31

TECHNOLOGY review

HARDWARE SOFTWARE FIRMWARE PROCESS ENGINEERING MATERIALS RESEARCH

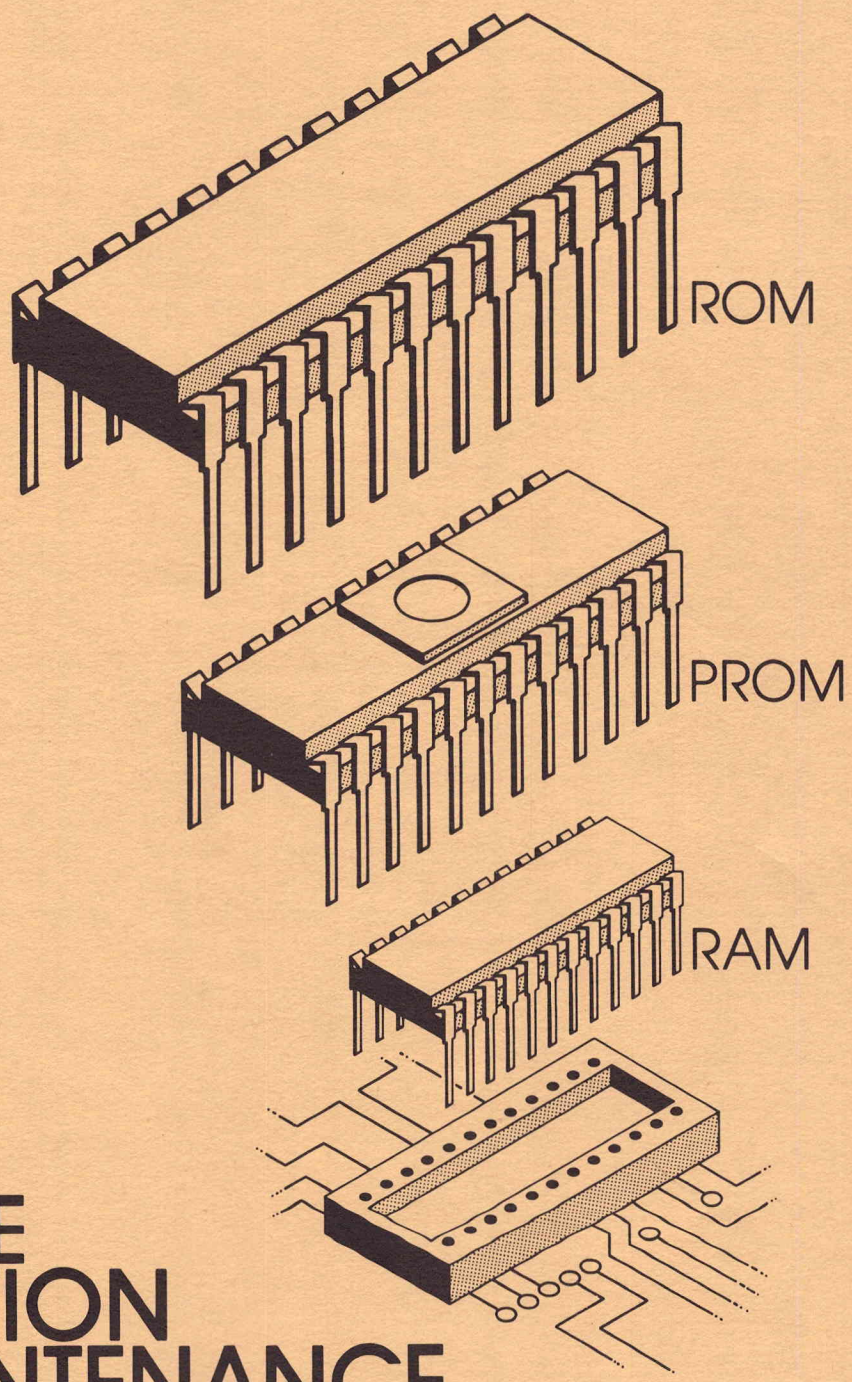

FIRMWARE PRODUCTION AND MAINTENANCE

CONTENTS

PUBLISHING OR PRESENTING A PAPER OUTSIDE OF TEK?

All papers and articles to be published or presented outside Tektronix must pass through Technical Communications Services for confidentiality review. TCS helps Tektronix employees write, edit and present technical papers. Further, the department interfaces with Patents and Licensing to make sure that patent and copyright applications have been filed for all patentable and copyrightable material discussed in the paper or article.

For more information and for assistance in producing your paper, call ext. 6795. □

Volume. 1, No. 1, October 1979. Managing editor: Burgess Laughlin, ext. 6795, d. s. 19-313. Associate Editor: Laura Lane. Cover and graphic design: Joe Yoder. Graphic assistant: Jackie Miner. Typesetter: Jean Bunker. Published by Technical Publications (part of Technical Communications Services) for the benefit of the Tektronix engineering and scientific community in the Beaverton, Grass Valley, Wilsonville and Vancouver areas.

Why TR?

Technology Review serves two purposes. Long-range, it promotes the flow of technical information among the diverse segments of the Tektronix engineering and scientific community. Short-range, it publicizes current events (new services available and notice of achievements by members of the technical community).

Contributing to TR

Do you have an article or paper to contribute or an announcement to make? Contact the editors on ext. 6795 or write to 19-313.

How long does it take to see an article appear in print? That is a function of many things (the completeness of the input, the review cycle and the timeliness of the content). But the *minimum* is six weeks for simple announcements and as much as 14 weeks for major technical articles.

The most important step for the contributor is to put the message on paper so that the editor will have something with which to work. Don't worry about organization, spelling and grammar. The editors will take care of those when they put the article into shape for you.

TECHNOLOGY REVIEW combines Engineering/Software News

Beginning with this issue, **Technology Review**, a new in-house publication, combines and replaces **Engineering News** and **Software News**. Combining these two publications offers several advantages:

- Reducing the number of in-house publications by one (a small step toward reducing in-basket clutter).

- Addressing one audience, the Tektronix technical community. Many articles address both the **Engineering News** and **Software News** audiences. For example, articles discussing firmware are important to both electronic engineers and software designers.

- Recognizing that the technical community's interests are diverse but interrelated. While specialization is important for furthering technology, specialists need to talk to each other. Having one publication for such diverse disciplines as electronic engineering, materials research, mechanical engineering, process engineering, firmware engineering, and software design, enables specialists in these fields to recognize the problems and follow the developments in other fields.

- Cutting costs. With one publication instead of two, there will be one press run, one cover, one mailing, and one set of boiler-plate information (such as the mailing coupon and credits).

Technical Publications, a department of Technical Communications Services, publishes **Technology Review.** The Technical Publications staff includes Burgess Laughlin (managing editor), Laura Lane (associate editor), Joe Yoder (graphic designer), Jackie Miner (graphic assistant), and Jean Bunker (typesetter). Technical publications also publishes **Forum Reports** and supports Engineering Activities Council presentations and publications.

For more information about **Technology Review** or about the Technical Publications department, call Burgess Laughlin, on ext. 6795. □

If you are on the EN or SN mailing lists, you will automatically receive Technology Review.

FIRMWARE PRODUCTION AND MAINTENANCE

Doug Bingham, Copier/Plotter Imaging Product Development, ext. 2676 (Wilsonville).

The following article is based on a presentation that Doug gave in Forum 13 in April 1979. (Sponsored by the Engineering Activities Council, Forums enable Tektronix engineers and scientists to present directly to management what engineers and scientists consider important in technology.) The topic for Forum 13 was "Managing Firmware Through Its Life Cycle." For a copy of Forum Report 13, call 6795 (Technical Communications Services).

After a complex firmware-based product is developed and released to the marketplace, any maintenance problems which arise must be solved quickly, efficiently, and inexpensively.

Reacting to these problems is often difficult because firmware in roms (read-only memories) can't be easily modified. This article shows designers how to avoid these problems by designing in hardware and firmware maintenance features, and reviews the formal modifications procedure to be followed after a product is in production.

PROBLEMS

Many different problems may arise during a product lifetime. First, there are **design and implementation errors.** These occur when a design team correctly specifies a product, but the design fails to meet the specifications.

Second, there are **design deficiencies.** These occur when the product was not specified correctly. This may be due to inadequate market research, the product entering an unexpected market, or misunderstanding customer needs. As a result, many customers either need an additional feature or want an existing feature to operate differently. Design deficiencies can be just as serious as implementation errors and may be more difficult to correct.

Third, the product may have **insufficient flexibility** to allow modifications, OEM special products, more options, and upgrading or downgrading during its market lifetime.

MAINTAINING FLEXIBILITY

The most basic hardware feature that makes firmware more flexible is having compatible control memories (ram, prom, and rom). As figure 1 suggests, each system memory socket should accept ram during program development (when the firmware is changing rapidly), and later prom (when firmware is frozen for evaluation, demo units, or early production).

