root file. Third, a DBUS program is employed to build the empty disc files that constitute the framework of the data base; each such disc file is called a data set. Fourth, data is entered into the previously created data sets by means



```

BEGIN DATA BASE SCHOOL

LEVELS:
  4 CLASS#
  5 TEACHID#
  6 STUDENT#
  8 TEACHER#
  10 CLASSUG#
  11 #TEACHID#
  12 STUDENT#
  14 TEACH#
  15 ZADMIN#

ITEMS:
  SCHL-TEACH, X12(8,15)
  SCHL-CRSE, X6(4,15)
  NO-SEC, X2(4,14)
  DESCR, X14(4,14)
  CREDITS, X2(4,14)
  SEATS, P4(4,10)
  FEES, P4(4,10)
  MEET-RGNTS, X6(4,10)
  SEMESTER, X2(4,12)
  ROOM-NO, X6(4,14)
  SCHL-CRSE-ID, X8(4,15)
  SCHL-STUDENT, X8(6,15)
  COM-CODE, X2(6,12)
  GRADE-LEVEL, X2(4,10)
  ENROLL, X2(6,12)
  MONTH, X2(6,12)
  DAY, X2(6,12)
  ABSENT, X2(6,12)
  COMMENTS, X10(8,15)
  SCORES, SP4(6,12)

SETS:

NAME: TEACH-MSTR, A(8,14)
ENTRY: SCHL-TEACH(1)
CAPACITY: 3000

NAME: SECTION-MSTR, A(5,12)
ENTRY: SCHL-CRSE-ID(2)
CAPACITY: 10000

NAME: STUDENT-MSTR, MANUAL(6,12)
ENTRY: SCHL-STUDENT(1), ROOM-NO, COM-CODE, GRADE-LEVEL
CAPACITY: 30000

NAME: COURSE-MSTR, MANUAL(4,14)
ENTRY: SCHL-CRSE(1), NO-SEC, DESCR, CREDITS, SEATS, FEES,
MEET-RGNTS, SEMESTER
CAPACITY: 10000

NAME: STUDENT-TEST, DETAIL(5,12)
ENTRY: SCHL-STUDENT(STUDENT-MSTR),
SCHL-CRSE-ID(SECTION-MSTR), COMMENTS, SCORES
CAPACITY: 100000

NAME: COURSE-SEC, DETAIL(5,11)
ENTRY: SCHL-TEACH(TEACH-MSTR(SCHL-CRSE)),
SCHL-CRSE(COURSE-MSTR),
SCHL-CRSE-ID(SECTION-MSTR),
ENROLL, ABSENT, MONTH, DAY
CAPACITY: 300000

END.
  
```

Summary of Processed Schema

DATA SET NAME	TYPE	LEVEL	FLD	PT	ENTR	REC	CAPACITY	BLK	BLK	DISC
		R	W	CNT	CT	LGTH		FAC	LGTH	SPACE
TEACH-MSTR	A	8	14	1	1	8	16	300	31	498
SECTION-MSTR	A	5	12	1	2	4	19	1000	20	382
STUDENT-MSTR	M	4	12	4	1	13	23	3000	22	508
COURSE-MSTR	M	4	14	8	1	17	27	1000	14	379
STUDENT-TEST	D	5	12	4	2	18	26	10013	19	496
COURSE-SEC	D	5	11	7	3	17	29	30004	13	378
TOTAL DISC SECTORS INCLUDING ROOT: 10015										6927

```

NUMBER OF ERROR MESSAGES: 0
HIGHEST LEVEL WORD: 15
ITEM NAME COUNT: 20
DATA SET COUNT: 6
ROOT LENGTH: 676
BUFFER LENGTH: 508
TRAILER LENGTH: 256
ROOT FILE SCHOOL CREATED.
  
```

Fig. 4. An example of a schema that defines an IMAGE/3000 data base.

of QUERY or applications programs.

When an applications program uses the data base, it must enter its data requirements into a data buffer, which defines the elements of the data base the program needs. When the program wants to access the data base, it communicates with IMAGE/3000 using DBML to access the DBMS subsystem. Using the program's data buffer and the data base root file to locate the desired data, DBMS gets the data for the applications program. If updating is required, DBMS will take the new data from the data buffer and put it into the data base.

IMAGE/3000 Data Base Organization

Within an IMAGE/3000 data base are three basic structures: data items, data entries, and data sets.

The *data item* is the smallest accessible data element. Each data item is a value and is referenced by a *data item name*, which is a character string defined in the schema by the data base manager. Usually, many data item values are referenced by the same data item name.

Data Item Name from Schema	Data Item Values
NAME	SMITH, JONES, JOHNSON, GREEN, MEADE, DILLION
CITY	SAN JOSE, DENVER, PRESCOTT, AJO, GILROY, TRACY
STATE	CALIFORNIA, ARIZONA, COLORADO

A *data entry* is an ordered collection of related data items and is defined by an ordered listing of the data item names in the schema. Data entries are stored in physical locations on a direct-access storage device, such as a disc.

Data Entry Definition from Schema	NAME	CITY	STATE	DEGREE	SEX
R_i	SMITH	PRESCOTT	ARIZONA	BA	M
Data R_{i+1}	JONES	DENVER	COLORADO	MBA	M
Entries R_{i+2}	JOHNSON	SAN JOSE	CALIFORNIA	MS	F
R_{i+3}	GREEN	TRACY	CALIFORNIA	BA	F

A *data set* is a collection of data entries sharing a common definition. All data entries within a data set are of the same length (the maximum length in IMAGE/3000 is 4094 bytes). A *data set name*, a character string defined in the schema, references any or all of the data entries of a data set. The number of data entries in a data set is limited by available disc space.

Data Set Name: PERSONNEL

Data Entry Definition from

Schema: NAME, CITY, STATE, DEGREE, SEX

Smith Prescott Arizona BA M

Jones Denver Colorado MBA M

Johnson San Jose California MS F

Green Tracy California BA F

Physical
Storage
Locations
on Disc

A *data base* is a named collection of related data sets. It is referenced by a *data base name*.

Search Items and Chains

Data entries in a data set can be referenced by one or more data items known as search items. Search

items of a data set are specified in the schema. Data entries of a data set are linked together in subsets based on the values of their search items. For example, the data item STATE might be specified as one of the search items for a data set. A general reference to all data entries of the data set having the same value of STATE is then possible. If all fifty states of the United States of America are represented, the data set would be logically divided into fifty subsets, with all the entries of each subset containing the same value of STATE.

Each data entry of a data set is distinguished by an entry number. An entry number is a unique integer between 1 and N, where N is the data set's capacity or the total number of available data entry storage locations, as defined in the schema. Each storage location is initially empty. Whenever a new data entry is added to a data set, it is assigned one of the unused entry numbers and written in the corresponding storage location. Entry numbers are the means by which data entries with like search item values are linked together.

IMAGE/3000 Chains

For each search item, pointers are maintained by IMAGE with each data entry, along with the data item

Why Data Base Management Systems?

Historically, the cost/performance ratio of computer hardware has improved, while the cost of programming has increased. Just as a major reason for the development of high-level languages such as FORTRAN and COBOL was to reduce coding time and cost, the primary objective of information systems development is to reduce the time and cost of writing programs to store and retrieve information in a computer system.

In data management, there are certain functions that are repetitive, time-consuming, and error-prone, and if each programmer performed these functions every time he created a new program it would be extremely inefficient and uneconomical. For example, in an earlier approach to data processing using a batch mode of operation, all the data for a particular application had to accompany the applications program. Thus data in the personnel file might be repeated in the skills file, the payroll file, and the medical file. When these files are combined into a common data base, redundant data is eliminated, storage costs are lower, and the data is internally consistent, requiring only one standard procedure for updating or modification. Also, data base management can be performed independently, freeing the applications programs from this task.

Data base management concepts are an evolution of earlier EDP techniques and not a radical new method. The innovation in the data base approach is that the definition and control of the data base are independent of the applications. The common data base of logically connected files or items of data is then accessible to all programs of the proper security clearance by means of special software that permits more efficient data processing and therefore easier systems development and lower programming costs.




Fig. 5. IMAGE/3000 data sets are of two types, master and detail. An important purpose of master data sets is to serve as indexes to detail data set chains.

values for that entry. These pointers contain the entry numbers of the preceding and succeeding data entries within the data set that have the same search item value. All data entries having the same search item value are linked together in this way, and are referred to collectively as a chain.

In a set of employees, for example, engineers form a subset of the employee set, based on the search item JOB-TITLE. All engineers would be linked together to form an engineer chain. Thus, members of a data chain have in common the value of a specified search item.

A chain may be in sorted order if each new data entry is inserted into the chain at a point determined by the value of a specified data item called a sort item. Sort items are defined in the schema.

