TECHNOLOGY report

BRIGHT FLAT ADDRESSABLE

DISPLAY DEVICES

MOVE INTO THE NEXT DECADE

Tektronix
COMMITTED TO EXCELLENCE

CONTENTS

Why TR?

Technology Report serves two purposes.
Long-range, it promotes the flow of tech-
nical information among the diverse seg-
ments of the Tektronix engineering and
scientific community. Short-range, it pub-
licizes current events (new services avail-
able and notice of achievements by mem-
bers of the technical community at
Tektronix).

HELP AVAILABLE FOR PAPERS, ARTICLES, AND PRESENTATIONS

If you're preparing a paper for publication or presentation out-
side Tektronix, the Technology Communications Support (TCS)
group of Corporate Marketing Communications can make your
job easier. TCS can provide editorial help with outlines, abstracts,
and manuscripts; prepare artwork for illustrations; and format
material to journal or conference requirements. They can also
help you "storyboard" your talk, and then produce professional,
attractive slides to go with it. In addition, they interface with
Patents and Trademarks to obtain confidentiality reviews and
to assure all necessary patent and copyright protection.

For more information, or for whatever assistance you may need,
contact Eleanor McElwee, ext. 642-8924. □

WRITING FOR *TECHNOLOGY REPORT*

Technology Report can effectively convey ideas, innovations,
services, and background information to the Tektronix techno-
logical community.

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 re-
view 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 we will have something to work with. Don't worry
about organization, spelling, and grammar. The editors will take
care of that when we put the article into shape for you.

Do you have an article to contribute or an announcement to
make? Contact the editor, Art Andersen, 642-8934 (Merlo Road)
or write to d.s. 53-077. □

Make Them Flat, Make Them Bright, Make Them Addressable, And, When You Can, Make Them Inexpensive
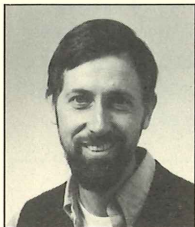
Display Devices Move Into the Next Decade
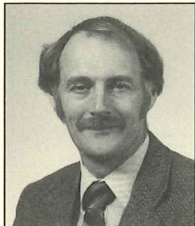
Harry Anderton is the Product Line Manager for Test and Measurement CRT, part of Display Devices Operations. Harry joined Tek in 1978. He received his BS in physics from Bristol University. His PhD in physics is from Emanual College, Cambridge, England.

Steve Blazo has recently joined Accessories Engineering as a Technologist. He is a Principal Engineer. Steve joined Tek in 1969 from Field Emission Corporation (now Hewlett Packard, McMinnville). His BS in physics is from MIT. His MS is from the Oregon Graduate Center.

The most wide-reaching and economically important display technology so far invented has to be the cathode ray tube (CRT). Since Karl Braun invented the CRT in 1897, perhaps a billion or more have presented pictures, graphs, and data. Almost ninety years after Braun's first tube, we now see CRT technology reaching maturity and suffering competition from alternate displays such as flat panels and projection systems. But the CRT will be with us for some time to come.

This article examines the progress of the alternative displays and discusses the CRTs being developed for *niche* markets such as computer-aided design and the aircraft cockpit.

The CRT is nearing technical maturity. Manufacturing costs for "commodity" CRTs are low—$20 for a B&W 9-inch CRT. These costs are mainly labor, materials and overhead, not engineering. Performance advances are being made only for niche markets such as CAD workstations, oscilloscopes, and military displays, and then only at great expense. Only small CRT-performance advances – less than a factor of two – are possible without complex new technologies such as microchannel plate multipliers (MCP) and multibeam cathodes.

Limited performance advances are not unusual for a mature technology. With other mature technologies—such as detergents, automobiles, and furniture—the battles for competitive supremacy takes place in quality, manufacturing, and marketing, not in the lab.

Meanwhile the major alternatives to the CRT, the flat panel and the projection display, are scarcely out of the lab. Early alternatives such as the plasma display, the liquid crystal calculator display, and the CRT projection system have attained significant production volume but have yet to dislodge the CRT from its volume markets.

This is not because the CRT is so good—It's big and heavy; it requires high voltages. The CRT remains King because the newer technologies do not yet have the features nor the low cost necessary to dethrone it.

Who, however, can deny the advantages of carrying a computer in one's briefcase, or hanging the TV set on the living room wall. These advantages will be vigorously pursued.

In this article, we will focus on advances in flat panel and projection displays.

Table 1
Display Component Market Value 1985

Display Device	$M World
All	6500
Cathode-Ray Tubes	
All	5620
Industrial	165
Magnetic:	
Shadow Mask Color	4150
Monochrome	195
Other Color	2
Line Scan	9
Multibeam	1
Electrostatic:	95
Mono Accelerator	11
Post Deflection Accelerator	75
Post Deflection Multiplier	9
Flat Panel	
Liquid Crystal	534
Electroluminescent	22
Plasma	89
Vacuum Fluorescent	147
Cathode Ray Tube	5
Projection	
Cathode Ray Tube	34
Liquid Crystal	1
Oil Film	4

Niche CRTs—CRTs in the Cockpit, the War-Room, and Workstation

Sixty-five percent of all displays use either the single-beam monochrome CRT or the shadow-mask color CRT. Consumer TV and personal computers are,by far, the largest applications.

Although CRT development is still being pursued, advances will be achieved only at great expense and used in only a fraction of the existing market. Consumer products are now so cheap that upgrading the television display would price the end product outside the consumer's expectations. High-definition TV (HDTV), for example, has been practical for many years, but HDTV has still failed to find a major application.

We must look to niche markets to see the real developments in CRTs. These niches will be small in respect to the major markets, but the niches still represent many millions of dollars. The major market developments will be in alternative technologies that posesses significant size, weight and power advantages over CRTs.

Noteworthy CRT-niche developments over the past year include:

- High-resolution shadow mask, high-resolution monochrome, and multibeam for CAD displays

- High-resolution color for avionics displays

- Complex electrostatic CRTs for measurement instruments

- High-brightness monochrome for liquid-crystal color displays

There have also been developments in flat CRTs and projection CRTs that we will discuss.

High-resolution shadow-mask CRTs

The most significant recent advance in high-resolution shadow-mask CRTs has been the 2448 by 2048 dot, 19-inch shadow-mask system developed by Hitachi.[1] The system uses a 0.15-mm shadow-mask pitch, digital correction of misconvergence to 0.1 mm, and 240 MHz video bandwidth. It requires 130 kHz horizontal deflection—about eight times the US TV standard—and a matching low-loss yoke. The layout of the deflection and convergence yokes is shown in figure 2.
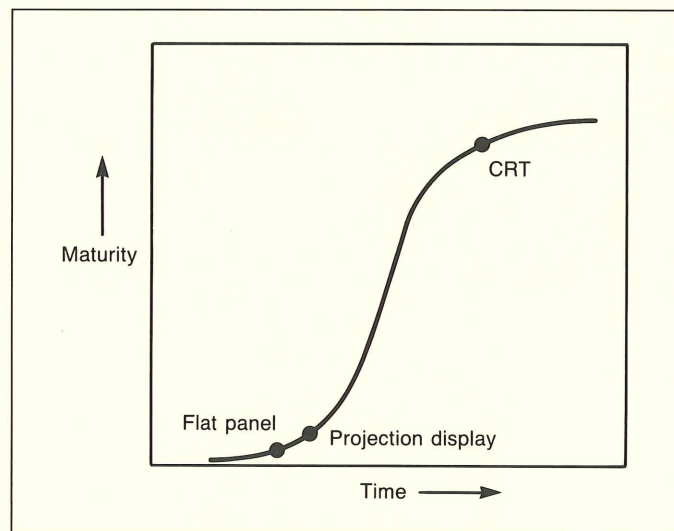