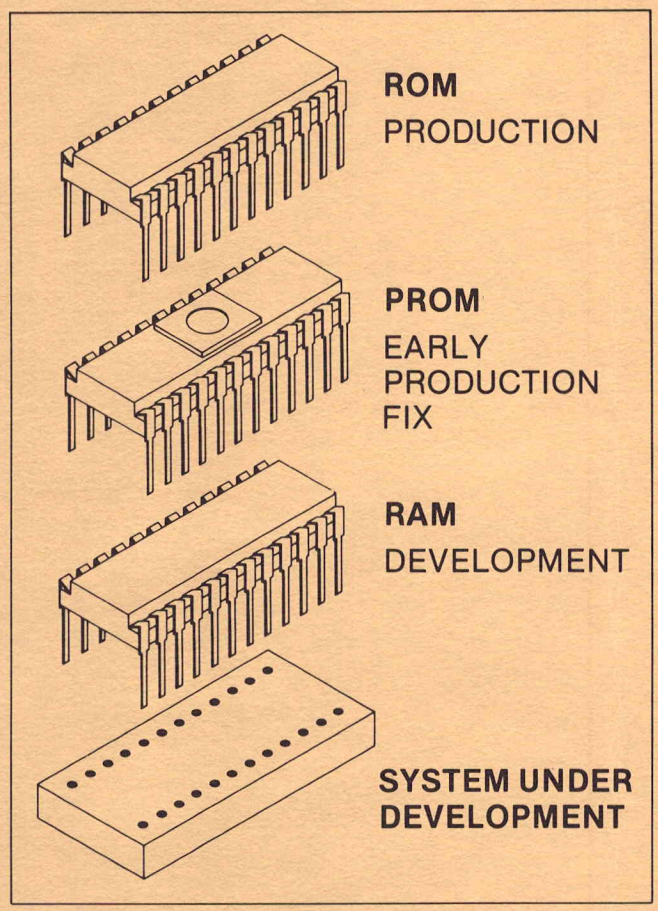

Figure 1. For the greatest design and maintenance flexibility, each system memory socket should accept ram during program development (when firmware is changing rapidly), and later prom (when firmware is frozen) either for evaluation, demo units or for early production or for fixes after production starts. Then, in the same socket, a designer can substitute roms when they are available for regular production.

Finally, in the same socket, a designer can substitute roms when they are available for regular production. For bus-oriented systems, these ram/prom sockets could be on a separate plug-in board which provides the same opportunity to use ram and prom during development, and then switch (during regular production) to rom in the built-in rom sockets.

After the firmware is in system roms, a hardware patch module makes the firmware modifiable. As figure 2 shows, a field programmable-logic array (fpla) device watches addresses that the processor places on the system address bus to access locations in the system roms.

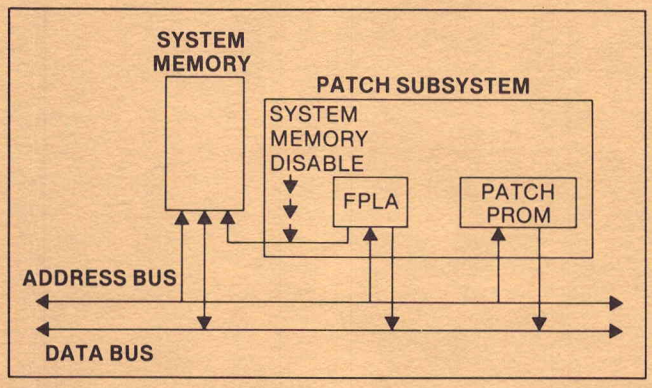

Figure 2. This firmware patch subsystem enables engineers to modify individual bytes in the system memory to fix a problem. If additional program steps are required in the fix, extra steps in the patch prom may be linked into the system memory by a jump instruction.

When the address of a location to-be-changed appears on the bus, the fpla recognizes this address and outputs nine bits of information. One bit disables the main system memory so that it does not respond to the address. The fpla then places the remaining 8 output bits on the system data bus and they return to the processor. To the processor, the contents of that location appear to be modified.

Refer to figure 2. If the firmware modification requires additional firmware instructions, then three addresses programmed in the fpla change three bytes in the system rom into a jump instruction. This instruction transfers control to a patch prom (contained in a hardware patch module along with the fpla) containing the extra firmware instructions. After these instructions are executed, control is transferred back to the system roms with another jump instruction.

This hardware patch module, containing the fpla and the patch prom, can be either built in or contained on a separate plug-in board. A plug-in board has a major benefit: the user obtains updated firmware simply by plugging in the board. If the plug-in slot is accessible from outside the system without disassembling the system, Manufacturing can mail the small update boards to customers for installation without requiring a service call.

In addition to hardware features, firmware flexibility is heavily influenced by the design of the system roms and by the use of ram in the system.

ROM INDEPENDENCE

Roms are the basic unit of firmware development (just as coding modules are the basic unit of software development). One reason firmware designers focus on individual roms is to closely watch the space remaining in each rom. Designers must strive to avoid a situation in which a problem occurs, but there is no room to fix it in that rom, and the fix must be implemented in other roms, scattering functions around in the system. Another reason is maintaining rom **independence** (one rom in the system can be changed without affecting the others).

One way to maintain rom independence is to carefully partition system firmware modules. Complete modules (and, ideally, groups of related modules and subroutines used in common by these modules) should reside in the same rom, to reduce rom-to-rom communication. This grouping produces math roms, utility roms, transform roms, and others.

Another way to maintain rom independence is assigning some specific random access storage memory to each rom with some other storage accessible to all the roms. If one rom changes, the use of the locations within the group assigned to that rom can change without affecting other roms in the system.

REFERENCES

Rom-to-rom references are necessary in any system, but some reference methods are more flexible and provide more rom independence than others. To maintain rom independence, designers should avoid direct rom-to-rom references because routine entry points have fixed locations. These fixed entry points can be awkward to maintain and a source of errors if the rom is changed.

The most basic reference method that provides some flexibility is using rom-based dispatch tables: if routine A in rom no. 1 needs to call routine B in rom no. 2, routine A transfers control to a jump table in rom no. 2. This jump table then passes control to routine B. The jump table must have a fixed location, usually at either end of the rom, but the location of routine B can move around within the rom as the rom changes.

Another reference method is having the jump table in system ram. This removes the jump table's rom space requirements, and also eliminates any fixed points in the system. At power-up, the initialization module fills this ram jump table from information in each rom's standard header.

As shown in figure 3, a standard rom format should consist of a checksum, the standard header, the code contained in the rom, and some empty space, in that order. When the rom is released for production, this empty space should be at least 10% of the size of the rom.

This empty space is for fixes and expansions of the code in that rom. As firmware changes may eventually exceed the patch capability, this space is required even if the hardware patch facility is available. If the patch capability is exceeded, some fixes must be taken out of the patch, put back in the related roms, and new roms procured.

Continued on page 6

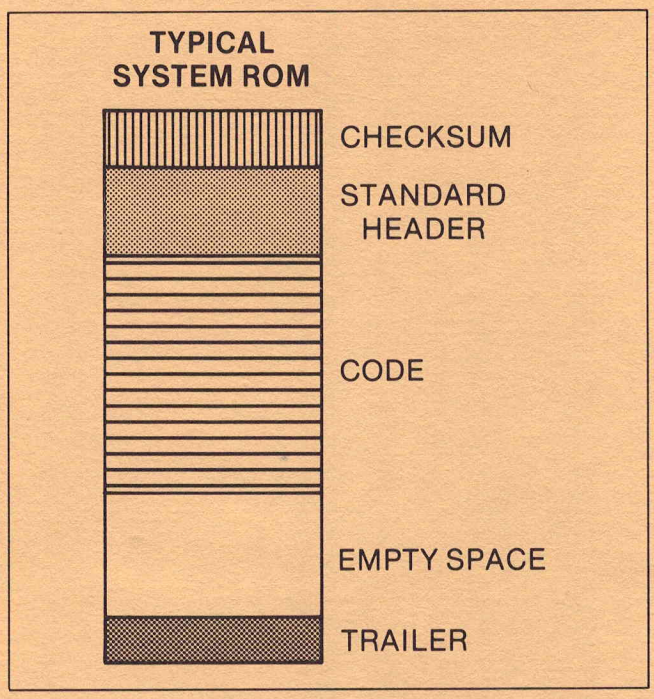

Figure 3. A standard rom format includes checksum, a standard header, code, and some empty space, in that order.

Continued from page 5

USING MAINTENANCE FEATURES

Some maintenance features are also useful during firmware development. First, independent roms allow use of absolute code and listings as valuable debugging aids. Second, limiting rom-to-rom references and separating ram assignments promotes independent rom development. Once designers have defined the overall system, several designers can work in parallel on different roms in the system.

This parallel operation allows all designers on the project to use a development library of working roms that represents the system's current state of development. Refer to figure 4. The development library is usually contained on local storage associated with the development system. A designer works on a new module or an updated module on a remote firmware development system.

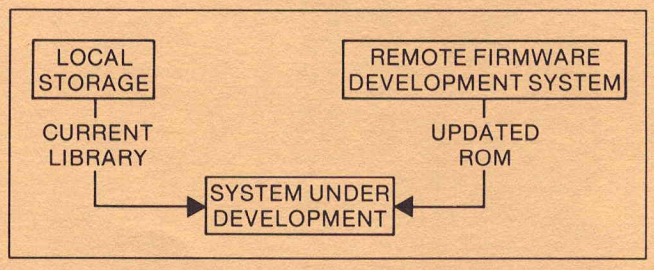

Figure 4. A firmware development library is usually contained in local storage associated with the development system. A designer works on a new module or an updated module on the remote firmware development system and then down-loads the new rom, overlaying the old rom in the system to obtain a new system for debugging.

The designer performs editing and assembly for only the rom containing the new module. He then loads the system library from local storage, and loads the updated rom information for the module on top of the old library. This procedure provides an updated system without having to assemble and link the entire system.