IMAGE/3000 Data Set Types

There are two types of data sets in IMAGE/3000: master data sets and detail data sets (Fig. 5). Detail data sets contain "line item" information. For example, in the detail data set PERSONNEL, each person's location, educational experience, and similar information is stored. An important purpose of master data sets is to serve as indexes to detail data set chains.

The data entries of a master data set have just one search item and unique search item values. Thus in a

master data set with the search item STATE, there are at most fifty entries, one for each of the United States that appears in a related detail data set.

Data entries of a master data set contain pointers to corresponding chains in related detail data sets. For example, if the specified search item is STATE, the "New York" entry in the master data set will contain the entry numbers of the beginning and end of the "New York" chain in each related detail data set. The "New York" entry (and all other entries) in the master data set may also contain master information about the state, such as population. This information then does not have to be duplicated with each "New York" entry of the detail data set or sets.

For each search item defined in a detail data set, an existing master data set is specified by name. This association establishes a master-detail data set relationship. A master data set may be related to more than one detail data set, and a detail data set may be related to more than one master data set.

IMAGE/3000 Data Access

All data base operations are accomplished through the facilities of the DBML that interface with a host language such as COBOL, FORTRAN, or SPL. The DBML is structured so that each command consists of a DBMS procedure call followed by a set of parameters.

In general, access to data within a data base is carried out on the data entry level. That is, each call to a DBMS procedure accesses some or all of the data items within a data entry. The table below presents an outline of the parameter set.

```
CALL <DBMS procedure> USING
  <base><dset><mode><status>[<list><buffer><arg>]
Where: <base>      is the data base of interest
       <dset>      is the data set of interest
       <mode>      various modes within each DBMS
                   procedure
       <status>     status area containing the results
                   of the execution of the DBMS pro-
                   cedure in the application program.
       <list>       is a list of data item names of
                   interest.
       <buffer>     is the address of a buffer in the
                   user's data area
       <arg>        is the search item of interest
       [ ]          indicates optional parameters
```

Some of the functions that can be accomplished through use of the DBMS procedures include adding a new data entry to a data set, deleting a data entry from a data set, reading some or all of the data items of a data entry, and changing the values of items of a data entry.

Master Data Set Access

Data in a master data set may be accessed in serial, directed, calculated, or chained fashion.

In serial access, DBMS accesses the data entry whose entry number is one greater than the last until a data entry is located and read or until an end of file is encountered. When a data set is accessed by a program for the first time, a search begins at entry number one. Reverse serial access is also possible.

Directed access is accomplished by an applications program's supplying an entry number. If present, the data entry with the specified entry number is read. If no such data entry exists, the program is notified by an exceptional condition return.

Calculated access is based on a search item value. It involves mapping the applications-program-supplied item value into a primary entry number by means of what is known as a key transformation. The data entry at that location is then accessed to determine if it contains the matching search item value. The key transformation may map more than one search item value into the same primary entry number. When this occurs the search item values are called synonyms. If the data entry at the location specified by the primary entry number does not contain the desired search item value but does contain a synonym, an exhaustive search of all synonyms with the same primary address is made to locate the desired data entry, if it exists in the data set. To eliminate ambiguity, data entries of a master data set must

have unique search item values. Calculated access is used to retrieve a selected data entry or obtain the chain head pointer of a detail data set chain from the master data entry without an exhaustive search of the entire master data set. Obtaining the chain head pointer is generally done as a prelude to accessing the data entries of a chain in a detail data set.

In chained access of a data entry in a master data set, the applications program may read the next data entry in a synonym chain in either the forward or the backward direction.

Detail Data Set Access

The first data entry in a detail data set is assigned entry number one and subsequent data entries are assigned entry numbers 2, 3, 4 and so on in sequence. However, DBMS keeps track of deleted data entries and always reallocates deleted entry numbers. This modified sequential allocation is called serial allocation; it applies only to detail data sets.

Data entries are logically linked to other similar data entries in the detail data set, as well as to the master data entry to which they belong. Detail data entries may be retrieved directly through chains related to the master data set.

The data entries of a detail data set may be accessed in serial, directed, or chained fashion. Serial and directed access to data entries in a detail data set are identical to that for master data sets. Chained access to data entries is only applicable to detail data sets having one or more search items. When a new data entry is added to such a data set it is linked into the existing chain of data entries whose search item values match that of the new data entry. If the new data entry has more than one search item this linking process is done for each search item.

Chained access to data entries of detail data sets enables rapid access to all data entries having a common search item value but, in itself, does not assist in locating the initial occurrence of a value. Retrieval of detail data entries in a given chain is generally preceded by a calculated access to the corresponding data entry in a master data set to obtain the chain head pointers. Once the application program has located itself on a chain, it may read the next data entry in either the forward or backward direction of the chain.

QUERY/3000 Subsystem

QUERY/3000 is a self-contained subsystem that interfaces with the DBMS of IMAGE/3000. A major development problem was to design a suitable language that would respond to spontaneous and unanticipated inquiry concerning data in an IMAGE/3000 data base. It was critical that the user be given a communication tool that relates to his problems rather than to the underlying programming problems, and that the language be the user's natural lan-

guage or something very close to it.

By entering English-like commands to QUERY/3000, the user can access data in the data base without learning complicated programming languages. The QUERY language, QL, allows the user to specify the information he wants using logical Boolean expressions of key-value pairs. The user may write procedures, have them executed, and, if desired, have them stored for repeated use at a later time.

QUERY/3000 Functions

QUERY/3000 is composed of six functions (see Fig. 6): QUERY language (QL), QUERY interpreter, retrieve subsystem, report writer subsystem, update subsystem, and auxiliary functions.

The QUERY interpreter initially displays a standard HP 3000 program identity message, an initialization message, and a prompt character on the user's terminal. The prompt character tells the user that he has correctly requested QUERY/3000 and the QUERY interpreter is active and awaiting further input. The QUERY interpreter is the switching module of QUERY/3000. It interprets requests from the user, passes control to the appropriate module for further processing, and outputs the results to the user.

The retrieve subsystem does both inter-record processing, in which specified data entries are selected from the data base, and intra-record processing, in which the selection criteria are more restrictive. Data entries satisfying the selection criteria are extracted from the data base and their entry numbers are placed in a selection file. The number of data entries meeting the retrieve criteria is then displayed to the user. The user may now allow the report writer or update subsystems to execute, or may limit the number of data entries retrieved by issuing another FIND command.




Fig. 6. QUERY/3000 functions make it easy for the non-specialist to retrieve and update data in IMAGE/3000 data bases.

QUERY Language		
Utility Commands	Main Commands	Stored Procedure Commands
DEFINE	FIND	CREATE
EXIT	REPORT	DISPLAY
HELP	UPDATE	DESTROY
FORM		ALTER

Fig. 7. QUERY/3000 commands.

The report subsystem provides flexibility in the format of reports. Reports may include page headings, column headings, and page numbers. Data item values may be subtotaled and totaled, and data entries may be sorted by multiple categories.

The update subsystem allows on-line update to the data base. The updated information is usable immediately after updating.

The auxiliary functions allow the user to create procedures for changing specific data item values in a data entry, for locating data entries that qualify according to specified search conditions and for writing reports about data entries that were retrieved. The user can also list, modify, or delete stored procedures.

QUERY-Language Commands

The QL commands, shown in Fig. 7, may be used in either an on-line or a batch environment.

The HELP command is a tutorial aid; it may be used whenever the user wants to know the required syntax for a particular request.

The EXIT command may be used anytime an input response is expected; it causes immediate termination of QUERY/3000 and returns control to MPE/3000. It is the proper method of notifying the system of exiting from QUERY/3000.

The purpose of the FORM command is to display the structure of a data base. When used, it lists the data item names, data set names, and relationships defined in the schema. These names may then be used in other commands.

The DEFINE command is used to inform QUERY/3000 of the data base of interest, the data sets of interest, a SPEC-FILE containing stored procedures, and the output device name. These remain valid for any other QL command, but may be changed at any time. If a required define type is not present in the define table, QUERY/3000 will display an error message and the user must supply the required define type before being allowed to continue.

Data Retrieval

Data entries are retrieved from the data base as

specified in the FIND command. When there is more than one comparison for each data entry, each comparison must be connected logically to the next by AND or OR. The FIND command is of the general form:

FIND <a data item name> <with a specified relation to> <a specified data item value> END
A list of available data item names may be obtained by using the FORM command. The type of comparison to be made is indicated by the relational operators listed below:

RELATIONAL OPERATOR	MEANING
IS IE	=
ISNOT INE	≠
ILT	<
INLT	≤
IGT	>
INGT	≥
IB	$n_1 \leq a \leq n_2$, where a is a data item name defined in the schema, and n_1 and n_2 are specific limits.