Figure 1. CRT technology is mature, but it is not dead by any means. Its technical maturity will confine improvements in CRT performance improvement to niche markets that are less sensistive to the costs of trying to advance a mature technology. On the other hand...flat panel and projection display systems are just getting started.
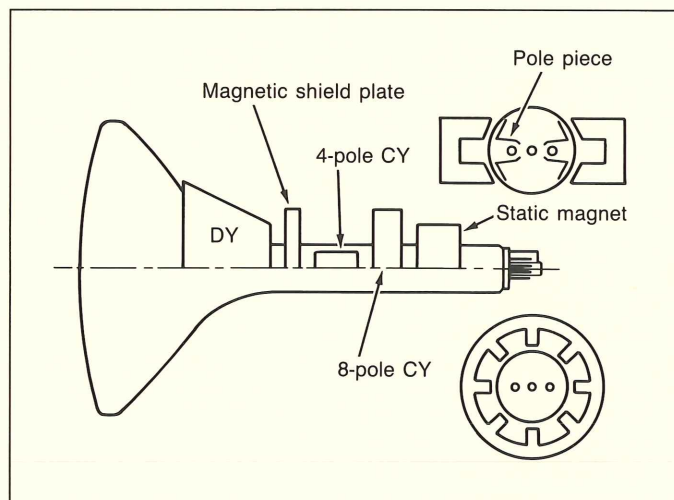

Figure 2. Hitachi has made the most significant recent advance in high-resolution shadow-mask CRTs. This tube's horizontal system can run at 130 kHz. It's vertical is capable of 240 MHz. Convergence is better than 0.1 mm thanks to digitally controlled correction.
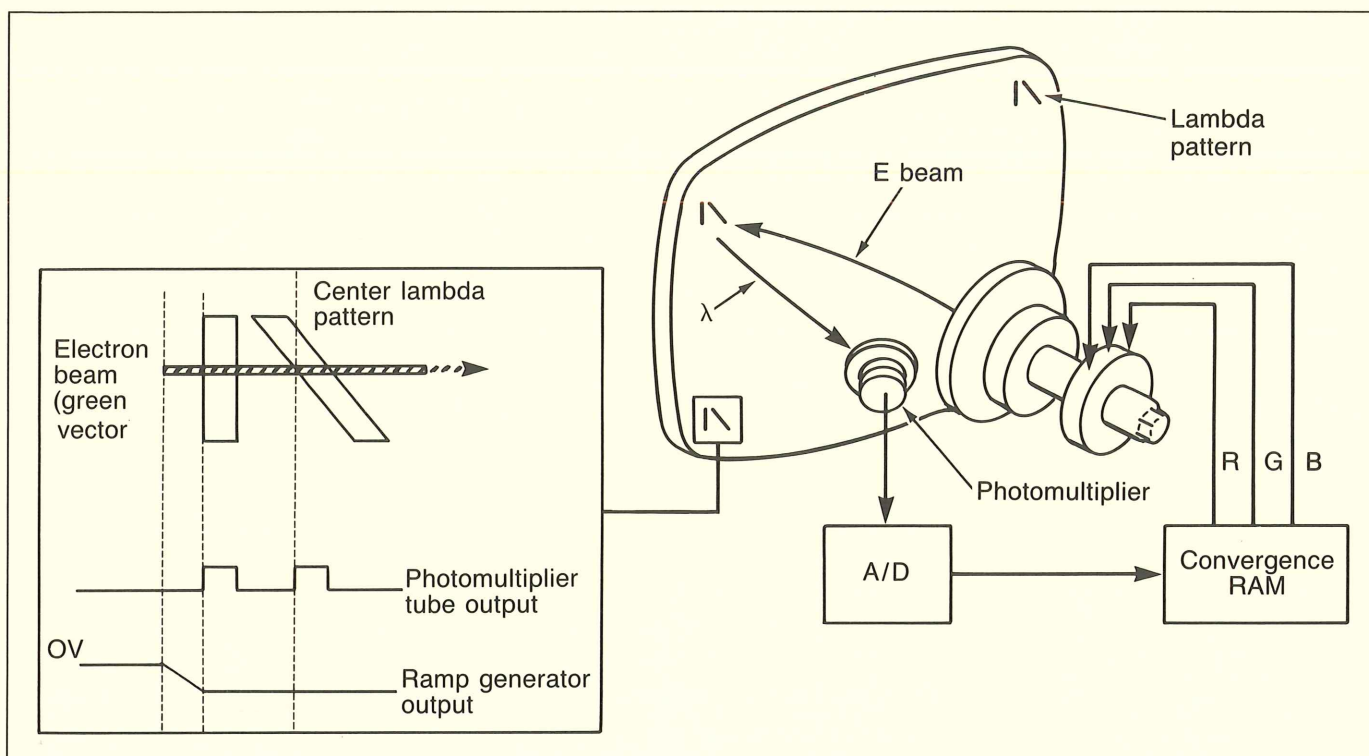
Figure 3. The chevron CRT dynamically corrects convergence errors by employing error feedback through A/D converter to digitally control RGB.

Other interesting developments in high-resolution shadow-mask CRTs over the past few years have been the self-converging yoke,[2] and the chevron feedback convergence system[3] shown in figure 3.

High-resolution monochrome CRTs

CAD often requires large, high-resolution color displays; so does the mapping of dense information. War-room and aircraft cockpit environments demand more than the ubiquitous 19-inch glassware can present.

Consequently, for some of these high-resolution applications, monochrome displays to 30 inches are being perfected. High-efficiency-yoke technology is being used to achieve the high horizontal-line frequency such displays require. This technology may yet make field-sequential color practical, but such CRTs will meet vigorous competition from projection systems, which do not suffer the CRT's size and weight disadvantages. The bulb for a 30-inch-diagonal, flat-face CRT, for example, can easily weigh more than 120 pounds.

Multibeam CRTs

The high-efficiency yoke is one way to overcome the bandwidth and duty-cycle problems associated with large, high-resolution CRTs. Another way is to parallel information channels within the device itself; as many as 16 beams have been reported. Examples are the Litton Line-Scan CRT[4] and the IBM Multibeam.[5]

Although these devices are still in the early stages, they promise significant advantages in that their drive electronics can be simpler and their cathode power loading is less. This suggests that this parallel technique will be pursued to a successful conclusion soon, especially since the liquid-crystal switch[6] offers color from a monochrome CRT without the complexity of a high-resolution target.

There are problems of course. Using the same electron optics to deflect, focus, and position several beams does pose significant challenges. However, today's dynamic- and asymmetric-correction techniques teamed with the sophisticated electron optics that computer-aided design can create should solve these multibeam problems as satisfactorily as problems of the shadow-mask color CRT have been solved.

Color avionics displays

Boeing, by placing a battery of color CRTs in the cockpit of the 757 and 767, has started the trend that is now putting color displays into every cockpit. In aircraft—particularly military aircraft—size, weight, and environment strongly shape the design. Despite its obvious limitations, the shadow-mask CRT is receiving great attention. Tektronix has developed such CRTs. These 5- to 7-inch CRTs readily withstand the heavy shake and shock requirements (figure 4).

Electrostatic CRTs

The electrostatic CRT is primarily used in instruments such as the oscilloscope. Here using electrostatic plates allows for low-power deflection of the electron beam across the CRT face over bandwidths much broader than possible with magnetic-deflection systems. This low-power wide-bandwidth advantage is paid for by gun complexity and poorer electron-optical performance, that is larger spots and smaller deflection angles. As

the beam current is increased to view low-duty-cycle fast events, the trace broadens. Such trace broadening reduces measurement precision.

Two recent technologies enable electrostatic CRTs to increase beam brightness without spreading the spot. The microchannel electron-multiplier (figure 5) and the electron-bombardment induced semiconductor (figure 6), both allow for orders of increased writing rate after the beam has been deflected. Consequently, the need for high-current electron beams and widely spaced deflection plates is eliminated and the electron spot can remain small.

Despite its writing-speed advantages, the electrostatic deflection gun still cannot accommodate deflection angles above 20 degrees, so it is not suitable for most display applications.