After debugging the new module, the designer updates the system library and makes it available to other designers, minimizing the coordination effort between them. A software evaluator can acquire a current copy of the development library to work on. In addition, separating individual roms encourages designers to focus on the space available in each rom and take steps to avoid running out of space.

Using independent roms meshes extremely well with modular programming techniques and enables designers to develop software in a straight-forward, top-down fashion because they can create a representative system very early in the design process using dummy or skeleton routines in the individual roms.

EARLY PRODUCTION

As figure 5 shows, shortly after a product enters production, new problems appear because the product is working in a real-world environment for the first time.

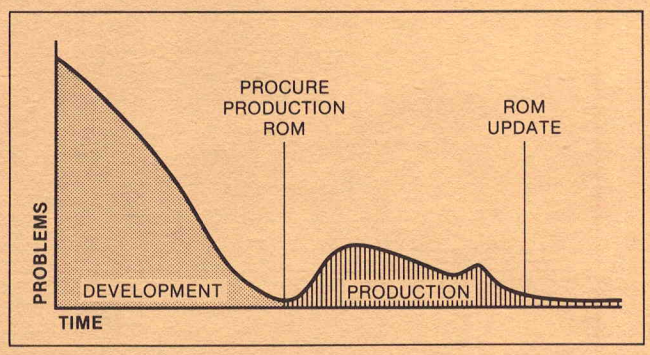

Figure 5. Even if a product is introduced into production without errors, problems will usually appear when the product is working in a real-world environment for the first time. The problem discovery rate will decline as the product matures, although surges of problems may appear when the product is introduced into new applications. When the discovery rate has declined sufficiently, Manufacturing may procure some new roms to retain some patch capability.

During this period, the new product's credibility is at stake, both in the eyes of new customers and the sales force. Firmware flexibility enables Engineering to respond to problems quickly and efficiently. If maintenance features are designed-into the product, hardware patches or proms can solve problems until new roms are available. Since rom manufacturing turnaround time is usually long, patches and the proms help Manufacturing respond quickly to problems. At the same time Manufacturing can accumulate fixes until producing new roms is cost effective and timely.

MID-PRODUCTION

As the product reaches mid-production, the number of new problems usually declines. However, surges of problems may appear as the product moves into new applications.

A key feature of midproduction is that maintenance personnel, rather than the original designers, fix problems. The firmware documentation must be very informative because maintenance people are less familiar with the product than the original designers. During this midproduction period, standard modified products and OEM specials appear; they too require complete documentation.

The extent to which a problem fix should be distributed is not always obvious. By midproduction there are many units in the field. Minor problems do not warrant blanket changes to all units.

Usually, we update all units in a product line only for major specification deviations or safety hazards. Updates are provided to particular customers who are experiencing problems. Service people upgrade units with the latest firmware when the units come to the service center for warranty repairs.

A flexible firmware architecture also allows new options and changes to the system to be produced (via new design projects) without major redesign of the basic system.

LATE PRODUCTION

A marketing goal during late-production may be to extend the product's market life with added options, enhancements, or a stripped-down version of the product at a reduced price. These efforts widen the market for the product and maintain sales for awhile.

A product can also be upgraded so it will be downward-compatible with a newer member of its product family. The flexibility features we've outlined enable designers to easily make rom changes during different stages of a product's market life to fix problems as well as react to changing market needs.

MODIFICATION PROCESS

The product modification process is a formal procedure for fixing problems and implementing changes once a product is in production. The modification process needs to be formal because so many people in the company are involved in changing even one rom. The procedure is complex because it includes coordinating engineering, manufacturing, marketing, service support, customer service, and purchasing groups, as well as outside rom vendors.

Outside vendors may need to be scheduled months in advance and the uncertainty of their delivery adds to the coordination problems. Most business units have a modifications group or have access to one.

A modification group handles the coordination of a firmware change. However, deciding when to make a fix and selecting one of several alternative fixes to a problem involves a product maintenance team having representatives from manufacturing, engineering, marketing, service support, customer service, and modified products groups.

These representatives should meet whenever anyone discovers a serious problem, or at least six months after a product goes into production and possibly each year thereafter. They review system problems, the remaining patch capability, and accumulated problems too minor for a unique patch.

The development of a rom fix should be managed just like a miniature product development. Once Engineering identifies the problem, the maintenance team should: (1) decide whether it should be fixed; (2) specify any constraints on the fix (for example, identify other problems that should be fixed at the same time considering their urgency/cost tradeoffs); (3) make sure the appropriate groups properly design, evaluate, document, and archive the fix; and (4) make sure the fix enters Manufacturing properly and then is available to the field.

Engineering should give field service people a bug list describing the remaining problems, their impact on the customer's use of the product, and how to work around these problems. Field Service also needs a description of the various firmware levels including a version chart showing which roms constitute a certain level of firmware. □

PLASTIC PART DESIGN SEMINAR

If enough Tektronix engineers are interested in attending, Education and Training will schedule a Plastic Part Design seminar for October 29, 1979. This seminar will cover engineering applications for: (1) Relating the plastic to the product, and (2) Designing for injection molding (the process and the equipment; tooling; and products and design factors).

Glenn Beall, the instructor, has spent eleven years with General Electric Co. as Test Facilities Design Engineer and was Abbott Labs Plastic Products R & D section head. He has extensive experience in plastics-design education and in plastics design.

The $2500 seminar cost will be divided among the attendees; the seminar will not be held unless there are ten or more people wishing to attend.

If you want to attend this seminar, call Barbara J. Sabo, ext. 7703 (Beaverton), or send a memo to 74-434. □

FIVE NEW MEMBERS JOIN ENGINEERING ACTIVITIES COUNCIL

In July, five new members joined the Engineering Activities Council. At the same time, other members, who have been with the Council since the end of 1977, left the Council (members serve for about 18 months). The new members are David Keith, David Kreitlow, Cathy Lin-Hendel, Anita Massey, and Wayne Thomas.

EAC CHARTER

The Council's primary objective is to provide engineers with a forum in which to present directly to management what engineers themselves consider to be important in technology. To meet their charter, the Council has sponsored 14 Engineering Forums ("Engineers Talk to Managers").

In each forum, four or five engineers discuss their viewpoints on the problems and progress of new technology. The forum chairpersons are Council members. They select forum speakers from the engineering

Bruce Ableidinger	LID, LDP Engineering	1742
Tim Flegal	LID, TM 500 Engineering	1533
Jim Tallman	LID, 7000 Series	7076
Dave Armstrong	SID, Advanced Service Instrument Development	5244
Jim Besemer	SID, DSI Firmware	7604
David Keith	SID, DSI Hardware	1916
Wayne Kelsoe	SID, Portables Engineering	6255
Ron Bohlman	Technical Support, CAD Development (Town Center)	221
Lynn Saunders	Technical Support, Software Engineering Research	5616
Hal Cobb	Tek Labs, Hybrid Circuits	6564
Elske Cordell	Tek Labs, Hybrid Packaging	7079
Cathy Lin-Hendel	Tek Labs, Display Research	6388
Mike Rieger	Tek Labs, Signal Processing Research	6907
Tom Woody	Tek Labs, Display Devices	7147
Wayne Thomas	Communications, TV Engineering	1819
Jim Zook	Communications, TV Engineering	7457
David Kreitlow	IDD, Copiers, Plotters, and Imaging Systems	2879
Anita Massey	IDD, Graphic Computing Systems	2431
John Moore	IDD, IDP Engineering	2648

Table 1. Bill Walker, executive vice president, selects Council members from a list of candidates who are nominated through three channels: (1) by managers of engineering departments, (2) by current Council members, or (3) through the candidates' own initiative (*Technology Review* articles announce Council openings). The current Council members are listed here.

Engineering Activities Council members: *front row, left to right,* Tom Woody, Jim Tallman, Elske Cordell, Hal Cobb, Wayne Kelsoe, Anita Massey, and Burgess Laughlin (support). *Back row, left to right,* David Kreitlow, Bruce Ableidinger, Dave Armstrong, John Moore, Wayne Thomas, Jim Besemer, Ron Bohlman, Lynn Saunders, and Mike Reiger. Council members not shown are, Tim Flegal, David Keith, Cathy Lin-Hendel, and Jim Zook.

	FORUMS	CHAIRPERSONS
1.	General Purpose Interface Bus	Robert Chew, Paul Williams
2.	A-D and D-A Converters	Bob Nordstrom, Mike Boer
3.	Video Display Techniques	Steve Joy, Phil Crosby
4.	New Technologies: I	John Addis, Bob Burns
5.	New Crt Technologies	Bob Oswald, Cal Diller
6.	Creative Microprocessor Hobby Projects	Dave Chapman, Joyce Lekas
7.	Managers Talk to Engineers	Mike Boer, John Mutton
8.	Microprocessor Design Pitfalls	Robert Chew, Paul Williams
9.	New Technologies: II	Hock Leow, Binoy Rosario
10.	Packaging	Bob Burns, Cal Diller
11.	Creative Microprocessor Hobby Fair Projects: II	Steve Joy, Hock Leow
12.	Reliability	Tim Flegal, Mike McMahon
13.	Managing Firmware Throughout Its Life Cycle	Lynn Saunders, Jim Tallman, Dave Armstrong
14.	Hard Copy Technology	Hal Cobb, Bruce Ableidinger

Table 2. In the last two years, the Engineering Council has sponsored 14 forums in which engineers presented (to corporate, divisional, and departmental managers) engineers' views of the problems and progress of new technology at Tektronix. The forum presentations are described in *Forum Reports*, which are distributed over the *Technology Review* mailing list. For a copy of a *Forum Report* listed here, call ext. 6795. (Forum Report 14 will be published in November.)