When multiple logical connectors are used, ANDs are satisfied first and then ORs. Parentheses are not allowed.

Examples:

1. Retrieve all data entries of persons who are between the ages of 21 and 25, live in the state of Iowa, and have blood type AB.

FIND STATE IS "IOWA" AND AGE IB "21", "25" AND B-T IS "AB" END

2. Retrieve all students majoring in German or Spanish, whose cumulative grade average is greater than 2.9.

FIND MAJOR IS "GERMAN", "SPANISH" AND AVGRADE IGT "2.9" END

Reports

The REPORT command is an extension of the FIND command. Its purpose is to generate a report of the data entries retrieved by specifying either a report procedure, the name of a report procedure stored on a SPEC-FILE, or the keyword ALL, which prints the data item name and data item value for each data item in every data entry retrieved, without any report formatting or data editing.

If the user wants to print the retrieved data on an output device in a formatted report, he may specify a report procedure. There are six statements that enable him to do this. These specify the header information to be printed at the top of each page, the order in which data is to be sorted, the data to be printed in each report column and the way in which it is to be punctuated, and any summaries to be printed of information contained in parts or all of the report.

Fig. 8 shows an example of a QUERY/3000 inquiry, report definition, and report.

```

QUERY/3000 QUERY
NEXT?

DEFINE
DATA=BASE * SCHOOL
LEVEL = ZADWINZ
MODE = 2
DATA-SETS = COURSE-SFC
SPEC-FILE =
OUTPUT = TEWM
OUTPUT =
NEXT?

FIND MONTH IS "MAY" AND DAY IS "4", "7", "8", "9", "5" AND
SCHL-CRSE-ID IS "CHEM1", "ENG2", "HIST5", "MATH2", "MATH3", "SHOP3" END
IN ENTRIES QUALIFIED
NEXT?

REPORT
M1="DAILY", 18
M1="ATTENDANCE", 47, SPACE B1, SPACE A2
M1="WEEK OF MAY 5, 1974", 20
M2="PAGE", 63
M2="PAGE", 66, SPACE A2
M3="DAY", 4
M3="COURSE", 15
M3="TEACHER", 38
M3="ENROLLED", 55
M3="ABSENT", 65, SPACE A2
S1="DAY 5", SCHL-CRSE-ID
D, DAY, 3
D, SCHL-CRSE-ID, 20
D, SCHL-TEACH, 40
D, ENROLL, 52
D, ABSENT, 63
T1=" ", 65, SPACE B2, SPACE A2
END

DAILY ATTENDANCE REGISTER

WEEK OF MAY 5, 1974 PAGE 1

DAY COURSE TEACHER ENROLLED ABSENT
5 CHEM1 BASS 18 2
5 ENG2 JOHNSON 30 5
5 HIST5 CORCORAN 28 2
5 MATH2 WHITE 15 1

6 CHEM1 BASS 18 2
6 MATH3 BROWN 25 3
6 SHOP3 DOLAN 25 4

7 CHEM1 BASS 18 3
7 ENG2 JOHNSON 30 8
7 HIST5 CORCORAN 28 5
7 MATH2 WHITE 15 4

8 CHEM1 BASS 18 0
8 MATH3 BROWN 25 3
8 SHOP3 DOLAN 25 6

9 CHEM1 BASS 18 1
9 ENG2 JOHNSON 30 7
9 HIST5 CORCORAN 28 2
9 MATH2 WHITE 15 0

NEXT?
EXIT

```

Fig. 8. An example of a QUERY/3000 request, response, report definition, and report.

Update Commands

The UPDATE command may be an extension of the FIND command if data items are to be deleted or replaced. The UPDATE command may also be used to add a new data entry. The update procedure allows three different types of update statements: ADD, DELETE, and REPLACE.

The ADD statement is used to add a data entry to the data set. The user need not include all values for every data entry; data items values not included will be declared null by the system. The system will prompt the user with the data item names of the data set when

requesting input. If a data item name is a key or search item, a value of null will not be accepted.

The DELETE statement in the UPDATE command is executed after the FIND command has been executed and the data entries have been retrieved. All retrieved data entries will be deleted from the data base.

To replace specified data item values, the REPLACE statement of the UPDATE command is used. It is executed after a FIND command has been executed and the data entries having data item values to be replaced have been selected. If the REPLACE statement is applied to data entries that do not contain the data item to be replaced, the system will not modify the retrieved data entries.

Stored Procedure Commands

The CREATE command is used to store a FIND, REPORT, or UPDATE command into a SPEC-FILE. FIND, REPORT and UPDATE commands may all be stored in one SPEC-FILE.

The DISPLAY command will display a stored procedure generated with a CREATE command. The keyword LIST will print all the procedure names in the SPEC-FILE.

The ALTER command allows the user to make modifications to a stored procedure in the SPEC-FILE. In-

dividual lines may be deleted, replaced, or inserted.

The DESTROY command causes the destruction of a stored procedure generated with a CREATE command. This command will not destroy any data sets of a data base.

Acknowledgments

The author wishes to acknowledge the following people who were involved in the design and implementation of the IMAGE and QUERY projects for both the HP 2100 and HP 3000 computers, and whose efforts made these projects a reality: Lee Johnson, Fred White, Lee Bollinger, Jon Bale, Ed Estes, Bob Mayer, Bob Brown, Carol Fuquay, Mary Berner, and Phil Taylor. 

References

1. "Feature Analysis of Generalized Data Base Management Systems," Association for Computing Machinery, New York, May 1971.
2. Report of the CODASYL Data Base Task Group, ACM, April 1971.
3. J.K. Lyon, "An Introduction to Data Base Design," John Wiley and Sons, Inc., 1971.
4. D. Lefkowitz, "File Structures for On-Line Systems," Spartan Books, 1969.

SPECIFICATIONS

IMAGE/3000 QUERY/3000

DATA ITEM NAMES PER DATA BASE: 255
DATA SETS PER DATA BASE: 99 (The space occupied by any single data set cannot exceed the capacity of any one disc drive. However, the total data base is limited only by the total available storage.)
CHARACTERS PER ITEM NAME: 16
CHARACTERS PER DATA SET NAME: 16
DATA ITEMS PER DATA ENTRY: 127
MAXIMUM DATA ENTRY SIZE: 4094 bytes
KEYS PER DETAIL DATA SET: 16
DETAIL DATA SETS PER MASTER DATA SET: 16
ENTRIES PER CHAIN: 65000
ORDERING INFORMATION:
32215A IMAGE/3000 Data Base Management System. Includes 800 bpi magnetic tape and manual (additional manuals optionally available). Price in U.S.A. \$10,000.
32215A-100 Same as 32215A but on 1600 bpi magnetic tape.
32216A QUERY/3000 Data Base Inquiry Facility. Includes 800 bpi magnetic tape and manual (additional manuals optionally available). Price in U.S.A. \$1,000.
32216A-100 Same as 32216A but on 1600 bpi magnetic tape.

IMAGE/2000 QUERY/2000

DATA ITEMS PER DATA BASE: 100
DATA SETS PER DATA BASE: 20 (If more data sets are needed it is possible to build up to 255 data bases on the system.)
MAXIMUM DATA ITEM SIZE: 126 bytes
MAXIMUM DATA ENTRY SIZE: 512 bytes
DATA ENTRIES PER DATA SET: 32767
LINKING: For fast retrieval of related information a data set may be linked to as many as five other data sets.
ORDERING INFORMATION:
24376B IMAGE/2000 with QUERY/2000. Includes tapes and manuals (additional manuals optionally available). Price in U.S.A. \$4250.
MANUFACTURING DIVISION: DATA SYSTEMS DIVISION
11000 Wolfe Road
Cupertino, California 95014 U.S.A.



Richard E. McIntire

Dick McIntire has been with HP since 1969, serving as the first project manager for IMAGE and QUERY, laboratory section manager for data base management systems, and special projects manager in data systems marketing. He is currently conducting feasibility studies of in-house on-line data base applications using IMAGE. Dick spent three years in the U.S. Navy before enrolling at Arizona State University to study mathematics. He received his BA degree in 1964, and for the next five years designed and implemented computer programs for satellite navigation and oil-field applications. Resuming his studies after joining HP, he received his MS degree in mathematics in 1971 and the MBA degree in 1973, both from the University of Santa Clara. He's a member of ACM and the Mathematical Association of America, and an associate professor in computer science at California State University, Sacramento. Bachelor McIntire lives in Los Gatos, California, and enjoys hiking and flying in his spare time.