Gun base connection

1 Heater
2 Cathode, red
3 Grid 1, red
4 Grid 2, red
5 Grid 2, green
6 Cathode, green
7 Grid 1, green
8 N/C
9 Grid 3, focus
10 N/C
11 Cathode, blue
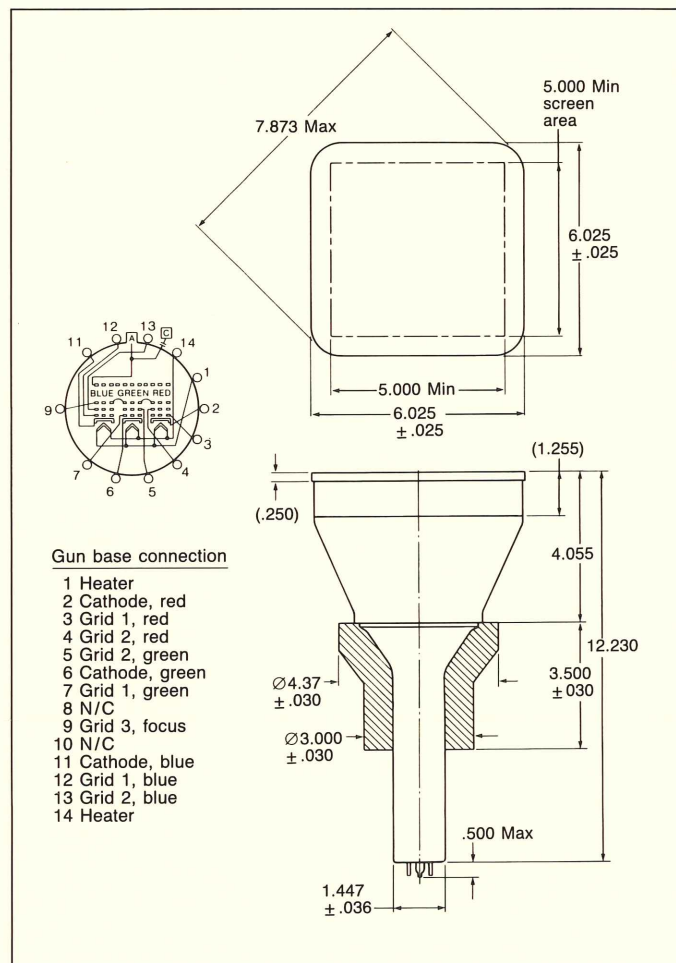12 Grid 1, blue
13 Grid 2, blue
14 Heater

Figure 4. The Tektronix' Avionics CRT is a prime example of the "niche market" CRT. The seven-inch tube shown here produces high-resolution displays bright enough to be readable in direct sunlight. The CRT is very tough, surviving vigorous environmental testing.
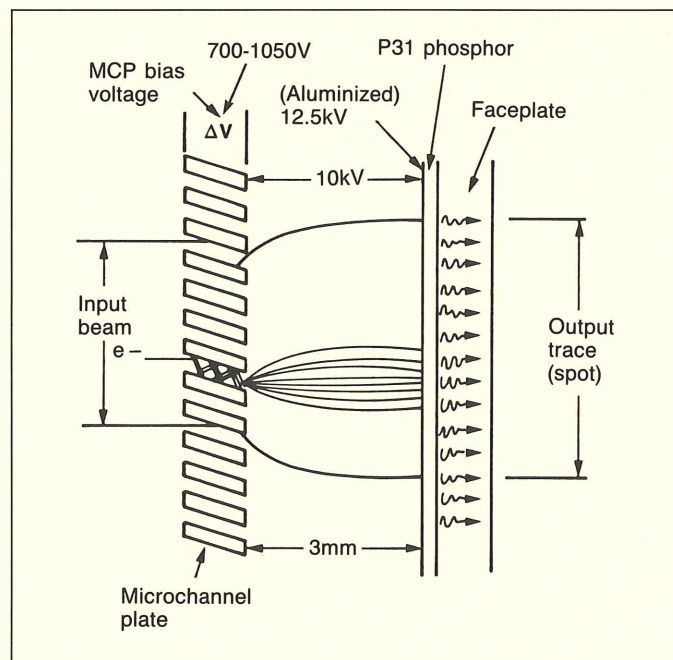

Figure 5. The microchannel plate in this CRT increases writing rate without the trace broadening inherent in increasing beam current. Instead of changing a grid voltage in the electron gun structure, intensity is increased by changing the voltage on the electron-multiplying microchannel plate. The spot diffusing electron optical properties of the widely spaced deflection plates of a conventional electrostatic tube are not present.
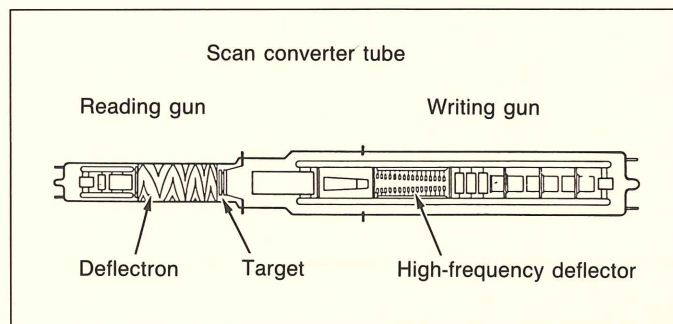

Figure 6. Like the microchannel plate CRT (figure 5), this CRT features a finely focused high-writing-rate spot. The fine spot is achieved by avoiding the use of widely spaced deflection plates. An electron bombardment induced semiconductor forms part of a scan-conversion system.

Flat Displays

For a decade or more, the industry has been predicting flat, hang-on-the-wall TV "within ten years." Although industry has made great gains in flat panels, the ten-year timeline appears little changed.

Instead of displacing CRTs, the flat-panel industry, generally, has been creating products for applications where CRTs are not suitable—hand-held instruments, portable computers, dashboards, battery-operated instruments, sunlight-readable displays, and many more.

Most of these flat-panel displays are monochrome. And they have no grey scale. However, the search for a practical, flat display capable of full-color TV imagery has begun to show some results. For example, Suwa Seikosha has recently demonstrated a 4-inch diagonal, liquid-crystal display, addressed by an active matrix of poly Si thin-film transistors. By operating a color-filter matrix as a light valve, Seikosha achieved full color.

The flat-display market has expanded at about 30 percent per year. Rapid growth will no doubt continue for many years. The major technologies in this industry are liquid crystal, plasma, electroluminescence, vacuum fluorescence, and the flat CRT.

Liquid crystals

Liquid crystals do not generate light; they modulate ambient light. This light may be natural, or can be provided artificially.

The most commonly used liquid crystal displays employ the twisted nematic-field effect, where the display consists of two glass plates with rows and columns of transparent conductors. Sandwiched between the conductor sides of the glass plates is a liquid crystal film 10 to 50 microns thick. The molecules are aligned parallel to the glass plates by evaporating a thin layer of SiO into the glass at an oblique angle or by rubbing the surface.

On assembly, the plates are mounted with the two plates aligned perpendicularly to one another. As a result the molecules exhibit a 90 degree twist across the cell (figure 7). This structure gives light that is plane polarized in the direction of one rubbing axis a 90 degree twist as it passes through the cell. When a voltage is applied across the cell the molecules line up perpendicularly to the glass surfaces, and plane-polarized light passes through unaffected. By placing polarizers on both sides of a display the display designer gives the end-user the option of either a white-on-black or a black-on-white display. Since the effect is entirely due to the field, the power needs can be as low as a few microwatts per square centimeter.

The twist cell typically responds slowly, about 50 milliseconds today. The soft threshold limits the number of lines that can be addressed in a simple matrix, but this limit increases every year; Data General presently markets a computer terminal with 640 by 200 pixels.
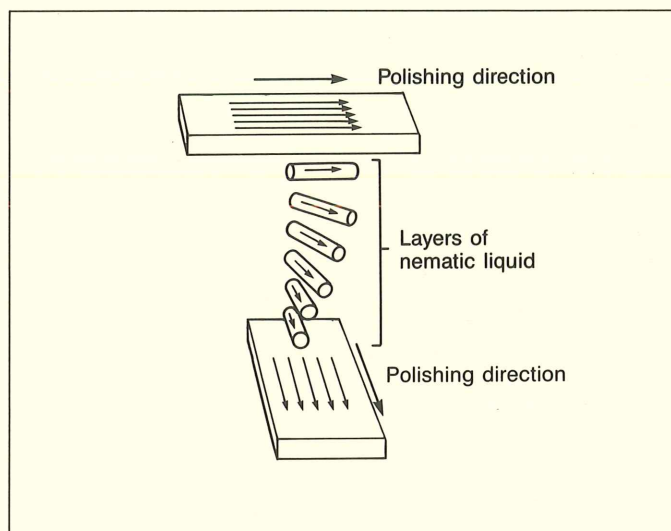

Figure 7. The alignment of molecules in a twisted nematic liquid cystal cell prevents light that is plane polarized in the direction of one rubbing axis from passing through the cell. Applying a voltage across the cell causes the twisted molecules to untwist and line up perpendicularly to the glass surfaces; this allows the plane-polarized light to pass through the cell.