Continued from page 8

community. The audience consists of approximately 125 corporate, divisional, and departmental managers. Attendance is limited by the capacity of the auditorium, but the forum presentations are published as **Forum Reports** (to add your name to the **Forum Report** distribution list, mail your name, delivery station, and payroll code to FR List, 19-313).

FOR MORE INFORMATION

To suggest a forum topic or if you have questions about the Council, call one of the members (table 1 lists their phone numbers).

Table 2 lists past forums. Technical Communications Services has published **Forum Reports** for forums 1-14. For a copy of any of these **Forum Reports,** call ext. 6795. □

CYBER TESLA VERSION 2 RELEASED

Microprocessor Software Support has developed a Version 2 TESLA compiler (CYBER2) for Cyber (the A machine). Version 1 will continue receiving support.

Microprocessor users may find the compiler helpful in debugging (at an algorithmic level) microprocessor programs without having to download. A symbolic debugger enables the user to interactively (1) establish breakpoints and single step a program, (2) dump and change the content of variables, and (3) trace program flow using the TESLA source listing as a guide for setting breakpoints, single stepping, and referencing variables.

The compiler is a procedure file on user number TESLA:

GET,CYBER2/UN=TESLA

(CYBER2 is the Version 2 TESLA compiler.)

The procedure file accepts the name of the TESLA source file to be compiled. The compiler then produces a TESLA source listing, an assembly listing, a relocatable file, a load map, and an absolute file.

A write-up documenting the Version 2 TESLA compiler and supporting software is also available. The Microprocessor Software Support group developed the write-up (CYB2MAN) with ASCII BARB, allowing printing of upper- and lowercase characters. Because high-speed line printers can't handle ASCII BARB output, users must dispose output to a remote job entry (rje) printer by entering:

GET,CYB2MAN/UN=TESLA

DISPOSE,CYB2MAN/EI

and claim it at an rje printer.

For more information, call Sue Anne Smith, ext. 1890 (Walker Rd.). □

MINIMIZING CROSSTALK AND LINE-RINGING

The August/September Engineering News carried an article entitled "Wire-Wrap for High-Speed Logic Circuits." The article discusses ways to *minimize* (not *prevent*, as the article incorrectly states) crosstalk and line-ringing in wire-wrap prototype high-speed logic circuit boards. For a copy of the article, call Technical Communications Services, ext. 6795 (Beaverton).

WIRE-WRAP NEWSLETTER

The CAD/Wire-Wrap Group publishes, as the need arises, the **CAD Wire-Wrap Newsletter** which describes developments in CAD wire-wrap at Tektronix and related news such as the use of component carriers.

To receive the **CAD Wire-Wrap Newsletter,** send your name, delivery station and payroll code (required for processing the mailing list) to LIST, Glenna Jones, 50-126. □

PATENT RECEIVED: No. 4,146,844

FEED-FORWARD AMPLIFIER

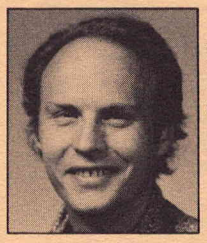

Patrick Quinn, Bipolar IC Design, ext. 5683 (Beaverton).

The circuit shown here uses a new technique to provide linear amplification at higher speeds than feedback techniques provide.

The circuit, known as the "Cascomp," is a high-precision amplifier which uses an error amplifier to provide first-order correction of thermal distortion and amplifier nonlinearity. The Cascomp amplifier requires only a few parts to provide precision of about 0.1 percent at frequencies of several hundred megahertz. The Cascomp also has overdrive characteristics similar to a simple differential amplifier.

The Cascomp is an inherently stable amplifier in which the correction amplifier (Q5, Q6) senses the main amplifier's (Q1, Q2) base-to-emitter signal distortion (as replicated at the emitters of Q3, Q4) and injects an error correction signal into the output nodes to provide an error-free differential output signal.

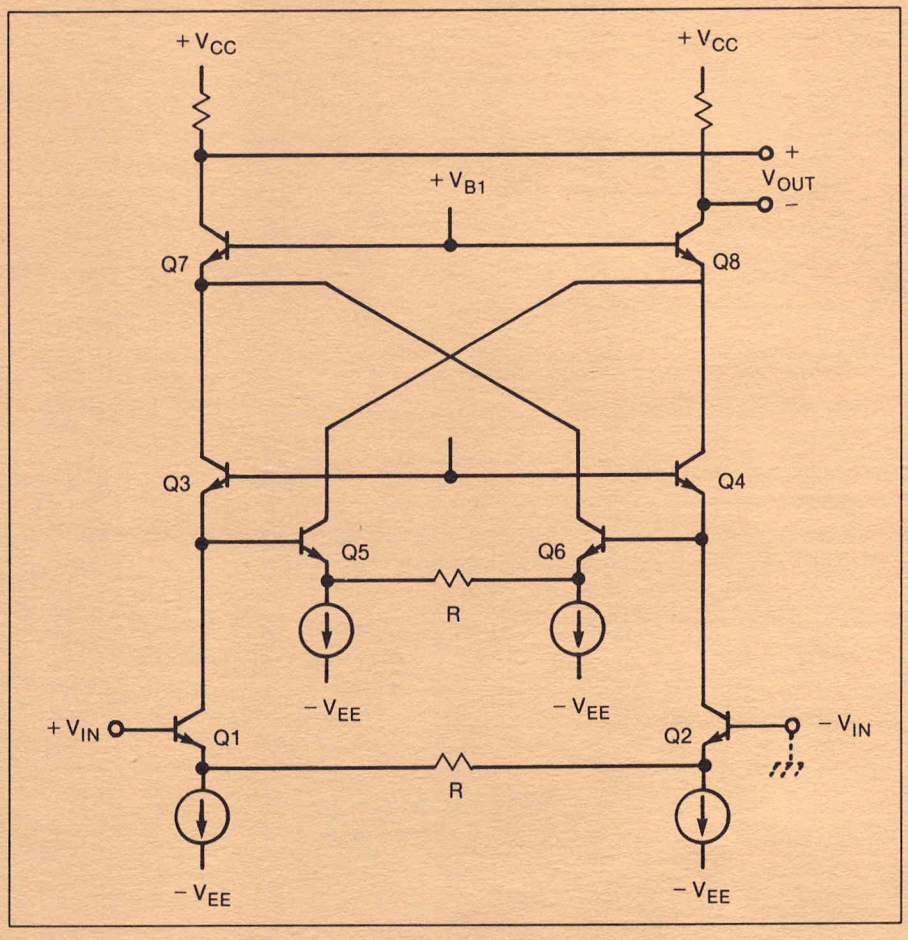

The second set of cascodes (Q7,Q8) are optional. Base voltage V_{B2} (and V_{B1}) need to be correctly chosen for minimum thermal distortion. □

Z80 OPERATIONS CLARIFIED

Because Z80 Microprocessor operations are not what programmers might expect, STS Engineering has developed a matrix which helps to end the confusion. The matrix shows all possible source and destination modes and the operations that can be performed between them.

Both axis of the matrix show:

- **All Z80 registers:** A, B, C, D, E, H, L, AF, BC, DE, HL, IX, IY, SP
- **Indirect address modes:** (HL), (BC), (DE), (IX+d), (IY+d), (SP)
- **Memory addressing:** (NN)
- **Immediate modes:** N (8-bit constant) and NN (16-bit constant)

Source modes are shown on the x-axis, destinations on the y-axis, and the operations that can be performed between them are shown at the intersection of the two. Italic type highlights instructions that set processor condition codes.

For a copy of the matrix, call Gerd Hoeren on ext. 1227 (Walker Road), or write to d.s. 92-701. □

PROFILE

DATABASE AND UTILITY SUPPORT

Among their other responsibilities, the Data Base and Utility Support group (part of the Scientific Computer Center) designs and develops data base management systems. DBUS responsibilities are carried out by group members *(left to right)*Glen Fullmer, Ingrid Palm, Greg Harris (project leader), and John Burley.

data base: an integrated file of data used by many processing applications.

data base management system: a comprehensive software system to build, maintain, and access a data base.

Data Base and Utility Support (DBUS), a department of the Scientific Computer Center, provides support to data base and data base management system (dbms) users. The group's responsibilities include dbms design, technical support for all Cyber users with business applications, and implementation of data-base and business-related jobs.

DATA BASE MANAGEMENT SYSTEMS

DBUS offers assistance in preparing and reviewing feasibility studies for a proposed dbms application, following up with assistance in designing the most efficient data base. As a dbms project progresses, they provide aid in detailed systems design, programming, and testing.

Occasionally, a group needing a data base management system is not large enough to justify its own full-time dbms staff. In this situation, DBUS will develop a system from the design stage through its implementation procedures, and then train that group to use the system.

Whether simply aiding in the design of a dbms application or actually developing it, DBUS offers continual consultation.

SYSTEM 2000

In designing most dbms applications, DBUS employs **System 2000,** a software package purchased by Tektronix from MRI Corporation (a division of Intel). System 2000 is a general-purpose data base management system operating on the Cyber system.

Another of the group's functions is the maintenance of System 2000 which includes implementing all software updates, testing those updates, and working closely with the Scientific Computer Center's Operations and System groups to implement data-base and business-related jobs. Implementing these jobs requires scheduling computer resources, estimating future resource requirements, and establishing backup and recovery procedures.