Quality Frequency Counters Designed for Minimum Cost

Lowest possible cost was a major design objective for these counters. But so were high quality, excellent performance, ease of service, and efficient production.

by Lewis W. Masters and Warren J. O'Buch

NEW TECHNOLOGY AND HEIGHTENED user interest in price/performance tradeoffs have joined forces in recent years to generate a growing roster of lower-cost instruments. When the decision was made to add a minimum-cost HP frequency counter to this roster, a "design-to-cost" approach was taken. The broad objective was to produce an instrument that would have traditional HP quality at a specified selling price. This was translated into five specific objectives: low cost, good performance, high quality, ease of servicing, and efficient production. All five objectives had to be properly balanced.

Performance

The four performance characteristics considered most important were frequency range, sensitivity, resolution, and accuracy. It became clear that two counters were required to satisfy the largest number of customers. One would be a minimum-cost counter that would work from audio frequencies to about 80 MHz with moderate sensitivity, resolution, and accuracy. A second, higher-cost counter was indicated for those applications which demanded higher frequency (to 225 MHz), greater resolution, and greater accuracy. These two counters eventually became the 5381A and 5382A, respectively.

Low Cost

Once the marketing group had specified the new counters' selling prices and had established performance targets, the project was turned over to engineering. Because it was a "design-to-cost" project, every design decision required analysis of its effect on total cost. In many instances a complete conceptual design of a counter was required before all costs could be determined. For example, going to one fewer display digits would have required re-optimization of the display system, the decade counter system, the time base select logic, and possibly the power supply.

These cost estimates were more than mere summaries of parts cost; they also included detailed labor and warranty cost estimates. The best design, then, had to have balanced costs in all areas; lowered parts cost was not considered if it was offset by increased labor or service cost.

Quality

While the objectives of good performance and low cost presented many design challenges, the quality requirement provided many discussions but very few compromises. The counters had to be rugged and dependable. The two instruments were going to carry the HP label and therefore the tradition of reliability and quality established by the more expensive HP counters had to be continued in these two products costing only a few hundred dollars.

Efficient Production

It was obvious to all involved with the low-cost counter project that the only way the cost objective could be met and the performance and quality standards maintained would be to employ the most efficient production and testing techniques available. Since the best source of this information was the production personnel themselves they were involved during the early stages of the instruments' design. All aspects of the production process from printed-circuit board construction to board loading, from component selection to final assembly, and from pretesting to final testing had to be explored for the most efficient methods. The printed-circuit boards were designed as suggested by the production experts to make them easier to load and solder. Analysis showed that time could be saved by eliminating pretesting for all components except the display board and by designing dedicated testers to speed final testing.

Serviceability

The quality objective implied that the counters had to carry the standard HP warranty of one year on parts and labor. To achieve this objective, it was necessary to assure that when a counter did require repair, servicing could be accomplished easily, rapidly, and at low cost. As a result, the service engineer was involved during the initial design states. His early involvement meant not only that the instrument would be easily serviced but that the operating and service manual would help keep both the need for servicing and the attendant costs to a minimum. The service engineer made many recommendations for test points and for redesigning the printed-circuit boards to make the test points easier to find and more accessible. In the manual, test points are emphasized on schematic diagrams, and waveforms and truth tables are incorporated into the schematics. The manual contains information for users not familiar with common electronic equipment—for example, explanations of BNC connectors and shielded cables—and more applications information than is usually provided.

New Counters

The new low-cost HP counters are Models 5381A and 5382A, Fig. 1. They represent a significant contribution to the state of the art, not in sophisticated circuit design or esoteric performance, but in the careful optimization of all facets of the design.

A glance at the front panels of the new counters



Fig. 1. Models 5381A (top) and 5382A Counters measure to 80 MHz and 225 MHz, respectively.




Fig. 2. Large-scale integrated circuits reduce the number of components, thereby reducing assembly costs and simplifying troubleshooting.

reveals their simplicity. There are no gate lights, annunciator lights, sample rate controls, multiple inputs, time measurement capabilities, or BCD outputs. It is amazing how fast costs can increase as things get more complicated. More integrated circuits inevitably mean more printed-circuit boards, more power supply capacity, longer assembly time, lower reliability, and more difficult troubleshooting and repair.

Great care was taken to keep hand-wired connections to a minimum. In the design of the timebase and decade counter, two approaches were considered—many inexpensive integrated circuits or a few expensive ones. On an initial cost basis the former looked slightly better, but because it meant increased board-loading labor and more difficult troubleshooting, the LSI approach was chosen. Two low-power MOS LSI IC's replace 25 TTL MSI IC's (Fig. 2).

Another cost-cutting philosophy was to minimize research and development. Not having to repay a large research and development expenditure and not spending a lot of money on special tooling or custom integrated circuits mean lower prices. The new counters use as many standard commercial components as possible and borrow many parts from other HP counters. For example, the timebase and decade counter IC's and the sturdy cast-aluminum case come from the 5300A. The only parts unique to the 5381A and 5382A are the front and rear panels and the power transformer. Using standard parts also shortened the time from inception to introduction and is expected to increase reliability.

A low selling price does not require the use of low-quality parts. It was found that using lower-grade parts often would not save money. Since the parts selected are used in many other instruments the vol-

ume is high enough that prices are lower than for lower-grade parts used in smaller quantities. Also, non-standard or lower-quality parts would probably raise the warranty costs more than they might reduce initial costs. All resistors have 5% tolerance, electrolytic capacitors are of an industrial grade, and tantalum capacitors are used where possible.

Display Digits

The resolution of a direct-count frequency counter is dependent only on the gate time selected, not on the number of digits in the display. However, it is more desirable and convenient to have many digits; this gives maximum resolution without the ambiguity inherent in overranging.

At first, it seemed that more digits would mean higher costs because of the additional light-emitting diodes, integrated circuits, power supply capacity, and labor that would be required. However, when five-digit and seven-digit versions of the 5381A were compared, the costs were surprisingly close. This happened because seven digits make certain design economies possible. The seven-digit counter makes effective use of a proprietary HP six-decade counter/latch/multiplexer LSI integrated circuit. This IC not only costs less than the equivalent TTL IC's, but also saves assembly labor, board space, repair time, and power.

The seven-digit counter also uses a simpler gate-time selection method. A five-digit 80-MHz counter should have five gate times, but a five-position gate-time selector would have to include a costly rotary or pushbutton switch. With seven digits the two fastest gate times could be eliminated with no loss of versatility, and an inexpensive, reliable, three-position

slide switch could be used. Because of the small price differential and large gain in convenience, the seven-digit 5381A was chosen. The 5382A has eight digits for the same reasons.

Maximum Count Rates

To provide maximum versatility and value the highest count rates commensurate with cost were desired for the two counters. Count rates above 80 MHz require the use of ECL (Emitter Coupled Logic) circuitry. The special ECL power supply and many signal level translations required to interface with the TTL control circuitry made this approach too expensive for the 5381A. Therefore, the 5381A uses the fastest available TTL counter decade.

The trade-off in the 5382A was somewhat different. Above 225 MHz costs skyrocket because the input frequencies become high enough to require special input amplifiers and impedance matched inputs. Furthermore, most customers who need higher count rates need it much higher, say 500 MHz or so. The 225 MHz cutoff was chosen to cover as many of the VHF communications bands as possible without incurring inappropriate costs.

Input Sensitivity

At first, very high input sensitivity sounds like a good idea. However, there are at least two drawbacks: it is expensive and it can lead to measurement problems for many customers. To illustrate the latter point, suppose one has a counter that has 10mV rms sensitivity and wants to measure a 10V rms signal. That 10V signal will have to have better than a 60-dB signal-to-noise ratio to guarantee correct counting (60 dB below 10V is 10mV). In practical situations, few signals are as clean as this. For these reasons we chose to provide moderate basic sensitivity and to spend extra money on frequency-compensated attenuators to help cope with noisy signals.

Crystal Timebase Oscillators

A precision crystal oscillator is the heart of any frequency counter; the absolute accuracy of the numbers displayed is directly dependent on the oscillator accuracy. Both counters' standard crystals are specified to have a low aging rate; this cost more but was justified by the longer periods between recalibration—a good price/performance trade-off.

The crystal in the 5381A may have four times as much temperature variation as that in the 5382A. This is not because it is an inferior crystal, but because its lower resonant frequency is more difficult to control. It was used because the lower frequency greatly simplifies the oscillator-timebase circuit.

The added expense of the 5382A crystal oscillator-timebase was justified by the 5382A's greater frequen-



Fig. 3. Amateur radio is one of the applications for Model 5382A. VHF television is another. Model 5381A is suitable for commercial mobile radio and citizens band.

cy range and resolution. For extra accuracy in such jobs as checking radio transmitters for FCC compliance, the 5382A/Option 001 is available. This is a high-stability temperature-compensated crystal oscillator (TCXO) that has an appropriately low aging rate.