Contrast ratio and viewing angle are typically sacrificed by the twist-cell technique and user satisfaction remains a serious question. The use of a 270 degree twist instead of the common 90 degree twist much improves matrixability. Recently, a display of 540 by 270 addressable dots with good contrast and viewability has been reported.[7] The use of ferroelectric chiral liquid crystals allows bistable storage and rapid (about 100 microsecond) switching, because these crystals have permanent dipole moment. Storage could substantially improve display performance, so we and others are actively researching this technology.

Another approach is the *dye* liquid crystal cell where a pleochroic dye, one in which the molecules exhibit anisotropic absorption, is dissolved in a liquid crystal host. Light attenuation can be continuously varied by changing the orientation of the liquid crystal molecules, providing grey scale. The dye cell needs no polarizer.

One solution to the matrix addressing problem is the active-substrate liquid crystal display, which uses two-dimensional arrays of transistors, one for each pixel. This active solution allows faster switching and grey scale; it also allows the display of much more information. But large, defect-free transistor arrays are difficult to make, limiting the commercialization. However, research is active, especially in Japan.

Thermally addressed *smectic* liquid-crystal displays require both heat and an electric field for activation. A smectic liquid crystal that freezes either clear or cloudy (depending on whether or not a voltage is applied during cooling) can provide a storage display. The image normally is viewed by projection and very high resolution can be achieved. IBM has reported on a 64-million pixel display that uses an array of scanned solid-state lasers to supply the heat.[8]

The advantages of liquid crystal displays—low power, low voltage, passivity, and low cost—outweigh their disadvantages. The use of liquid crystal displays is therefore rapidly increasing.

Plasma displays

Gas-discharge displays produce light by cold-cathode discharge. Gas-discharge devices have a sharp threshold. This allows many lines of numerous pixels to be addressed without crosstalk. A typical plasma-display panel has a matrix of gas-discharge cells defined by two sets of orthogonal electrodes (see figure 8). These electrodes are deposited on two parallel glass substrates. The space between the substrates is filled with a neon-argon gas mixture. A spacing of 0.004 in. (0.1 mm) and a gas pressure of 100 torr are typical. Plasma panels will operate in both AC and DC modes.
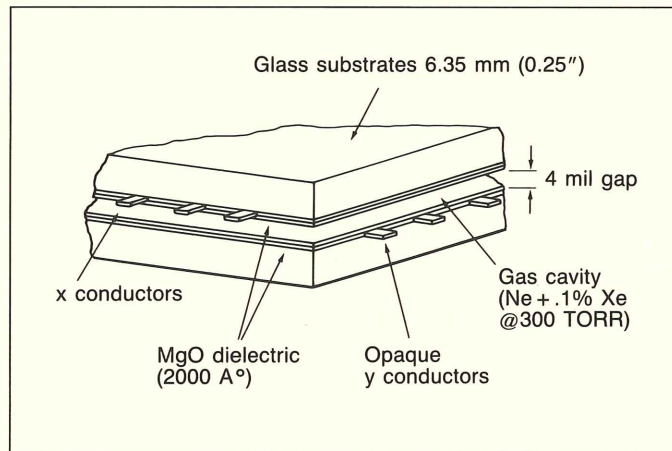

Figure 8. In a typical plasma-display panel, two sets of orthogonal electrodes define a matrix of gas-discharge cells. Displays of this type are free of crosstalk, even when many lines of many pixels are displayed. This absence of crosstalk is a direct result of the sharp thresholds exhibited by gas-discharge cells. Plasma panels will operate in either an AC or a DC mode.

For AC operation, the electrodes are covered with an insulating film. A charge deposited on the insulator during the gas discharge sustains the polarity-reversed applied voltage on alternate half-cycles to maintain gas break-down. Once initiated by a writing pulse, the gas discharge can be maintained with a reduced sustaining voltage and the panel exhibits memory. The insulator eliminates the need for the current-limiting resistors required for bare electrodes.

AC plasma panels have no grey scale because they are bi-stable devices. They do not need display refreshing. In addition, the brightness of an individual pixel is independent of display area. AC plasma display size is limited only by manufacturing constraints such as the ability to manufacture large, fine electrode grids.

AC plasma technology is the basis for some of the largest non-projection displays. Photonics Technology sells display panels measuring one meter diagonally; they plan even larger panels. The recently announced IBM 3290 Terminal uses an AC plasma display of 340 by 274 millimeters (13.4 by 10.8 inches) and 960 by 768 line resolution. This IBM terminal can display 69 lines at 160 characters/line. Numerous halftone patterns are available to give the appearance of grey scale. The display subassembly itself costs $4,500, including associated electronics.

DC plasma panels were widely used in the past, manufactured by Buroughs under the trade names Nixie and Self-Scan. Burroughs recently dropped these products and is concentrating on the AC panel for future products. DC plasma panels typically have no memory and must be refreshed. Display brightness is inversely proportional to the number of lines addressed since each line is *on* only during its fraction of the raster time. Despite these limitations, several companies are looking at DC plasma for TV applications because grey scale is easy to achieve through current or duty-cycle modulation.

Sony is very active in DC plasma and has reported a high-resolution panel that includes a trigger electrode plus the usual x-y matrix. This trigger electrode greatly increases drive-voltage sensistivity and thus reduces the cost of drive electronics.

Large plasma displays, both AC and DC, can be fabricated with good viewability. But high power consumption and cost plus the lack of color (at present) bar widespread use of plasma displays.

AC electroluminescence

AC electroluminescence (ACEL) is light produced by applying a strong electric field to a luminescent material. Various structures will work, but thin-film electroluminescence (TFEL) is the most popular structure (figure 9). Light is believed to be generated by the impact ionization of the manganese ion. The dielectric layers prevent current runaway and promote uniform current distribution. The advantages of TFEL was first demonstrated by Inoguchi[9], who produced electroluminescent devices and ran them more than 20,000 hours at high brightness up to 1000 footlambert (a typical home TV emits about 100 footlamberts) without degradation. The encapsulating feature of the dielectric-film layers is partly responsible for this long life.

In addition to a high peak brightness, an attractive aspect of AC electroluminesence is the extreme nonlinearity of its brightness/voltage curve. This allows many lines to be addressed in a simple matrix, without crosstalk. Peak brightness of several thousand foot lamberts is readily attainable, making displays of several hundred lines practical. A major disadvantage of ACEL devices is their high capacitance and low light-generation efficiency. These, together with high drive-voltage requirements, mean expensive drive electronics.
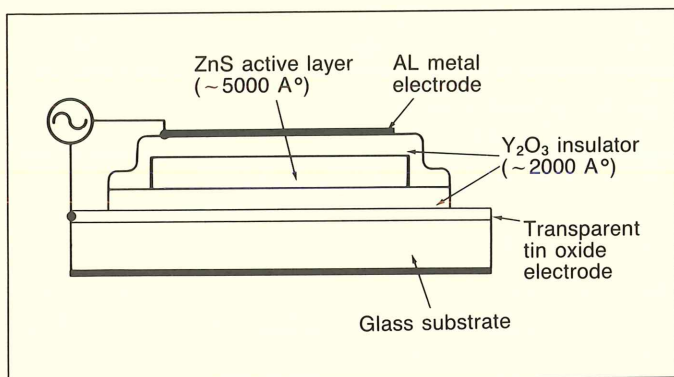
Figure 9. A thin-film electroluminescence device, such as shown, is believed to emit light through the impact ionization of the manganese atom. They are bright and longlasting. One such device ran for more than 20,000 hours at up to 1000 footlamberts without degradation. Like the plasma panel, an AC electroluminescent device is free of display crosstalk.

The size of electroluminescent displays has been limited by the difficulty of manufacturing pinhole-free dielectric layers capable of operating near breakdown voltages. Nevertheless, each year sizes have been increasing and the cost of the panel itself (exclusive of electronics) should be quite low eventually. Electroluminescent panels of 5 by 7 inches with a resolution of 256 by 512 are presently in use. The markets served by plasma panels and electroluminescent panels substantially overlap. It will be interesting to see which technology will dominate.

Vacuum fluorescent

A vacuum fluorescent display typically consists of a matrix-addressed triode structure enclosed in a glass envelope (see figure 10). When positive potentials are applied between the grid lead and the phosphor-coated anode, the phosphor will glow at the intersection and nowhere else. Multiplexing is therefore very easy with vacuum fluorescent devices. Multiple colors can be produced, at a loss in resolution, by using separate red-, green-, and blue-emitting phosphor dots.