DBUS is also a focal point for Tektronix' System 2000 user community. The group is a central resource for System 2000 procedures, guidelines, definitions, and data-base standards. This activity reduces entering and storing redundant data by separate departments.

Serving as a coordination point, DBUS (1) gathers and disseminates all information of concern to System 2000 users and (2) keeps in contact with the National System 2000 Users Group. These communications make it possible for Tektronix' System 2000 users to learn from other users around the country as well as here at Tektronix.

CYBER

Besides supporting dbms, DBUS provides "utility support" for various other Cyber applications. These applications include information systems such as the Nichols 5500 Management Information System (N5500).

DBUS also supports off-line data entry, a new and rapidly growing function helping all Cyber users to get better response time during the day, while helping the Scientific Computer Center to more fully utilize the Cyber systems at night. During prime usage hours, the data entry stations collect data (independently of the Cyber) and, at night, transmit that data into the Cyber to update the data bases and produce the desired reports.

FOR MORE INFORMATION

DBUS is available for dbms services as described above. Even if you're simply curious about dbms, give DBUS a call. Call Greg Harris, Data Base and Utility Support project leader, at ext. 5102 (Beaverton). □

DBMS CLASS

If you're interested in a class (through TEP or MRI) on any aspect of data base management systems, contact Greg Harris, ext. 5102 (Beaverton).

APD STUDIES THIN GOLD FINISH FOR PROTOTYPE ECB CONTACTS

Abigail Cooke, Advanced Process Development (in Electrochemical Support), ext. 7830 (Beaverton).

Advance Process Development is examining several finishes for etched circuit board (ecb) contacts. Among these finishes, "thin" (20 micro-inch) gold holds promise for *prototype* circuit board contacts. Figure 1 identifies circuit board terminology, and table 1 summarizes contact finish alternatives.

Etched circuit board contacts have one prime requirement: continued low contact resistance throughout their lifetime. Thus the contact finish must have low initial contact resistance, sustained low contact resistance after use (wear resistance), and corrosion resistance to maintain the clear initial surface.

Requirements for prototype ecb contacts are less stringent than for production ecb contacts because prototype instruments have much shorter lifetimes. A prototype ecb contact finish must have low electrical resistance, but does not have to be as durable or as corrosion-resistant as a production ecb contact finish.

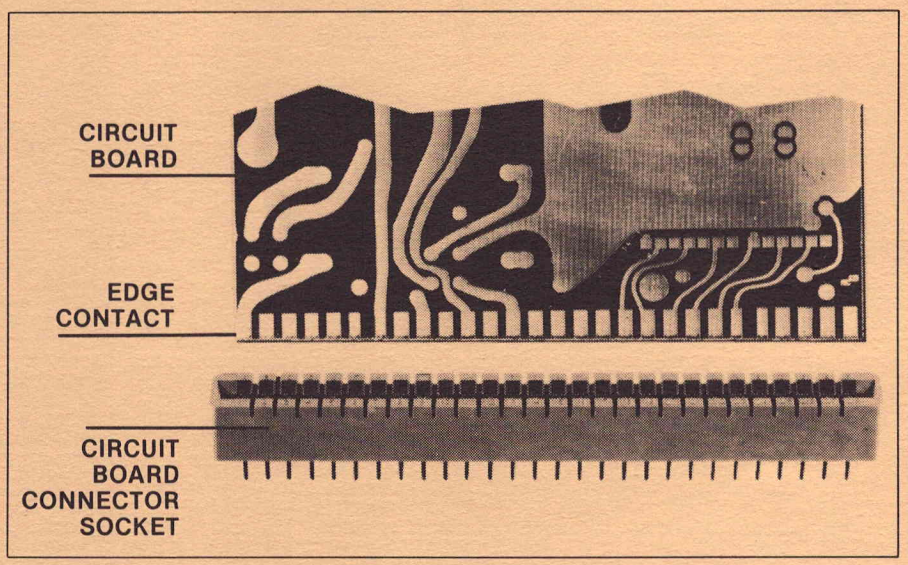

Figure 1. Circuit board nomenclature.

THIN GOLD

"Thin" gold is an electrochemically-deposited layer of pure gold 20 micro-inches thick.

Other circuit board manufacturers besides Tektronix use thin gold for applications such as edge fingers (for

CONTACT FINISH	COST PER SQUARE INCH	CORROSION RESISTANCE OF BULK METAL	NORMAL FORCE REQUIRED FOR CONTACT RESISTANCE UNDER 3 MILLIOHMS	CONTACT RESISTANCE WITH REPEATED USE
Copper	0¢	Poor	Not attainable in use	Higher
Tin	< 4¢	Good	450 Grams	Higher
Thick Gold (50 or 100 Micro-inches)	13¢ or 26¢	Excellent	50 - 100 Grams	Unchanged
Thin Gold (20 Micro-inches)	5.2¢	Good	100 - 150 Grams	Higher

Table 1. Etched circuit board contact finish alternatives include copper, tin, and various thicknesses of gold.

assembly only), pads for components soldered to a circuit board, and switch pads. Many connectors used at Tektronix have only 30 micro-inches of gold on the edge card connector which mates with the circuit board. For high-reliability applications, connectors are available which have up to 100 micro-inches.

A major advantage of using thin layers of gold is that they consume less gold. Based on a $250 per troy ounce gold price, Tektronix will spend about $3.25 million on gold in FY000, if we use as much as we did in FY900 (about 13,000 troy ounces).

Gold costs about 25 cents per square inch for 100-micro-inch layer, 12.5 cents for 50-micro-inch, and 5 cents for 20-micro-inch.

WEAR RESISTANCE

APD tested edge contacts that had thin gold deposited on nickel (which, in turn, was deposited over copper, as figure 2 shows). Gold has much less electrical resistance than nickel. When the gold layer wears thin (through removal and insertion of the board), the resistance of the contact rises. Gold completely worn away produces a connection more prone to corrosion and thus a higher-resistance connection.

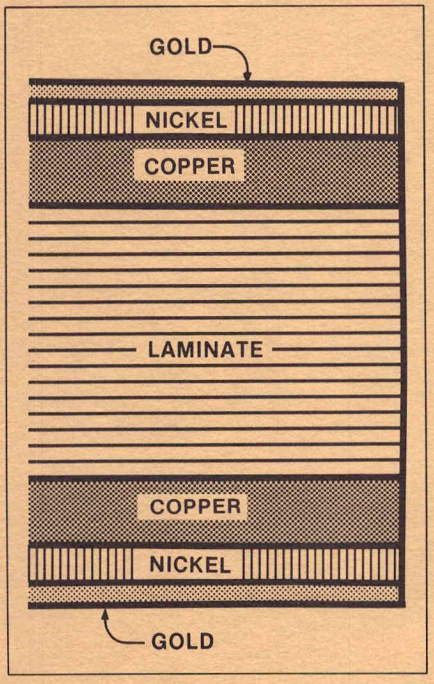

Figure 2. Cross-section of a thin-gold contact.

THE TESTS

To test the durability of thin-gold contacts, APD repeatedly inserted and removed test circuit boards from their sockets. The tested boards had contacts with varying amounts of gold. APD personnel, using a microscope, inspected the contacts after every ten cycles, and recorded the number of cycles when the underlying nickel was visible. The tests continued until 50% of the contacts showed nickel or until 400 cycles had been made, whichever came first. (The test limit was 400 cycles because a prototype circuit board is not likely to be removed and re-inserted 400 times during the life of the prototype.)

To test the corrosion resistance of thin gold, APD plated several boards with various thicknesses of gold over a standard nickel layer. After etching the boards, APD sent them to Environmental Labs for humidity-sulfide corrosion tests (described in September 1976 **Engineering News**; call ext. 6795 for a copy). APD examined the tested boards and compared the quantity of pores on each board.

In one respect, test results produced no surprises: thinner-gold contacts' first failure occurred sooner and subsequent failures occurred more often than thicker-gold contacts. Refer to table 2. "Failure" was defined as a circuit board contact on which the underlying nickel is visible.

But in another way, the test results were a surprise. With 20 micro-inch gold contacts, the first failure didn't occur for 310 insertion-removal cycles. That's acceptable for most short-life prototype instruments or test boards. After the testing, four contacts from each edge connector were examined. Copper was visible on all the connector contacts examined, indicating gold had worn away. This indicates connectors may be a weak spot in an electrical system.

CORROSION RESISTANCE

Humidity-sulfide gold-porosity tests of ecb's plated with various gold thicknesses revealed expected results: 20 micro-inch boards had more **pore corrosion** (spotting at openings in the gold plate) than thicker gold plated boards. Thicker gold plating tends to fill the pores.

APPLICATIONS

About 15% of the prototype boards passing through the Quick Board Line use thin-gold finish for contacts. Since thin gold was introduced in 904, almost all product groups have used thin gold for some prototype boards.

Because they do not have the same degree of wear or corrosion resistance as 50- or 100-micro-inch gold, thin-gold contacts are not suitable for prototype instruments undergoing humidity tests or for prototype instruments intended for lifetime testing.

If your application seems to fit thin gold's properties — low contact resistance, and few insertions or wiping cycles — try it. You'll be saving all of us money. □

GOLD THICKNESS	NUMBER OF CYCLES TO	
	FIRST FAILURE	END OF TEST
10	100	160 (50% failed)
15	120	220 (50% failed)
20	310	400 (18% failed)
25	310	400 (28% failed)

Table 2. Advance Process Development tested circuit board edge contacts show that thin-gold (20 micro-inch) contacts can provide adequate wear properties for many prototype boards.