Acknowledgments

The authors gratefully acknowledge the contributions of the following people: Al Bagley and Ian Band for their contributions to the final electronic design; Roy Criswell, Don Larke, Burt Olson, and Jim Feagin for making production happen; Dick Buchanan for making the instruments serviceable; and Ernie Drake for the unique and effective operating and service manual. 



Warren J. O'Buch (right)

Born in New York City, Warren O'Buch attended Columbia University there. He received his BS degree in 1964 and his MS in 1966, both in metallurgical engineering, then spent three years in the U.S. Marine Corps, serving as a company commander. A civilian again, Warren worked as a systems analyst before enrolling at Stanford University for his MBA degree, which he received in 1971. At HP since 1971, Warren has been a systems analyst, IC department accounting manager, division cost accounting manager, and low-cost-counter product manager. He has just been named division materials manager. A member of AIME, Warren lives in Cupertino, California with his wife of two years. Tennis, golf, and backpacking are his choices for recreational activities.

Lewis W. Masters (left)

Salt Lake City native Lew Masters graduated from the University of Maryland in 1966 with a BS degree in mechanical engineering. In 1969, he received the MS degree in electrical engineering from the University of California at Santa Barbara. With HP since 1970, he has designed two functional modules for the 5300 Measuring System, assisted in the design of two others, and designed the 5381A and 5382A Counters. He and his wife have recently acquired a home in Los Altos, California, so Lew is currently spending most of his spare time remodeling, building furniture, and gardening. Eventually, however, he hopes to find more time to indulge his interests in building hi-fi equipment and working on cars.

SPECIFICATIONS

HP Models 5381A and 5382A Frequency Counters

5381A

FREQUENCY RANGE: 10 Hz to 80 MHz

DISPLAY: 7 Digit (LEDs)

INPUT IMPEDANCE: 1 M Ω , <50 pF

SENSITIVITY: 25 mV (rms sine wave) 30 Hz to 20 MHz

50 mV (rms sine wave) 10 Hz to 80 MHz

INPUT ATTENUATOR: Three position ($\times 1$, $\times 10$, $\times 100$)

MAXIMUM INPUT LEVELS:

ATTENUATOR $\times 1$	dc to 40 Hz	200 V (dc + peak ac)
	40 Hz to 100 kHz	250 V rms
	100 kHz to 5 MHz	2.5×10^7 V Hz
ATTENUATOR $\times 10$, $\times 100$	>5 MHz	5 V rms
	dc to 40 Hz	200 V (dc + peak ac)
	40 Hz to 1 MHz	250 V rms
	1 MHz to 50 MHz	2.5×10^8 V Hz
	50 MHz to 80 MHz	5 V rms

ACCURACY: ± 1 Count $\pm$ Time Base Accuracy

GATE TIMES: Manually Selected 0.1 second, 1 second, 10 seconds

RESOLUTION: 10 Hz at 0.1 second gate time, 1 Hz at 1 second gate time, 0.1 Hz at 10 second gate time

TIME BASE:

Internal

FREQUENCY: 1 MHz Crystal

AGING: <0.3 ppm/month

TEMPERATURE: ± 10 ppm 0°C to 40°C

LINE VOLTAGE: ± 1 ppm for 10% line variation

External Input

FREQUENCY RANGE: 10 kHz to 2 MHz

SENSITIVITY: TTL Level or 2.5 V rms sine wave

MAXIMUM INPUT: 25 V rms dc to 2 MHz

5382A

FREQUENCY RANGE: 10 Hz to 225 MHz

DISPLAY: 8 Digit (LEDs)

INPUT IMPEDANCE: 1 M Ω , <40 pF

SENSITIVITY: 25 mV (rms sine wave) 30 Hz to 10 MHz

50 mV (rms sine wave) 10 Hz to 225 MHz

INPUT ATTENUATOR: Three Position ($\times 1$, $\times 10$, $\times 100$)

MAXIMUM INPUT LEVELS:

ATTENUATOR $\times 1$	dc to 40 Hz	200 V (dc + peak ac)
	40 Hz to 100 kHz	250 V rms
	100 kHz to 5 MHz	2.5×10^7 V Hz
ATTENUATOR $\times 10$, $\times 100$	>5 MHz	5 V rms
	dc to 40 Hz	200 V (dc + peak ac)
	40 Hz to 1 MHz	250 V rms
	1 MHz to 50 MHz	2.5×10^8 V Hz
	50 MHz to 225 MHz	5 V rms

ACCURACY: ± 1 Count $\pm$ Time Base Accuracy

GATE TIMES: Manually Selected 0.1 second, 1 second, 10 seconds

RESOLUTION: 10 Hz at 0.1 second gate time, 1 Hz at 1 second gate time, 0.1 Hz at 10 second gate time

TIME BASE:

Internal

FREQUENCY: 10 MHz Crystal

AGING: <0.3 ppm/month

TEMPERATURE: ± 2.5 ppm 0°C to 40°C

LINE VOLTAGE: ± 0.5 ppm for 10% line variation

External Input

FREQUENCY RANGE: 100 kHz to 10 MHz

SENSITIVITY: 250 mV rms, 1 k Ω input impedance

MAXIMUM INPUT: 25 V rms dc to 10 MHz

OPTION 001: Temperature Compensated Crystal Oscillator

FREQUENCY: 10 MHz

AGING: <1 part in 10^7 /month

TEMPERATURE: ± 1 ppm 0°C to 40°C

LINE VOLTAGE: 1 part in 10^7 for 10% change

Meets or exceeds F.C.C. requirements.

5381A and 5382A General Data

OPERATING TEMPERATURE: 0°C to 40°C

POWER REQUIREMENTS: 100, 120, 220 and 240 V rms (48 Hz to 440 Hz) (+5% - 15%) 20 VA max.

WEIGHT: Net: 4.75 lb (2.2 kg) Shipping: 6 lb (2.8 kg)

DIMENSIONS: 3.5 in H $\times$ 6.25 in W $\times$ 9.75 in D (89 mm $\times$ 160 mm $\times$ 248 mm)

PRICES IN U.S.A.: 5381A, \$249; 5382A, \$450; Option 001, add \$100 to price of 5382A (no external oscillator with Option 001).

MANUFACTURING DIVISION: SANTA CLARA DIVISION

5301 Stevens Creek Boulevard

Santa Clara, California 95050, U.S.A.

A Versatile Bipolar Power Supply/Amplifier for Lab and Systems Use

The elements of a power supply/amplifier can be combined to perform as several different instruments: a precision bipolar power supply, a fast-responding programmable power supply, a low-distortion current-limiting power amplifier, a power amplifier with programmable gain.

by Santo Pecchio

FUNCTIONAL ELEMENTS ASSEMBLED in an instrument to do one job can often be reconnected to do several other jobs, thus enhancing the usability of the instrument.

Such has been the case with the bipolar power supply/amplifier. The moving spirit behind the development of this instrument was automatic testing, which requires power supplies capable of fast response to external control and, preferably, with bipolar output. Although Hewlett-Packard has developed a number of digitally-programmed power supplies to meet this need, numerous applications do not need the high accuracy (0.01%) of these instruments. The bipolar power supply/amplifier evolved to fill the need for moderate accuracy at lower cost.

Bipolar Power Supply/Amplifiers (BPS/As) achieve an accuracy of 0.1% but they respond to control signals within 0.1 ms, much faster than the 50 ms or so of the usual laboratory power supply. However, because of the various elements built into BPS/As to fulfill this function, they perform several other jobs well, and thus become useful tools to have around the lab. An output that is bipolar and continuously variable through zero is a capability that the usual lab supply lacks but that the BPS/A has. This can be useful for checking semiconductor characteristics, for example. Extremely low output impedance makes the BPS/A useful as a direct-coupled amplifier and/or impedance converter that can drive a variety of loads. The BPS/A, within its bandwidth, can thus be used to increase the output capabilities of function generators and other signal sources that have limited drive capabilities. The fast response to programming controls, when working with a suitable resistance or voltage programmer, also makes them usable as waveform synthesizers.

What It Is

To give an idea of the varied capabilities of a bipolar power supply/amplifier, let us take a closer look at what it is. It is first of all a precision, low-noise, constant-voltage power supply with adjustable current limiting. The output voltage can be set by a front-panel control or by using an external resistance or voltage programmer, and it is bipolar with a smooth, continuous transition from one polarity to the other. Like other HP power supplies, two or more BPS/As can be operated in an auto-parallel or auto-series mode to obtain either higher voltage or greater current with one-knob or single-programmer control of all supplies.

The BPS/A is also a power amplifier with manually-varied or programmable gain. Its low output impedance, less than 2 m Ω , enables it to deliver full output with low distortion (total harmonic distortion is typically less than 0.1%).

Another capability is the ability to sink current, as well as supply current. This is particularly useful in certain kinds of tests, such as tracing hysteresis curves or performing other cyclical tests on inductive loads.