The technology has been applied to medium-resolution, medium-size displays, where it is a strong contender because of its superior viewability. It is also used to form the world's largest color television screen, the Sony Jumbotron.[10] In the Jumbotron, each RGB pixel is an individual tube. The "pixels" are assembled into modules and the modules into the full screen. Each resulting Jumbotron image is viewable in full daylight from a distance of a kilometer.
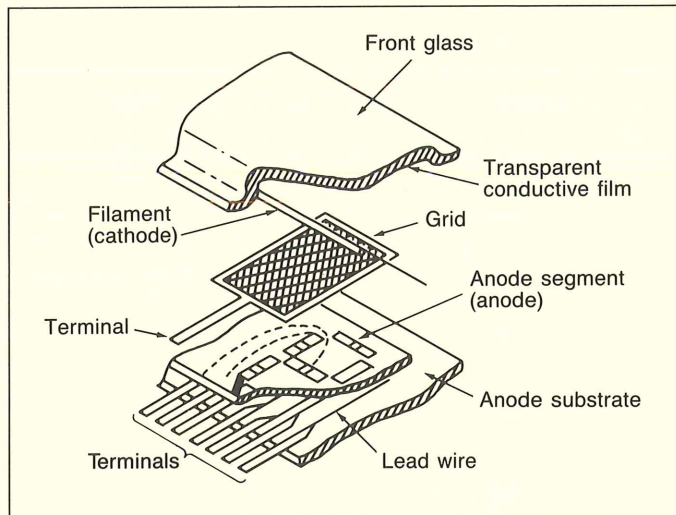

Figure 10. Vacuum fluorescence technology boasts superior viewability. Gigantic outdoor displays are built using a matrix of many individual elements such as the triode device shown here. Sony's Jumbotron can be viewed from one kilometer. In the Jumbotron, each RGB pixel consists of one device.

The flat CRT

CRT producers are adapting their products to meet new needs. Sony[11], RCA[12], Philips[13], and Matsushita[14]) have reduced CRT envelope depths by changing the placement of the electron gun. The exciting recent flat CRTs of Philips and Matsushita both include extra components that improve performance. The Philips CRT (figure 11) contains an electron multiplier, which reduces beam current requirements and simplifies the electron optics. The Matsushita flat CRT (figure 12) uses an array of cathode and gun structures to avoid the long path of the folded beam structure of the Philips and other "flat" tubes. Both techniques can produce color, and high quality examples should be available soon.

Projection Displays

Projection displays have seen considerable development in the past few years. Several types have been investigated, including CRTs, light valves, and laser scanners. Of these, CRT-based systems have made steady progress for the smaller, home-entertainment displays, while light valves have found applications in the larger (commercial) displays.

CRT projection systems

Several manufacturers make $3000-range home-entertainment systems with three to six foot screens. Brightness has been improved dramatically. And resolution, today, is fully compatable with 500-line NTSC television standards. By liquid cooling the faceplate and using phosphors that do not thermally quench, very high electron beam powers have been made practical.
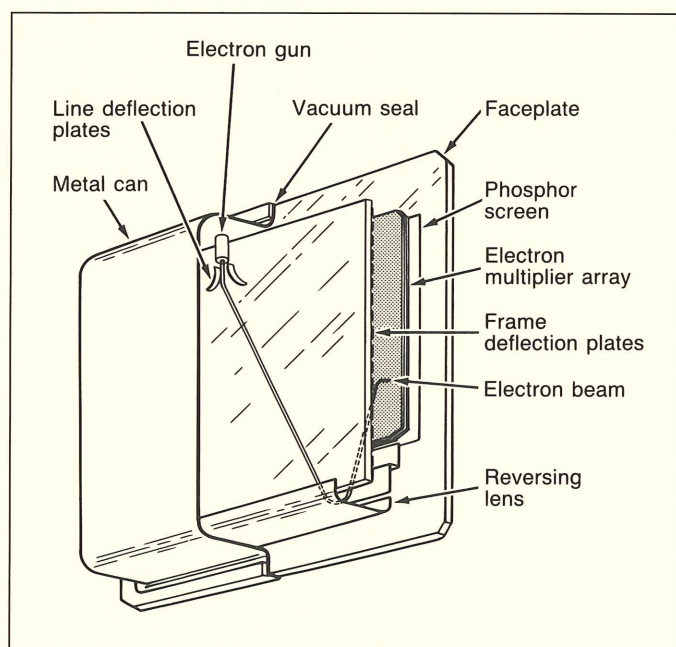
Figure 11. CRTs can be "flat." The gun-at-the-side-technique used here in the Philips flat channel-multiplier CRT is not new. But Philips has reduced beam current requirements by employing an electron multiplier. This has simplified electron optics and produced an exciting folded-beam device.
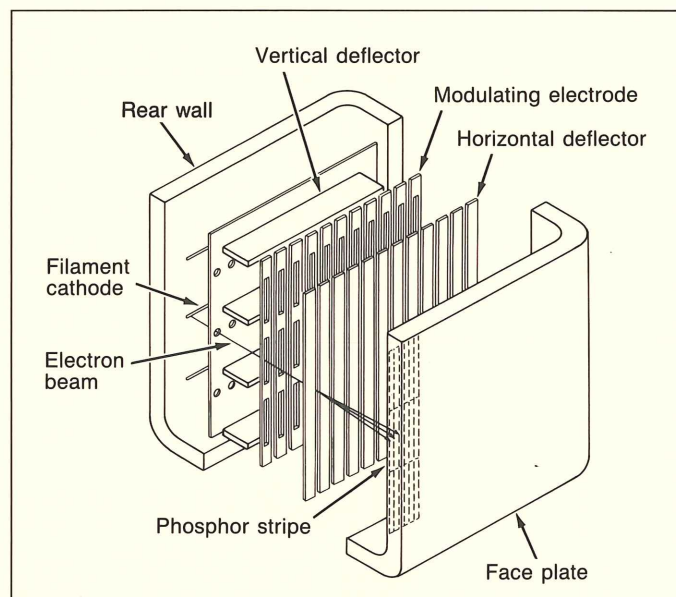

Figure 12. To avoid the folded beam used by Philips and others, Matsuchita uses an array of cathodes and gun structures. This flat CRT is, like the Philips CRT, capable of color and expected to be marketed soon.

Projectors today can emit light in excess of 300 lumens—comparable to a standard 35 millimeter slide projector. These higher-brightness systems, such as the Electrohome single-lens projector and the Sony HDTV receiver, cost considerably more than most home-entertainment models. CRT-based projectors have yet to compete effectively with the light-valve systems in mass-entertainment and commercial applications.

Light-valve projectors

In projection systems, the really exciting developments are in the electron-beam and laser-beam-addressed light valves. For many years, General Electric has lead this field. In their system an electron beam deposits a charge on a thin film of oil causing the film to ripple. A novel optical system (Schlieren) transforms the ripple pattern into a full-color projectable image. This image is very bright and has high resolution.

Although GE's impressive system can project a 1000-lumen 800-line image, it costs $80,000. GE's device is commonly used to project closed-circuit televised sporting events in auditoriums. Improvements are under way to increase resolution and reduce system cost. It appears GE intends to make a version of this system for home-entertainment.

A long-term competitor for the GE light valve is the Eidophor projector, which also uses an oil film but has separate CRTs and optical paths for red, green, and blue information.