WHEN SIZING PROCESSORS, THREE NUMBERS ARE BETTER THAN ONE

Jack D. Grimes, Graphic Computing Systems Engineering, ext. 2558 (Wilsonville).

First, a test. List the sizes of the following computers and microprocessors:

PDP-11 _____
6800 _____
PDP-8 _____
4004 _____
360/30 _____
9900 _____
9980 _____
8086 _____

The DEC PDP-11 minicomputer and the Texas Instruments 9900 microprocessor are 16 bitters, that's for sure. As for the others, most people quote the data path width, except for the IBM 360/30 computer which is 32 .... or is it 8?

The confusion arises because three numbers, rather than one, can describe computer hardware. When a single number is given, it could be any one of the three, depending on which you think is most important:

- **Data path width** for those interested in hardware cost tied to the bus width, potential factor of two in speed, or memory organization.
- **Address space** (in bytes) for those interested in the program capacity of the machine; or the difficulty of compiler code generation.
- **CPU register width** for those interested in numeric precision and high performance arithmetic.

Life was simpler about five years ago when we needed to distinguish among only a few machines and before the current generation of "big" little machines added to the confusion.

DATA PATH WIDTH

Data path width is the number most frequently used to classify microprocessors. Most people classify microprocessors such as the 8008, 8080, 6800, and Z80 as 8-bit machines, and the PDP-11, 9900, and 8086 as 16-bit machines. Data path width is important because it is directly related to the hardware complexity required to support a microprocessor bus in a product. In the typical microprocessor application, the data bus interconnects memory as well as peripheral interfaces. As a result, the width of the data bus determines the number of bus lines and influences the complexity of the circuit board layout and the overall hardware cost.

Some newer microprocessors, such as the 8085, multiplex 8 of the 16 address bits with the 8 data-bus bits, reducing the total number of address-plus-data lines to 16. Earlier microprocessors (such as the 8008) were criticized for this approach because of the external logic required to multiplex the bus. The 8085 circumvents this problem by relying on the many available MOS peripheral devices which provide the same conventions and standard functions such as rom, ram, and i/o as an integrated family.

A second reason for the frequent use of data path width is its relationship to the microprocessor's information processing capability. Since all instructions and operands must be transferred one byte at a time over the typical 8-bit data bus, the data bus is a key information path in the hardware system. In this context, a 16-bit data bus has approximately twice the processing potential of an 8-bit data bus, supporting the position that a "16-bit" machine is twice as fast as an "8-bit" machine.

Because of rising hardware costs, the trend toward doubling data bus widths cannot continue. Thirty-two bits is perhaps an upper limit for LSI-intensive systems, and 16 bits is the most cost-effective for many applications. Increases in performance must come from more clever and more problem-oriented designs rather than from simple increases in the width of the data bus.

The effects of data path width on memory organization are also interesting especially for machines patterned after the PDP-11 where the direct-address space is 16 bits. The data path is 16 bits wide and the machine is byte addressed. This gives the curious result that a memory system built from 64K-by-1-bit ram's (scheduled for 1979 availability) will generate a *minimum* memory system of 128K bytes with 16 parts. The system has twice as much memory as the processor can directly address!

ADDRESS SPACE

Address space or "addressability" of a processor is a rough but reliable indicator of the program capacity of the processor. "Addressability" refers to the number of addressable units, typically 8- or 16-bits long, that a program can directly reference. Using this measure, the 8008 is in one class with a 14-bit address.

A second class includes the 8080, 6800, and some "16-bit" processors such as the 9900 and the PDP-11. The microprocessors in this second class require approximately the same number (within a factor of 1.5) of bits to implement a function. For example, the PDP-

11 with its powerful instruction set, requires one-third fewer bits than the 6800 to represent a given function.

In terms of program capacity then, all machines with 16-bit address space are approximately the same. Some processors cheat in this area by employing various forms of bank switching (also called "memory segmentation"). These approaches to circumventing a direct-addressing limit in processors such as the PDP-11, the 8086, and the Z8000, extract a price when the user attempts to program these processors for a single application using a high level language compiler. These machines reflect the 16-bit addressing limitation in the size of the largest array that they can address, and in greater compiler complexity. By this addressability measure, the latest of the microprocessors to be introduced, the M68000, is by far the easiest target machine for a compiler to generate code for.

CPU REGISTER WIDTH

The width of the registers in the processor is the third parameter used to size processors. The width of any register within the processor is closely related to the way it is used. For example, the M6800 contains 16-bit registers and 8-bit registers. These are the processor resources available to the programs and limit the built-in numeric precision of the machine as well as its performance in multi-precision operations. As a contrast, the newly-announced M68000 contains 16 32-bit registers.

When an application requires 32-bit arithmetic operations and the instruction set provides built-in 32-bit operations, there is a reduction in the program space required to represent the operations. In addition, this space reduction will usually lead to a performance increase. Memory accesses frequently limit performance when many primitive suboperations are used to provide an equivalent operation via a subroutine, for example, a multiply operation.

The number and width of registers are important parameters because many applications require 16-bit and 32-bit or longer arithmetic words, and the speed of these operations is directly related to the CPU register length. Some machines can treat registers in pairs, as in the Z8000. In this machine, 32-bit operations can occur between pairs of 16-bit registers thereby achieving some of the same operational characteristics and performance of 32-bit registers, with fewer processor register resources.

SUMMARY

Applications are too divergent and software is much too important to merely group all "16-bit machines" into the same category. Specifically, such a grouping is an injustice to the latest generation of microprocessors (the 8086, Z8000, and the M68000) because it obscures important differences between the M68000 and the other two, in two of the three sizing parameters. □

	Addressing		Data Path Width	Register Width	Comments
	Direct Bits	Max Bits			
8080	16	16	8	8,16	
6800	16	16	8	8,16	
6502	16	16	8	8,16	
Z80	16	16	8	8,16	
LSI-11/2	16	16	16	16	Board level CPU
8086	16	20	16	16,20	Available since 1978
Z8000	16	23	16	16,32	Available 1979
M68000	24	24	16	32	Available August 1979
360/30	24	24	8	32	

Table 1. This table compares nine popular processors in terms of the three size parameters discussed in the article. The main point is that using any one of the three parameters is an over-simplified way of characterizing processors.

PATENT RECEIVED: 4,109,182

SIMULTANEOUS DELAYED SWEEP DISPLAY

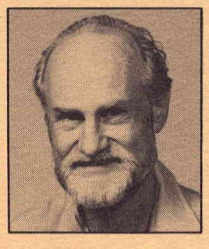

Oliver Dalton, SID Administration, ext. 6574 (Beaverton).

By providing a *simultaneous* display of two delayed sweeps, this invention improves the accuracy of time difference measurements. It further provides a direct readout of time difference between the two delayed sweeps.

FUNCTION

This invention is used in a dual-delayed sweep oscilloscope having both conventional delaying (main) and delayed timebase generators and produces two intensified zones of interest on the delaying sweep display.

Utilizing an alternating sweep sequence, a delayed sweep may then be displayed for each intensified zone. Two separately-controllable delayed-sweep pickoffs alternately gate the delayed-sweep generator; synchronized sweep switching controls the display. A scaled digital multimeter providing direct readouts facilitates accurate time interval measurements between the two delayed sweeps displayed.

THE DISPLAY

In the screen display on page 17, trace 1 is the delaying sweep showing two intensified zones set by the controls. The controls position the intensified zones on trace 1 and also, simultaneously, the selected zones (shown expanded in time) on traces 2 and 3.

Trace 2 displays the first pulse on the delayed sweep, while trace 3 displays the fourth pulse. In this example, the delayed sweep is running 10 times faster than the delaying sweep. The upper readout shows time difference, and the lower readout shows delayed sweep time/division.

The zone-separation control can be used to move trace 3 (and, simultaneously, the second intensified zone on trace 1) with respect to trace 2, thus any point on