Yet another capability, not of prime importance but nevertheless indicative of the versatility of these instruments, is the ability of the BPS/A to function as a programmable load. This derives from the current-sink capability of the instrument.

The Latest

A new series of bipolar power supply/amplifiers has now been developed (Fig. 1). As power supplies, these have higher resolution than earlier models and as amplifiers, they have wider bandwidths. Line and load regulation (source and load effects) are in the



Fig. 1. New Bipolar Power Supply/Amplifiers are made in two versions: one with front-panel meters and controls for bench use and one with blank front panels for systems use. All have 0.01% line and load regulation but respond quickly (<0.1 ms) to external voltage or resistance programming commands.

0.01% class. In addition, they have constant-current operation, rather than simple current limiting, with automatic crossover between constant voltage and constant current operation. This characteristic is particularly useful in semiconductor testing where the current at a breakdown voltage should not be exceeded beyond a chosen level. The new BPS/As also have a fixed gain mode that gives an amplification factor accurate within 0.1%, a useful feature when they are used to give a known output level when boosting the output of function generators.

There are three models in the new series with outputs as shown in the table. All are dual-range instruments that give a corresponding increase in resolution when operated on the lower range.

Model			6825A	6826A	6827A
Power Supply Output	Voltage	High range	-20 to +20V	-50 to +50V	-100 to +100V
		Low range	-5 to +5V	-5 to +5V	-10 to +10V
	Current		0 to 2.0A	0 to 1.0A	0 to 0.5A
Amplifier	Fixed Gain	High range	$\times 4$	$\times 10$	$\times 20$
		Low range	$\times 1$	$\times 1$	$\times 2$
		Bandwidth	dc - 40 kHz	dc - 40 kHz	dc - 30 kHz
	Variable Gain	High range	$\times 0 - \times 8$	$\times 0 - \times 20$	$\times 0 - \times 40$
		Low range	$\times 0 - \times 2$	$\times 0 - \times 2$	$\times 0 - \times 4$
		Bandwidth	dc - 15 kHz	dc - 15 kHz	dc - 15 kHz
Bipolar Power Supply/Amplifier Performance Characteristics					

As fixed-gain amplifiers, these instruments have bandwidths of dc to 40 kHz and a signal-to-noise ratio better than 65 dB. As variable-gain amplifiers, band-

width is dc to 15 kHz. The table shows the gains available in the fixed and variable gain modes. Maximum amplifier output corresponds to the maximum output in the power supply mode, e.g., 100 V p-p at 1A pk for the ± 50 V Model 6826A in the high range. Input impedance is 10 k Ω .

The new series of BPS/As is offered in two versions, one with full manual control from the front panel and with voltage and current meters, and one without meters or front panel controls for system use. Remote control of the voltage output, maximum current, and amplifier gain can be achieved with resistance control cards designed for the HP 6940A/41A series Multiprogrammers¹. These provide control of voltage and gain with 12-bit resolution and control of positive and negative current limits with 6-bit resolution.

A Look Inside

In discussing the design of the Bipolar Power Supply/Amplifier, it is helpful to recall that regulated power supplies fundamentally are precision power amplifiers with the reference voltage serving as the input signal and the voltage control functioning as a gain control.

Viewed as an amplifier, the conventional power supply has very low bandwidth, primarily the result of a fairly sizable output capacitor. This capacitor, in large degree responsible for minimizing noise, power supply ripple, and other short term fluctuations on the output, slows the power supply's response to changes in voltage level. Hence, to get the speed of response wanted for automatic test systems, the BPS/A has no output capacitor. Instead, the




Fig. 2. Simplified block diagram of Bipolar Power Supply/Amplifier. In the power supply mode, resistor R_V controls the output voltage over the full bipolar range; in the amplifier mode, it controls gain.

bandwidth of the amplifier circuits is made relatively wide so short-term disturbances in the output voltage, such as those caused by sudden changes in load, can be suppressed without the use of an output capacitor. Noise (rms) on the output in a 20 Hz to 20 MHz bandwidth is then on the order of 65 dB below the maximum output voltage.

In another departure from conventional power supply practice, the BPS/A uses two series regulators: one for positive outputs and one for negative outputs. As shown by the simplified circuit diagram of Fig. 2, the regulator circuit configuration resembles a single-ended push-pull amplifier operating in class AB.

As has been common practice, a precision reference voltage is developed by a temperature-compensated zener diode driven by an auxiliary power supply within the instrument. The manner in which the mixed reference is related to a variable output voltage, however, differs from conventional practice. This is explained with reference to Fig. 2. Here, the output voltage (E'_o) of the first amplifier is:

$$E'_o = -E_{in} \frac{R_V}{R_{in}}$$

When the instrument is in the power supply mode, E_{in} is the reference voltage. It is summed with E'_o at the input to the second amplifier, whose output E_o becomes:

$$\begin{aligned} E_o &= -E_{in} \frac{R_F}{R_A} - E'_o \frac{R_F}{R_B} \\ E_o &= -E_{in} \frac{R_F}{R_A} - \left(-E_{in} \frac{R_V}{R_{in}} \right) \frac{R_F}{R_B} \\ E_o &= E_{in} \left(\frac{R_V}{R_{in} R_B} - \frac{1}{R_A} \right) R_F \end{aligned}$$

When $E_{in} = 5V$, and $R_{in} = R_A = R_B = R_F = 10k\Omega$ as in the diagram:

$$E_o = 5 \left(\frac{R_V}{10^4} - 1 \right)$$

Now, when R_V is 0, E_o is $-5V$. When it is $10k\Omega$, E_o is 0 and when $20k\Omega$, E_o is $+5V$. Thus a bipolar output is derived from a single unipolar reference voltage. The single control adjusts the output voltage over the entire range, from the most negative value through zero to the most positive value. No polarity switching is required.

External resistances can be used in place of R_V . The

output voltage can thus be programmed by switching the resistors, and the output will be linearly proportional to the resistance.

Feedback resistor R_F is switched to either of two values by a relay inside the instrument, giving the instrument two output voltage ranges. In the high range, the nominal value of R_F is $100\text{k}\Omega$, 10 times its value in the low range. This increases the overall gain by a factor of 10, giving a range of $\pm 50\text{V}$ in the example just given.

Amplifier Mode

When the instrument is used as a variable-gain amplifier (non-inverting), the input signal replaces the reference voltage. Feedback resistor R_v then provides control of the gain and R_F gives a 10:1 change in gain. As in the power supply mode, external resistances may be used in place of R_v to enable remote control of gain.

When the instrument is in the fixed-gain mode, the input signal bypasses the input amplifier and goes directly to the output amplifier chain, as indicated in Fig. 2. Bypassing the input stage increases the upper bandwidth limit from 15 kHz to 40 kHz (and inverts the signal at the output). As before, resistor R_F provides a 10:1 change in gain.




Fig. 3. Current control circuit responds whenever the voltage drop across current-sense resistor R_S becomes equal to the voltage across either R_{2A} or R_{2B} .

Current Control

Control of the output current is derived from a low-value resistance R_S in series with the “-Out” bus, as shown in Fig. 3. The voltage developed across this resistance is compared to a reference voltage at the input to a current-control amplifier.

There are two current-control amplifiers, one for each polarity. When the instrument is in the constant-voltage mode (output current below the selected current level), these amplifiers are saturated, reverse-biasing the diodes in series with their outputs so they have no effect on the series regulators.

When the output current is great enough to generate a voltage equal to either reference, the appropriate current-control amplifier comes out of saturation, overriding the voltage-comparison amplifier and limiting the output voltage so the current is held at the selected level.

Whenever either current-control amplifier is active, it triggers the current-mode indicator circuit. This turns on the front-panel indicator to inform the operator that the instrument is operating in the constant-current mode, and it supplies a flag to a rear-panel connector for use by an external programmer.

Linear Current Programming

The action of the current regulator is such as to maintain the inputs to the active control amplifier at the same level so the voltage drop across R_S equals that across R_{2A} (or R_{2B}). Thus, whenever the instrument is in the constant-current mode, the voltage across resistor R_1 (or R_3) is maintained at a constant level. Consequently, the current through R_1 (or R_3) is held at a constant level so the voltage developed across control resistor R_{2A} (or R_{2B}) is a linear function of the resistance. External resistors can be used in place of control resistors R_{2A} and R_{2B} to permit external programming of the current.

Reverse Current

The ability to control current in either direction allows the BPS/A to sink current from an active load, e.g. the collapsing field of a transformer or other inductive device, at levels up to one half the full-rated output current. The current-control circuits function in a different manner, however, because the opposite series regulator would be in control.

As shown in Fig. 4, reverse current flows when the external voltage is greater than the BPS/A output voltage. The BPS/A in effect supplies a bucking voltage that reduces the reverse current that otherwise might flow. To prevent the reverse current from exceeding the selected limit, the BPS/A voltage adjusts itself accordingly.