Other light-valve systems use liquid crystals for the imaging layer. Hughes is selling a complex CRT-addressed reflective liquid-crystal system (figure 13). Other liquid-crystal systems have been reported by IBM and by Tektronix (figure 14). Such systems are technically difficult but hold great promise for inexpensive large-area displays.
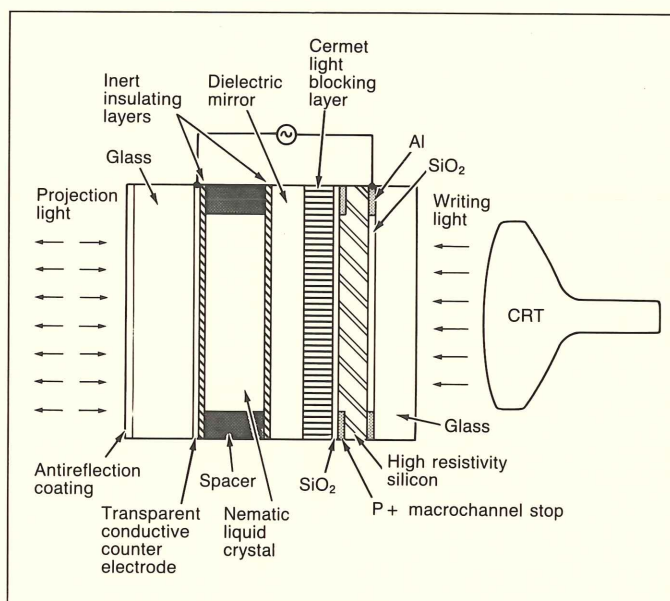

Figure 13. This Hughes light-valve controlled projection system is quite complex. It is being sold now. Other liquid crystal systems have been reported by IBM and Tektronix (see figure 14).
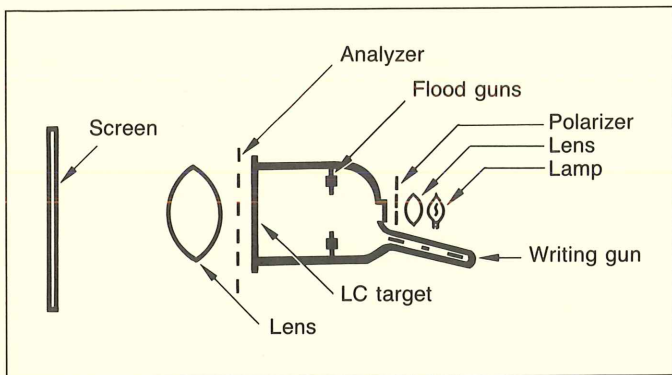
Figure 14. The Tektronix liquid crystal light-valve projection system. It, like its competitors, holds great promise for inexpensive large displays. It is, like its competitors, technically complex.

Another light valve being developed is the matrix-addressed liquid crystal. This technique uses a imaging layer similar to that of Suwa Siekosha's thin-film transistor-addressed liquid crystal. The matrix-addressed liquid crystal, although it suffers from structure problems similar to those of a shadow-mask CRT, appears to have the correct cost trade offs to take it into the mass market of home-entertainment receivers.

Conclusion

As the workhorse of display technologies, the CRT is still firmly entrenched. It holds its major markets because of its solid performance and cost advantages over alternative technologies. But the CRT is "mature" in that it is increasingly difficult to double its performance without incurring order of magnitude increases in complexity and cost. Therefore CRT advances will be limited to niche markets where performance is all important.

Meanwhile, both flat displays and projection displays are making steady progress towards being competitive with the CRT. They have the advantages of low weight, small size, and low-power consumption, considering the display areas achieved. Therefore the markets for flat and projection displays will increase more rapidly than the markets for conventional CRTs.

Thin-film, transistor-addressed liquid crystals, combined with tiny color filters, appear to be the *technology of choice* for many flat and projection displays. Over the next decade it will make its place in the major markets for displays.

For More Information

For more information call Harry Anderton, 627-7620 (50-221) or Steve Blazo, 253-5430 (C1-717). □

References

1. Ando K, Osawa M, Shimizu T, Maruyama T, Fukushima M, "A flicker-free 2448×2048 dot color CRT display," *SID 85 Digest,* (1985), pp. 338-340.

2. Kaashoek J, "Deflection system design for 110° shadow-mask tubes," *Philips Electron Applications Bulletin,* Vol. 32 Number 1, (1973), pp. 3-22.

3. Denham D, Meyer W, "A High Resolution Color Display with Autoconvergence," *SID Japan Display* 1983.

4. Wurtz J, "Multi-beam CRTs get the picture," *Electro-Optical Systems Design,* (September 1982).

5. Beck V.D., Piggins B.P., Uber III A.E., "Multiple beam cathode ray tube with apertured cathode and control grid," U.S. Patent #4301389, (1981).

6. Vatne R, Johnson P, Bos P, "A LC/CRT field-sequential color display," *SID 83 Digest,* (1983), pp. 28–29.

7. Scheffer T.J., Nehring J, Kaufmann M, Amstutz H, Heimgartner D, Eglin P., "24×80 Character LCD panel using the supertwisted birefringence effect," *SID 85 Digest,* (1985), pp. 120–123.

8. Dewey A, Anderson S, Cheroff G, Feng J, Handen C, Johnson H, Left J, Lynch R, Marinelli C, Schmiedeskamp R, "A 64 Million PEL liquid-crystal-projection display," *Society for Information Display Technical Digest,* Vol. 14, (1983), pp. 36–37.

9. Inoguchi T, Takeda M, Kakihara Y, Nakata Y, Yoshida N, "Stable high-brightness thin-film electroluminscent panels," *Society for Information Display Technical Digest,* Vol 5, ibid., (1974), pp. 84-85.

10. Ohkoshi A., "Ultra-large-screen color display," *SID 85 Digest,* (1985), pp. 18–21.

11. Cohen C, "Sony's pocket TV slims down CRT technology," *Electronics International,* (February 1982).

12. Credelle T, Anderson C, Marlowe F, Gange R, Fields J, Fischer J, van Raalte J, Bloom S., "Cathodoluminescent flat panel TV using electron beam guides," *SID 80 Digest,* (1980), pp. 26–28.

13. Washington D, Mansell J.R., Lamport D.L., Knapp A.G., Woodhead A.W., "Progress of the flat channel multiplier CRT," *SID 85 Digest,* (1985), pp. 166–169.

14. Watanabe M, Nonomura K, Ueda M, Kanehisa T, Moriyama Y, Nishida J, "A color flat-panel display using matrix drive and deflection system," *SID 85 Digest,* (1985), pp. 185–186.

From a Glow Many Years Ago, the CRT Came to Be

Herr Plucker was running some electrical discharge experiments, something he had been doing for some time. The centerpiece of his lab was a partially evacuated glass tube made by his friend Geisler. Today he was working with a Geisler's latest, a tube whose air content was down to just 0.00001 atmosphere. In 1859, this was a real hard vacuum, even for technologically advanced Germany, where so many fundamentals of science were being worked out.

After fussing with the peripheral apparatus and connecting his voltage, Plucker turned to look at the tube. It was glowing! Not all over, just part of the inside wall. It was fluorescing, emitting light.

Herr Plucker was sure that "rays" caused this glow, rays originating from the cathode within the tube. It was curious how he could move the glowing area with a magnet. With the magnet, he concluded, he infuenced the rays.

Plucker was the first to knowingly deflect a "cathode ray." In Plucker's movable glow lay the essentials for modern information display.

Plucker and Geisler were not alone in their interests. By the mid-nineteenth century, tens, even hundreds of ardent experimenters were building curious devices of glass and metal. With the help of these contraptions, nineteenth-century men and women were discovering and interlacing the elements of twentieth-century physics. They were laying the foundations for our information-based, technological society.

Without them we would not have the *Display*, at least as we have known it. But then, would we have missed it? Would we have had anything to display?

Thirty-eight years after the glow in Plucker's lab, another German invented the "practical" cathode-ray tube. He wanted to study high-frequency electrical currents. In 1897, if you wanted to study such things, to do science, you had to build your instruments. That's why Ferdinand Braun built the first cathode-ray oscillograph.

In essence, Braun's oscillograph was the cathode-ray tube itself. No electronic amplifier had been invented to drive the the vertical axis. (DeForest wouldn't be sticking a grid in the vacuum-tube diode to make the triode for a few more years.) So Braun couldn't build amplifiers. He couldn't generate a timebase electronically either. But with just deflection coils and a rotating mirror for a "timebase," Braun saw something remarkable. He saw current variations plotted in the time-domain—perhaps the first real-time plot of a fast variable versus time.

Braun saw something remarkable . . . current variations plotted in the time domain.

By adding a second set of deflection coils to his cathode-ray oscillograph, Braun could even plot Lisajou figures.

Just as in most modern displays, in Braun's tube the medium for displaying information was a layer of phosphor, his writing tool a beam of electrons. A tool focused, intensified, and moved with electron optics.

What made Braun's creation for studying high-frequency phenomena practical was the intense work of numerous others. For over half a century, others had been deriving the nature of the mysterious "cathode rays." Whatever they were, they had negative charge. They had properties similar to those observed in experiments with light. The rays could be focused on a layer of phosphor to creat a point of light. They could be deflected either magnetically or electrostatically. They could be used, as Braun proved, to plot one variable against another.