Continued on page 17

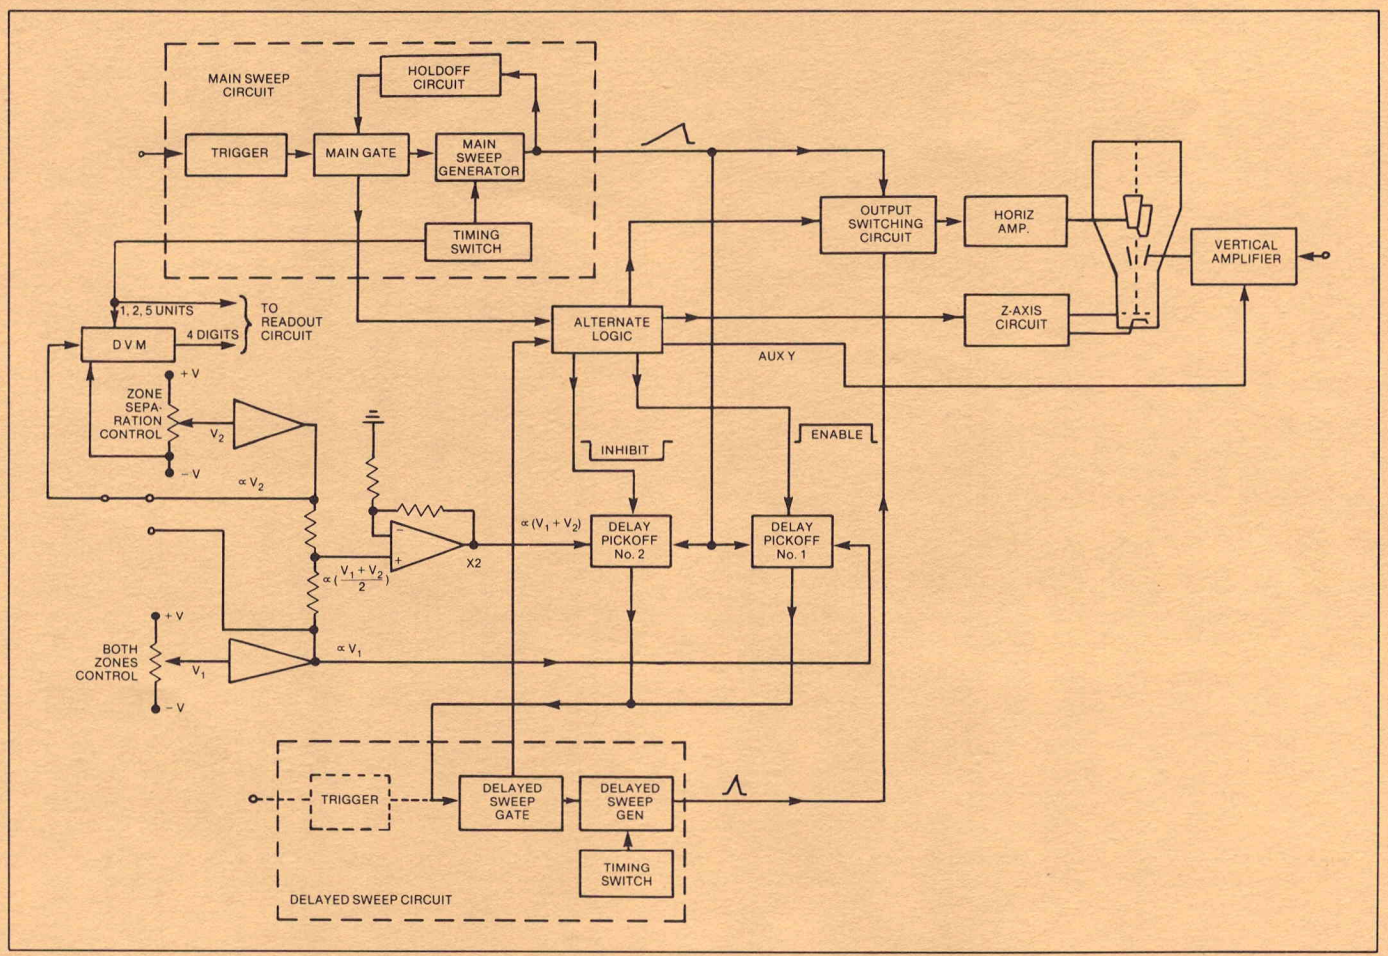

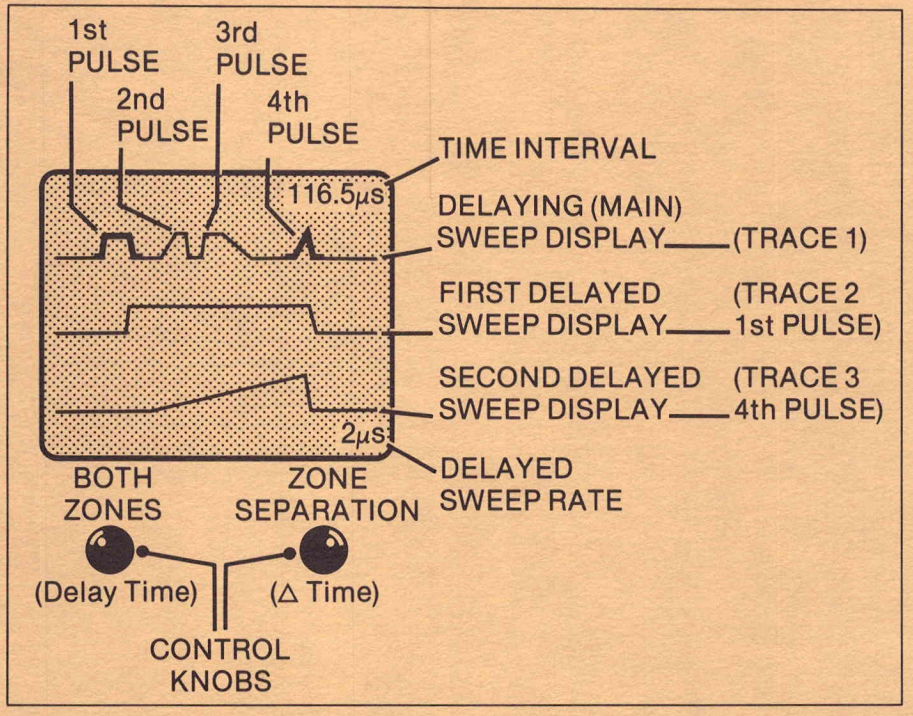

COMPUTER-AIDED MECHANICAL DESIGN ASSISTANCE AVAILABLE

The Scientific Computer Center (SCC) now has a focal point for mechanical engineering computer-aided design; Barry M. Ratihn recently joined the SCC from Oregon State University where he received his masters degree in Structural Mechanics.

Negotiations with Structural Dynamics Research Corporation (SDRC) are in progress to make mechanical engineering software available in-house. This software will be an integrated package for use on the Tektronix MEG 131 (a mechanical engineering graphics, minicomputer-based system) with the Cyber 175. The capabilities of this software are many; the SCC will sponsor (October 16-17) a one and one-half day SDRC seminar to review these capabilities. Call Margene Tung, Education and Training, on ext. 7703 (Beaverton) for costs and sign-up.

Barry will assist engineers who want to use finite element modeling in their projects. The analysis of some current designs has taken Barry about two days per design. The results of the analysis typically consist of deflection plots, stress plots and dynamic mode shape plots. In addition, Barry is available to answer any questions that you may have on the software and hardware needed for mechanical design and analysis.

Call Barry on ext. 5976 (Beaverton) or drop by d.s. 50-454. □

Continued from page 16

either trace can be aligned with any point on the other trace; the readout continuously displays the time difference. A vertical separation control (not shown) can superimpose the two delayed sweep traces, thereby minimizing operator error in taking a reading.

APPLICATIONS

This invention is used in 7B10/15 and 7B80/85 timebase plug-in units. □

technical standards

STANDARDS AID DESIGNERS

John David's article in the July 1979 **Engineering News** presented thoughtful insight into the pros and cons of standards. Except in a totally new situation where a unique solution is the only solution, standards can be of tremendous value.

Standards are developed in response to recurring needs revealed as projects evolve. The new standards then provide criteria for subsequent related needs. Repetitious needs are met faster, more efficiently, and at lower cost, thus freeing designers for more creative work and increasing company profits.

If you see a need for either a new standard or a revision to a standard, fill out a Technical Standards Action Request form. These forms are available from Technical Standards, ext. 241 (Town Center), d.s. 41-260, and at library stations. The Technical Standards group may call on you to provide detailed information and assistance in developing the standard or revision.

For standards assistance or for more information, call Technical Standards, ext. 241 (Town Center). □

NEW TEK SPS BASIC VERSION AVAILABLE

Tek SPS BASIC, a software package facilitating (1) instrument control and (2) acquisition and analysis of waveform data, has been expanded and enhanced in a new version (V02). V02 offers the following features.

•Advanced tasking and error handling capability that prevents errors in one task from terminating other tasks.

• Record i/o for random access to files on any directory-structured peripheral.

•Complete GPIB (IEEE-488 interface) capability.

•More free memory.

•Handling of multiple powerfails.

•Fast overlay commands that dramatically increase the execution speed of programs in which overlay segments increase efficient use of memory space.

•Initialization that frees users from having to answer load-time questions. Users may define a special file to modify the initialization routine to fit their applications.

•Choice of terminal driver that provides more flexible terminal operation.

•Time and date capability that allows scheduling routines and dating files.

•Non-resident command handling that loads commands faster from disk (up to five times improvement on floppy-disk based system).

•Improved graphics capability including x-y plot.

•Added commands for High-Level Support Package that allow bit-level variable manipulation.

Programs written to run under Version 1 of Tek SPS BASIC will generally run under Version 2 without revision. To handle the few instances of syntactic change between versions, users of V02 will receive a conversion program which automatically revises the text of BASIC programs to correct minor incompatibilities and issues warning flags to indicate code requiring user attention.

With the introduction of V02, Tek SPS BASIC V01 software will receive less support than it has in the past. **SPS Programming Update** will continue to report problems, but patches or new releases will not be issued.

For copies of Tek SPS BASIC V02 or questions about its features, call Cindy McMeekin, ext. 1148 (Walker Road), or write d.s. 92-801. □

PATENT RECEIVED: No. 4,135,201

DYNAMIC DAMPING FOR SECAM HIGH-FREQUENCY DE-EMPHASIS

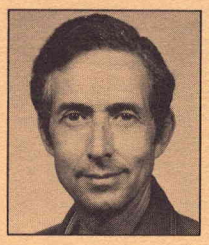

Doug Dickie, Data Acquisition Research (but formerly with TV Products Engineering), ext. 6555 (Beaverton).

Larry Nelson, formerly with TV Products Engineering, has left the company.