Referring now to Fig. 3, if the BPS/A output voltage were positive but there is excessive reverse current,




Fig. 4. Diagram illustrates that the opposite polarity output transistor controls the current when the voltage of an active load exceeds the BPS/A output voltage.

the negative current-control amplifier would be turned on because diode D2 is now forward-biased. This alters the reference input to the negative current amplifier. This action causes the negative output series regulator to take control of the output.

Fig. 4 also indicates how the BPS/A may be used as a programmable load. The BPS/A is programmed for the desired current and for a voltage lower than that generated by the active load. The current level is then maintained by the control action described above.

Acknowledgments

A number of people contributed to the finished product. Many thanks are due Jim Gallo, Group Leader, for technical guidance, Don Pauser for the mechanical design, Bill Russo for building many breadboards and testing prototypes, and production engineer Itzhak Priel for the transition from prototype to production. Thanks are also due Al Levine for considerable assistance in the design and test of remote control resistance cards.

References

1. J. Mickowski, "Multiprogrammer Magnifies Minicomputer I/O Capacity," Hewlett-Packard Journal, August 1972.



Santo Pecchio

Santo Pecchio has a BSEE degree earned at Fairleigh Dickinson University. Following World War II, during which he served as a Signal Corps radio operator on a signal communication ship in the Pacific theater, Santo went to work for a manufacturer of home TV receivers, first in the design and manufacture of deflection circuits, and later in the RF-tuner development section. He continued along these lines working on 60 MHz IF strips until he joined HP ten years ago where he was involved in the Model 3205A Telemetry Signal Generator. Then it was to dc as the BPS/A project leader. In his off hours, Santo likes to hike, ski and point in oils. He and his wife have three offspring—a son, 24, and two daughters, 21 and 17.

PARTIAL SPECIFICATIONS

Models 6825A, 6826A, 6827A Bipolar Power Supply/Amplifiers

MODEL (BENCH VERSION) (SYSTEM VERSION):	6825A 6830A	6826A 6831A	6827A 6832A
DC OUTPUT VOLTAGE (high range):	-20V to +20V	-50V to +50V	-100V to +100V
OTHER OUTPUT CHARACTERISTICS: (See table on page 16)			
LOAD EFFECT (for a current change equal to the current rating):	0.01% + 0.5mV	0.01% + 1mV	0.01% + 1mV
LOAD EFFECT TRANSIENT RECOVERY (time to recover within indicated change of output voltage following a change in output current equal to the current rating):	0.1ms/20mV	0.1ms/50mV	0.1ms/100mV
SOURCE EFFECT (for a change in line voltage between 104 and 127 Vac or 208 and 254 Vac):	0.01% + 2mV	0.01% + 5mV	0.01% + 10mV
PAR (ripple and noise within 20Hz and 20MHz, rms/p-p):	5/15mV	6/35mV	10/50mV
TEMPERATURE COEFFICIENT (output change per °C change in ambient):	0.01% + 1.5mV	0.01% + 3mV	0.01% + 6mV
DRIFT (change in output over 8-hr interval under constant line, load, and ambient):	0.03% + 5mV	0.03% + 10mV	0.03% + 20mV
REMOTE CONTROL RESISTANCE COEFFICIENT:			
VOLTAGE High range	500Ω/V ± 0.1%	200Ω/V ± 0.1%	100Ω/V ± 0.1%
Low range	2000Ω/V ± 0.1%	2000Ω/V ± 0.1%	1000Ω/V ± 0.1%
CURRENT	5Ω/mA ± 1%	10Ω/mA ± 1%	10Ω/mA ± 1%
REMOTE CONTROL GAIN COEFFICIENT: (variable gain on high range)	4R _L */10.24kΩ	10R _L */10.24kΩ	20R _L */10.24kΩ
CONSTANT CURRENT VOLTAGE COEFFICIENT:	2A/V ± 1%	1A/V ± 0.5%	1A/V ± 0.5%
POWER (all models):	104-127/208-254 Vac, 48-63 Hz, 1.2A		
TEMPERATURE RANGE, OPERATING:	0 to 55°C		
COOLING:	Convection		
WEIGHT:	16 lbs, 8.2 kg		
DIMENSIONS:	7 1/4 in W × 6 in H × 12 1/2 in D (20 × 16 × 32 cm)		
PRICES IN U.S.A.:	Models 6825A, 6826A, 6827A: \$800 Models 6830A, 6831A, 6832A: \$900		
MANUFACTURING DIVISION:	NEW JERSEY DIVISION Green Pond Road Rockaway, New Jersey 07866		

*See Fig. 2 on page 17

An Automatic Exposure Control for a Lab-Bench X-Ray Camera

Locating shorted bonding wires in IC's, checking pad alignment in multilayer circuit boards, and finding much other "inside" information becomes as easy as taking snapshots when you use an industrial x-ray system that has automatic exposure control.

by John L. Brewster

X-RAYS PROVIDE THE MOST enlightening way to find out what's inside an opaque object, such as a relay can, a potted circuit, or a multiple-layer circuit board. But, because of complicated exposure techniques and radiation safety requirements, x-ray inspection was nearly always delegated to x-ray technician specialists. The result was a substantial delay in wanted data or, more likely, a decision to use a less effective technique.

Hence, a cabinet x-ray system that was easy to operate (Fig. 1) met with prompt acceptance when introduced some years ago. Using Polaroid film with this system, engineers and technicians could obtain radiographs like those shown in Fig. 2 in a few minutes. Then, because of its ease of use, other disciplines began to take advantage of it too. For example, foresters use it for examining seeds, cones, and wood sections, museums for paintings and encrusted archeological finds, and criminologists for weapons, documents, and packages.

Although the ease of radiography with this system can be compared to the ease of oscillography with dedicated oscilloscope cameras, there are significant differences. As shown by Fig. 3, film exposure time can vary widely with only moderate changes in either the thickness of a subject or the voltage on the x-ray tube. Furthermore, there is a wide range of film speeds associated with industrial radiography, a range of 2^9 or about 500:1. Users whose subjects required only a relatively narrow range of exposures soon developed exposure charts that consistently gave good results. On the other hand, there were many others who were interested in a variety of subjects, including composites of materials like plastics, aluminum, and copper. They either wasted a lot of film searching for an optimum exposure, or were forced to resort to complex calculations.

What was needed for these users was an automatic

exposure control. This could give anyone consistent results in the same way the "electric eye" helps the novice photographer get consistently good results in general photography.

Defining the Detector

A project for developing an automatic exposure control for the x-ray system was thus initiated. In



Fig. 1. X-ray system developed by the Field Emission Corporation, now Hewlett-Packard's McMinnville Division, has a lead-shielded enclosure with electrical door interlocks to ensure safety for everyday lab applications. It uses a simple, self-rectified circuit for the x-ray tube. The basic unit, with a film-to-source distance of 24 inches, covers film sizes up to 8 × 10 inches. Film to source distances up to 66 inches can be obtained with cabinet extensions.




Fig. 2. Radiographs disclose "inside" information, such as the broken bonding wire in the integrated circuit package at left, and the proper alignment of internal traces in the multilayer circuit boards at right.




Fig. 3. Curves show the exposure time required to obtain a given density in a typical film as a function of object thickness (in this case aluminum) and x-ray tube voltage. Small changes in either parameter result in a large change in exposure time.

view of the typical customer applications that had developed for the basic x-ray unit, it was decided the control should be able to accommodate the following parameter ranges:

- X-ray tube voltages between 15 and 130 kVp.
- Spectral sensitivity the same as x-ray films.
- Film sensitivity ranges over a 500:1 range.
- Exposure times ranging from 10 seconds to 5 minutes.

It was planned to place the detector *underneath* the film (the x-ray tube is at the top of the cabinet and the subject and film are on a shelf at the bottom). This was primarily for convenience but it would also avoid the shadows or distortions in the x-ray image that result from placing a detector above the film plane.

The design that resulted is based on the air-chamber dosimeter, which has proven to be reliable as well as simple. An air-chamber dosimeter of the type carried in the pockets of radiation workers consists simply of two coaxial electrodes, with the outer electrode being semi-transparent to radiation. A capacitor, charged to a predetermined voltage at the beginning of a monitoring period, establishes a potential difference between the electrodes. An incoming x-ray photon ionizes the air in the space between the electrodes and the resulting charged particles carry current that partially discharges the capacitor. The capacitor voltage at any time is thus inversely proportional to the x-ray exposure accumulated since the capacitor was charged.