Without Plucker, Hittorf, Crookes, Hertz, Lenard, Perrin, Goldstein, Campbell-Swinton, Fleming, Hess, and other others, Braun could not have built his contribution to twentieth-century science, war, and entertainment. Without Braun's device, would there be a Tek, Sony, Apple, or RCA?
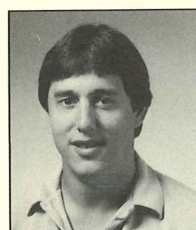
Art Andersen—Editor

Extending Structured Analysis to Become a Design Tool
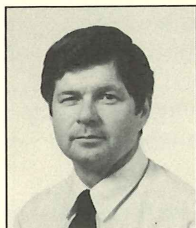
Walter Webb is a software engineering manager in the Software Development Products Division. He joined Tek in 1982. He has worked for Federal Electric Corporation, Aerospace Corporation, AC Electronics, and Autonetics. Walt has an MS in Systems Management from the University of Southern California, an MS in Electrical Engineering from the University of California, and a BS in Electrical Engineering from the University of Santa Clara.
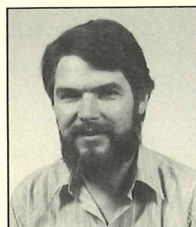
Rainer Wieland is a senior software engineer in the Software Development Products Division. Rainer joined Tek in 1978. He is project leader for the SA Tools Project. Earlier he was a system programmer with Motorola Microsystems. Rainer has an AB in Astronomy from Brown University and has done postgraduate work in the same field at the University of Michigan.

Chris Olson is a software engineer in the Software Development Products Division. He joined Tek in 1983, immediately after receiving his batchelor's degree in Computer Science from Washington State University. For the last two years Chris has been part of the team developing Structured Analysis Tools.

In an extension of the structured-analysis method of writing specifications, program architecture is determined by data-flow diagrams. These data-flow diagrams are defined before the program design process starts. (Interactive-graphic tools are used to define the data-flow diagrams.) The software design is the requirements definition, and the traditional structure-design step is bypassed. The resulting program closely reflects the structured analysis document. This approach was successfully employed in building *SA Tools*, a product of the Software Development Products Division.

Last year, Tektronix began selling SA Tools. SA Tools help software engineers use *structured analysis*.[1] SA Tools itself was specified employing structured analysis. The resulting specification – a structured spec – showed all processes and their data interfaces down to the mini-specification level. The next step was to develop a structured design by applying transform and transaction analysis on the structured specification. The procedure for doing this, however, is ill-defined.

When you carry out transform analysis, remember that it is a strategy. You cannot unthinkingly follow its steps as you could those of an algorithm. From time to time, to stay on the right track, you must bring to bear your knowledge of what the system is supposed to accomplish. And, when you derive your first structure chart, you must use all the design criteria you have learned to improve it.

One day, transform analysis may become an algorithm. But if it does, the structure chart will disappear and the DFD will be implemented directly, for a machine can obey an algorithm much better than can a human being . . . perhaps, we shall see DFDs being executed on a horde of dynamically reconfigurable microprocessors. (From The Practical Guide To Structured Systems Design by Meilir Page-Jones).

Early in the project, the engineers were exposed to the *large-grain data-flow* technique.[3] With this approach, data-flow diagrams directly dictate the architecture of the program. Contrary to the popular convention of developing structure charts from the DFDs, structure charts are never drawn. By adapting the large-grain technique, a new model for program design was created. This was the model used to develop the software for SA Tools.

The New Model for Program Design

In the new model, processes in a data-flow diagram are defined either by a lower-level data-flow diagram or by a mini-specification (mini-spec). The DFDs exist in a hierarchical tree structure where the mini-specs are the leaves on the tree. A mini-spec describes a primitive process in structured English.

The new program-design model is based on the DFD: only mini-spec processes do any work (are executable) and they may execute in parallel. In addition, all data flows are thought of as being single-entry queues that are either full or empty. The essential question is "When should a mini-spec (module) be activated?" How this question is answered follows.

A module is able to perform its task when all of its input queues are full *and* when all of its output queues are empty. The destinations of some output data flows are external to the system; their queues are fixed as always being empty. This approach is similar to techniques proposed by others.[4,5] Some modules do not need all of their inputs to be present or all of their outputs to be consumed. Such modules require state-memory to track which data flows are present or which have been used.

Each module is responsible for filling its output queues and for emptying its input queues. Failure to do so results in a static (or deadlocked) system. In this manner each module can be executed independently based on the state of its inputs and outputs. If no queues change state (are emptied or filled) after all modules have executed, the entire program is done and halts. A single main program controls the execution of all modules.

This model is implemented by associating a boolean flag for each queue. A queue is full if its corresponding flag is *set,* empty if its flag is *clear.* The setting and clearing of these flags is performed using compile-time macros.

The specific application of this model to develop a program is described in the subsequent sections.

From Data Flow Diagrams to Code

Let's start from a structured specification consisting of data-flow diagrams, mini-specs, and a data dictionary. (We know that these documents are consistent and correct.) The examples and descriptions given are based on the use and output of Tek's SA Tools. (See *The Functions of SA Tools* in this issue.)

These tools are grouped into five categories: graphics editing tools, evaluation tools, correction tools, display tools, and auxiliary tools. The cumbersome tasks of drawing, correcting, and verifying the data-flow diagrams are simplified by using these automated tools.

To produce high-level code (C in this case) from the data-flow diagrams the steps shown in figure 1 are followed. Figure 1 itself is a DFD with circles representing automated steps and rectangles representing manual steps. The steps are outlined in Table 1.
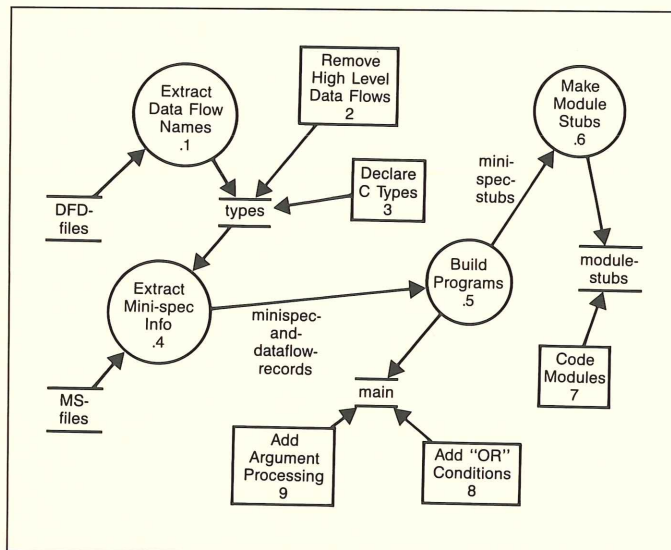

Figure 1. This data-flow diagram (DFD) represents the processes by which the new model converts lower-level DFDs to code. Circles represent automated processes; rectangles represent manual.

Table 1	
Step	Description
Extract Data Flow Names	The data flow names are extracted from all DFDs, sorted and saved in the file *types.*
Remove High Level Data Flows	The data flow names not attached to a mini-spec are deleted. The output data flows that terminate on the boundary of DFD 0 (the top level DFD) are also deleted.
Declare C Types	The C language declarations for the remaining data flows are defined.
Extract Mini-spec Info	The mini-spec body is converted to C language comments. Records indicating the name and type of each data flow and the names and parameters of each module are created.
Build Programs	The main program containing the code to call each mini-spec, the data declarations for all data flows with their associated queue states, and the module stubs containing the parameter declarations and mini-spec body as comments are all created.
Make Module Stubs	The module stubs are split into separate files.
Code Modules	The mini-spec inserted as comments into each module is converted to code.
Add "Or" Conditions	This step is optional. It consists of adding logical "or" conditions to the "if" statements preceding a module's call from the main program.
Add Argument Processing	This step is optional. It consists of adding code to process any command line parameters that are needed by the program.

An example

The technique in Table 1 was used to produce the SA Tools. Each DFD and mini-spec must be in a separate file. A directory must be created containing only those DFDs and *mini-specs* directly involved in generating code. For this example the files for the SA Tools' *lookdd* command are in the current directory:

dfd0	dfd2	dfd9	dfd9.1	ms2.1
ms9.1.1	ms9.1.2	ms9.1.3	ms9.1.4	ms9.1.5

To generate code for one of the other list commands, a different set of lower-level DFDs and mini-specs would be used with the same top-level DFD. The leveled DFDs for the *lookdd* command are shown in figures 2 through 5.