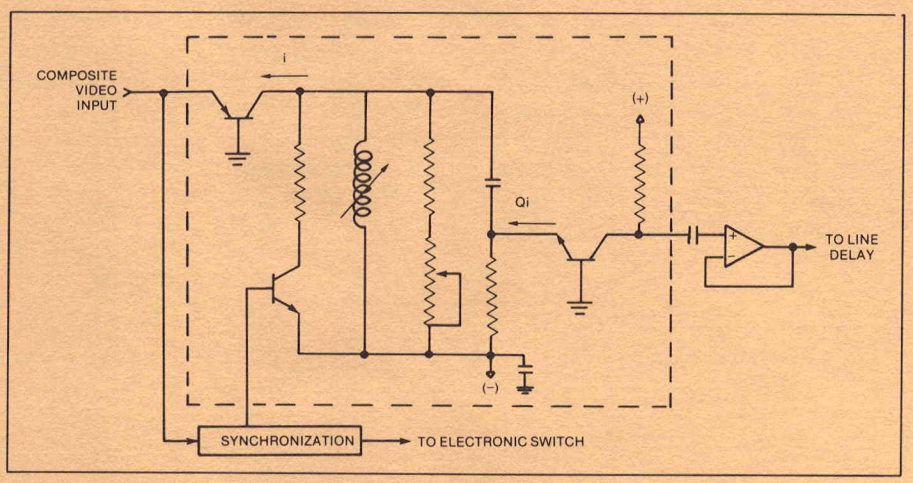

Used in a SECAM (Sequential Color and Memory) tv decoder de-emphasis network, the circuit shown here minimizes the false subcarrier produced when sync pulses ring this network.

For normal operation, SECAM tv monitors require a high-Q (high ratio of reactance to effective resistance) subcarrier de-emphasis network at the decoding (receiving) end of the system (for example, a color picture monitor). In a properly operating monitor, the subcarrier starts after sync pulses appear; however, the sync pulses can cause ringing in the de-emphasis network (it is highly resonant). This ringing creates a false subcarrier. The damping circuit shown here reduces Q during sync pulses, thus damping ringing and minimizing the false subcarrier. □

PATENT RECEIVED: No. 4,151,444

VOLTAGE SWITCHING CIRCUIT IMPROVES COLOR DISPLAY

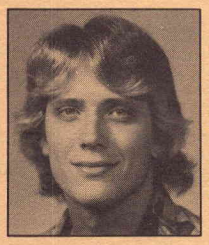

Tim Jenness, Information Display Division Technology Development, ext. 2846 (Wilsonville).

The schematic shown here illustrates a high-speed, high-voltage switch that supplies the anode voltage for a beam penetration crt. This single-gun crt is capable of producing high-resolution color pictures unsurpassed by any other crt technology.

The crt uses combinations of linear and non-linear phosphors such that color shifts are electron beam energy-dependent. Changing the beam's accelerating voltage varies the beam's kinetic energy.

This circuit offers numerous advantages:

Low cost: All driver transistors operate at low voltages. This simplifies the control-signal interface and uses fewer driver transistors. Other manufacturers employ transistor "stacks" (series-connected transistors) for switching on the high voltage side.

High reliability: Because the transistors operate at low voltages, they are less likely to break down.

High speed: While other circuits require up to 200 microseconds to switch colors, this technique yields 10 microsecond switching speeds.

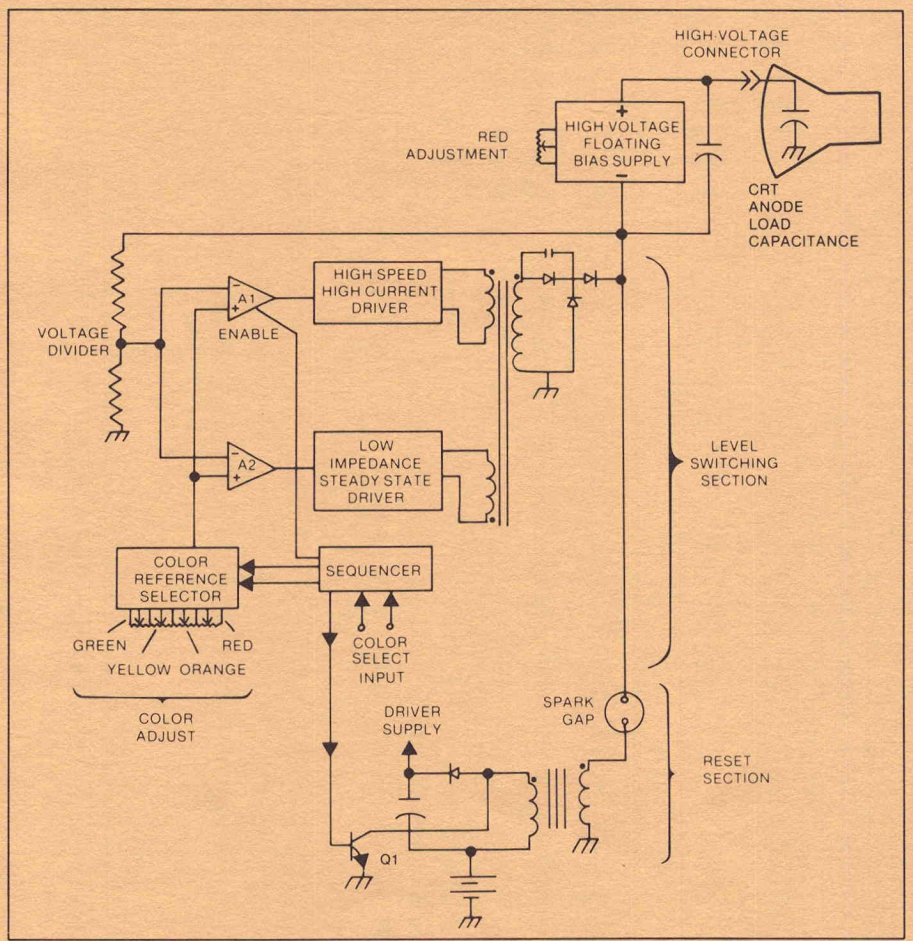

Energy conservation: During descending voltage changes, the energy unloads off the load capacitance and recycles back into a low voltage power supply.

Low weight: The entire package weighs only seven pounds; comparable switches weigh 150 pounds.

The switch circuitry consists primarily of two series-connected supplies. One is a **floating bias supply** adjusted for the lowest output level

Continued on page 20

Technology Review
MAILING LIST COUPON

☐ **ADD**
☐ **REMOVE**
☐ **CHANGE**

Name: _____
Old Delivery Station: _____
New Delivery Station: _____
Payroll Code: _____
(Required for the mailing list)

MAIL COUPON TO 19-313

Allow four weeks for change

"FINAL" ANALYZES DESIGNED-IN SERVICEABILITY/DEVELOPMENT COSTS TRADE-OFF

Carol Golding (Scientific Computer Center) has written a program, FINAL, that performs the analysis described in Jim Geisman's January 1979 **Engineering News** article, "Quick and Dirty Analysis of the Trade-Off Between Designed-In Serviceability and Development Cost."

FINAL implements the procedure described in Jim Geisman's article, thus allowing the user to enter numbers (such as unit cost savings and increases) for a given proposed product. The program pauses at the end of a page to allow the user to make hard copies of the numbers entered.

To obtain FINAL on Cyber (the B machine), use:

ATTACH,FINAL/UN=AB00CBG

If you have questions about using FINAL, call Jim Geisman (Marketing Planning and Strategies) on ext. 6023 or Carol Golding (Scientific Computer Center) on ext. 5976. For a copy of Jim Geisman's article, call ext. 6795 (Technical Communications Services) and ask for the January 1979 **Engineering News.** □

Continued from page 19

that the system delivers. The other is the **level-switching section.**

A **low impedance driver** controls an envelope of a carrier frequency equal to that of the transformer resonant frequency. This circuit maintains the proper voltage on the capacitive load.

A **reset section** unloads energy off the capacitive load and recycles it back into the **driver supply.** Resetting transistor Q1 applies a negative voltage (through the transformer) to the bottom of the **spark gap**. This voltage fires the spark gap. Once fired, the spark gap acts as a closed switch coupling energy through the transformer into the driver supply.

A **sequencer** provides all the timing and commands needed to satisfy the color information specified at the **color select input.** The sequencer addresses reference voltages (from the **color reference selector**) which are applied to error amplifiers A1 and A2. The error amplifiers receive feedback from the voltage divider feedback string.

This voltage switching circuit could be used to improve beam-penetration color refreshed displays and for switched color write-thru direct view storage tubes. □

Scratch Area

WANT TO REACH TR READERS?

A continuing feature of **Technology Review** is the *Scratch Area* column. This space is set aside for miscellaneous short notices such as new personnel introductions, calls for information, and calls for papers. To contribute to *Scratch Area,* send your input to the associate editor, Laura Lane, at d.s. 19-313 (ext. 6795). □

4025 PROBLEMS NEED DOCUMENTATION

Greg Harris (Data Base and Utility Support) reports that the 4025 product group has questions about the use of 4025's in the following modes:

- Form fill-out mode
- Buffered mode
- With tape for off-line data entry

If you use a 4025 for any of the above, contact Greg. The DBUS group needs feedback on (1) how many 4025's are used in the above modes, (2) how much time is spent in these modes, and (3) what problems have surfaced. Also, they are interested in any other interesting 4025 applications you may be doing.

The 4025 product group will use this information when they modify the 4025's firmware. Greg can be reached at ext. 5102 or 5833 (Beaverton), d.s. 50-454. □