Detector Design

Fig. 4 shows a cross-section of the detector that resulted from this approach. X-rays, after passing through the subject and film, go through the detector's thin conductive window, across the narrow air gap, and into the collector electrode. The collector is formed on a lead plate that prevents further passage




Fig. 4. X-ray camera detector consists of an aluminized mylar "window" electrode and a lead-backed molybdenum collector electrode with a narrow air space between. Electrons ejected from the molybdenum by x-ray photons ionize the air, giving rise to a current.

of the x-ray beam, which otherwise would cause spurious currents in the circuit connections beneath.

An electrometer monitors the voltage across the capacitance formed by the collector and the conductive window. This capacitance is charged to about 10 volts at the start of an exposure. X-ray photons reaching the collector eject electrons that ionize the air between the plates of the capacitor, creating a current that discharges the capacitance.

When the electrometer output falls about 5 volts, it triggers an SCR that activates a relay to open a contact, terminating the exposure by shutting off the x-ray tube.

This system functions analogously to an integrating light meter, terminating the exposure when the film has been exposed sufficiently to obtain the degree of darkening desired. To accommodate differing film speeds, kVp settings, and object "opacity", the gain of the electrometer is adjusted in a range of 1 to 50 by the FILM SPEED and LIGHT-DARK (fine exposure) controls. Additional range is provided by a lead shield that has a small hole through it. This can be slid into place over the detector to reduce the detector's effective sensitivity.

Exposure Control

With the development of the automatic exposure control, making radiographs has become simple and straightforward. The user first places the film and object in place, closes and latches the door, and resets the exposure-time indicator. Resetting the exposure-time indicator closes a relay (K1 in the circuit diagram, Fig. 5), charging the capacitance of the detector electrodes. He then presses the x-ray ON button, which turns on the x-ray tube and which also opens the relay.

Watching the EXPOSURE-PROGRESS meter, which monitors the electrometer output, the user turns the kVp control up from a low value until the meter needle starts to move. For high contrast photography of a fairly uniform object, he sets the kVp control no

higher than that required for very slow meter movement. For photography of an object that has a wide range of x-ray opacity, he sets the kVp control to a higher level where meter movement is fairly rapid.

The exposure is then allowed to continue, with no further attention required of the user. When the detector voltage falls to a level that causes the electrometer output to trip a unijunction transistor, which triggers the SCR, relay K2 is energized, turning off the x-ray tube current. This occurs when the EXPOSURE-PROGRESS meter reaches about 80% of full scale.

The film is then removed and processed. Should it turn out to be lighter or darker than desired, the LIGHT-DARK control can be readjusted accordingly and a new exposure made. Alternatively, the operator can note the time accumulated on the EXPOSURE-TIME indicator, set that value into the manual timer with the appropriate correction, and then make a new exposure using the manual mode.

Circuit Details

A diagram of the exposure-control circuit is shown in Fig. 5. Operation was outlined in the procedure described above. It is noted here that the EXPOSURE-PROGRESS meter simply monitors the output voltage of the operational amplifier, which is proportional to the accumulated exposure. The gain of the amplifier, and hence the sensitivity of the system is adjusted by the LIGHT-DARK and FILM SPEED controls.

The need to prevent leakage from the electrode capacitance required the use of a MOSFET for the electrometer input (transistor Q1, in Fig. 5). This gives an input impedance of 10^{15} ohms. To ensure low leakage, the input circuit is assembled in a lint-free area and enclosed entirely in a lead enclosure, one part of which is the lead backing of the collector electrode. Pressure relief in the enclosure to allow for changes in atmospheric pressure is provided by leakage of air through the connecting cable. Low humidity is maintained within by a desiccant.




Fig. 5. Detector circuit uses a high-impedance MOSFET input. Capacitor C1, which is formed by the detector electrodes, is charged through relay K1 at the start of an exposure and discharged by the x-ray radiation acting on the detector.

Spectral Response

To ensure accurate film exposures at all x-ray voltage settings, it is important that the detector's spectral response match that of the film closely. Fig. 6 is a plot of the relative sensitivity of a typical x-ray film as a function of monochromatic x-ray beam energy. Discontinuities occur where the beam energy matches the energy of orbital electrons in the silver-bromide molecules (like photographic film, x-ray film consists of small silver-bromide crystals suspended in a gelatin on a plastic film base). One discontinuity occurs for silver and one for bromine. X-ray absorption, and hence film sensitivity, is greatest just above the silver discontinuity. In actual practice, however, the effects of the discontinuities are smoothed out because the x-ray beam is not monochromatic but spans a narrow range of energy levels.

The obvious solution would be to make the collector out of silver, which would give it a response closely matching that of the film. But, since the detector is to be placed behind the film, it would receive less energy at those energy levels where the film is more absorptive, which is precisely where maximum detector output is needed.

The solution to this problem was to use a collector material that has a peak response between the silver and bromine discontinuities. The material chosen is molybdenum. With the energy spread that exists

in the x-ray beam, this gives high detector output at the same time there is high absorption by the film.

The detector "window", which serves as the second electrode, is made of a thin coating of aluminum deposited on a mylar film. So that x-ray absorption in the window is minimized, the coating is no thicker than that required to make it electrically conductive.

Another consideration involved detector sensi-




Fig. 6. Blackening of a typical x-ray film as a function of monochromatic x-ray energy level. The curve peaks where the x-ray energy matches the inner orbital electron energies of silver and then drops steeply below this level. The same phenomenon occurs with the bromine peak farther down.

tivity at the higher x-ray energies. The energy of ejected electrons increases with the incident x-ray energy so the mean free path of ejected electrons is longer at the higher energies. This results in more collisions with air molecules causing increased ionization and consequently greater current. The air gap was therefore made very short, on the order of 0.01 inch, to prevent this effect from increasing detector sensitivity at the higher x-ray energy levels.

As a result of these considerations, the detector gives consistent exposure control over the kVp range of the system. Thus, the power of x-ray analysis has been made available to scientists, engineers, and technicians without the need for special training.

SPECIFICATIONS

HP Model 43805 Shielded Cabinet X-ray System (with automatic sensor)

X-RAY TUBE: Beryllium window, 0.5 mm source.
VOLTAGE: 10 kVp to 110 kVp, 130 kVp with option 004.
CURRENT: Approximately 3mA continuous.

TYPICAL RADIATION OUTPUT

Tube Voltage kVp	Distance from Anode (inches) (cm)		Approximate Dose (Roentgens/minute)
110	12	32	300
110	25	63.5	70
130	12	32	400

COMPARTMENT SIZE: 14.5 in h x 18.25 in w x 15.5 in d (36.8 x 46.5 x 39.5 cm).

FILM-TO-SOURCE DISTANCE: 24 inches or 61 cm (up to 58.5 inches, 149 cm, with special configurations).

BEAM SIZE AT LOWEST FILM POSITION: 12.5 in or 32 cm diam. (up to 19 in, 46 cm, with cabinet extensions).

RADIATION SAFETY: Less than 0.5 mr/hr at 2 inches (5 cm) from any surface.

POWER: 110-120 V, 190-250 V, 50-60 Hz, 600 VA.

DIMENSIONS: 33.7 in h x 21.5 in w x 20 in d (85.6 x 54.6 x 50.8 cm).

WEIGHT: (Net) 400 lbs (181.44 kg).

PRICE IN U.S.A.: \$3485.

MANUFACTURING DIVISION: McMinnville Division
500 Linke Street
McMinnville, Oregon 97128

Acknowledgments

Special mentions for contributions during the project include Harry Cowherd for circuit design, Ted Morris for mechanical design, the late Frank Grundhauser for initiation and guidance of the project, and to Mike Hafner, Faxitron Group Leader, for advice and counsel.



John Brewster

A native of Phoenix, Arizona, John Brewster attended Arizona State and Pasadena College (AB in Physics, 1953). He earned his PhD at UCLA (1963) where he was a teaching and research assistant while working part-time on microwave gun design at Hughes Research Labs. He joined the Field Emission Corporation, now HP's McMinnville Division, in 1961. Initially he worked on high-power pulser systems, including project leadership on the Model 43706 Electron Accelerator. He became engineering manager in 1970. Interests include religious activities and travel in his pickup camper. He is married and has two teen-age daughters.

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

HEWLETT-PACKARD JOURNAL

JULY 1974 Volume 25 • Number 11

Technical Information from the Laboratories of
Hewlett-Packard Company

Hewlett-Packard S.A., CH-1217 Meyrin 2
Geneva, Switzerland

Yokogawa-Hewlett-Packard Ltd., Shibuya-Ku
Tokyo 151 Japan

Editorial Director • Howard L. Roberts

Managing Editor • Richard P. Dolan

Contributing Editors • Ross H. Snyder,
Laurence D. Shergalis

Art Director, Photographer • Arvid A. Danielson

Art Assistant • Sue M. Reinheimer

Administrative Services • Anne S. LoPresti

European Production Manager • Kurt Hungerbühler

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

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