Figure 2 is the top-level DFD for SA Tools' *list* commands. It contains no mini-specs and produces no executable modules.
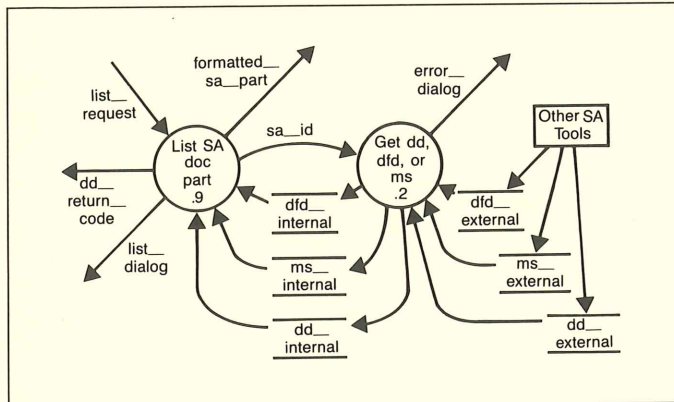
Figure 2. This is the top-level DFD for SA Tools' *list* commands. It contains no mini-specs and it produces no executable modules.
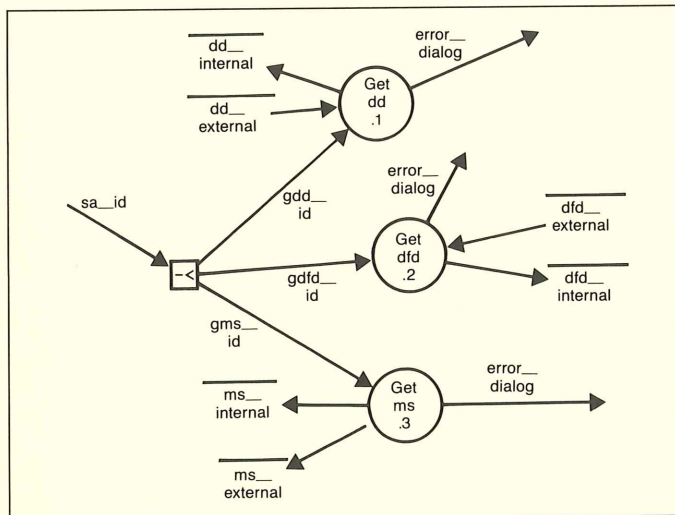
Figure 3. This DFD contains the input modules for the three types of files supported by SA Tools.
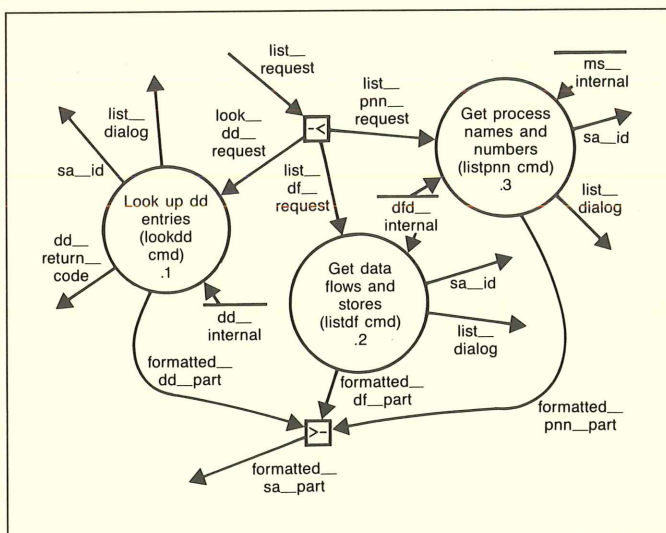
Figure 4. The *list* commands as separate entities. There are no mini-spec processes here as only *lookdd* command mini-specs will expand into code (see figure 5).

Figure 5. The processes of the *lookdd* command. Each process is a mini-spec having a corresponding code module.

From Data Flow Diagrams to Code – Step by Step

The automated steps are composed of standard UNIX commands, SA Tools' commands, and special programs developed to support this code-generation technique. The manual steps are performed using a text editor.

Step 1. *Extract data-flow names* – This step consists of the command sequence:
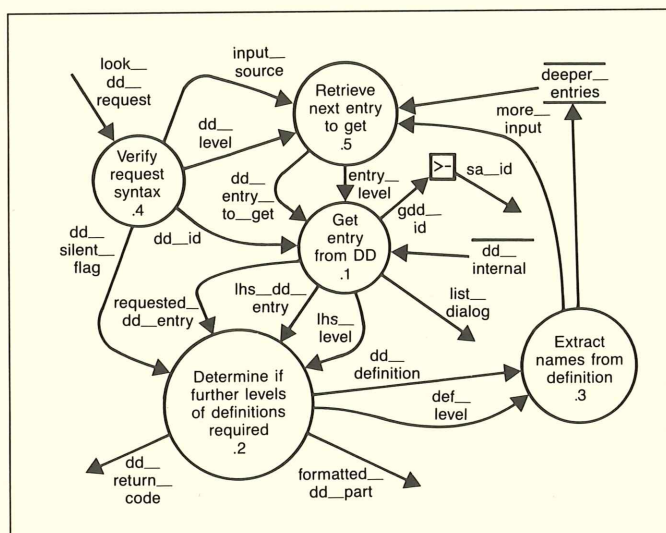
```
listdf dfd* | sort | uniq > types
```

Data-flow names are extracted from all the data-flow diagrams (listdf dfd*), sorted by name (sort), duplicates dropped (uniq), and saved in a file (>types).

Performing this step on the DFDs in figures 2-5 results in a list of all of the data-flow names used in the diagrams. The contents of the file *types* are shown below:

dd__definition	dd__entry__to__get	dd__external
dd__id	dd__internal	dd__level
dd__return__code	dd__silent__flag	deeper__entries
def__level	dfd__external	dfd__internal
entry__level	error__dialog	formatted__dd__part
formatted__df__part	formatted__pnn__part	formatted__sa__part
gdd__id	gdfd__id	gms__id
input__source	lhs__dd__entry	lhs__level
list__df__request	list__dialog	list__pnn__request
list__request	look__dd__request	more__input
ms__external	ms__internal	requested__dd__entry
sa__id		

Step 2. *Remove high-level data flows* – This step deletes all data-flow names that do not flow into or out of a mini-spec. Since only mini-specs are executed, the intermediate-level data-flow names can be discarded. The data-flow names that are outputs (terminate on a boundary point) of the top-level DFD must also be deleted. Such outputs are not cleared by the modules in the current program since they are consumed external to this system.

Step 3. *Declare C types* – In this step, the designer specifies the C data declarations for the remaining names in the file *types.* If the data dictionary was constructed correctly, this information should be readily available. The results of this step are shown below:

dd__definition	char*
dd__entry__to__get	char*
dd__id	char*
dd__internal	struct dd
dd__level	int
dd__return__code	int
dd__silent__flag	int
gdd__id	char*
look__dd__request	char*
requested__dd__entry	char*

Step 4. *Extract mini-spec info* – This step consists of the command sequence:

```
dfdtolist dfd*
```

For efficiency, this step and the next two are usually performed together from a UNIX shell script. We normally put the following into a shell script called *dtoc*:

```
dfdtolist dfd* | awk −f awk.script ; mkproc proc.c
```

The *dfdtolist* program uses the same DFDs as are used in the previous steps. The file *types* as well as all of the mini-spec files for the DFDs are also required by *dfdtolist*. These files need not be specified on the command line that invokes *dfdtolist*; the program accesses them directly. The output from this step consists of C comment blocks, dataflow records, and minispec records for each *mini-spec.* The output is directed to standard output so that it can be piped to the next step.

The comment blocks indicate information derived from mini-spec such as author, date, parent name, as well as the entire mini-spec body.

The *dataflow* record indicates the name and type of a data flow originating or terminating on the mini-spec's process bubble in the parent DFD. Each *dataflow* record has the format:

dataflow	name	flag	C-type

The *name* is the data-flow name. The *flag* indicates whether the data flow is an input or an output. Both fields are derived from the parent DFD. The *C-type* field is taken from the file *types.*

The minispec record indicates the name of a process and all of its required data-flow names. All fields are derived from the parent DFD. Each minispec record has the format:

minispec	file-name	function-name	parameters

file-name is the name of the file from which the mini-spec was read.

function-name is the name of the C function representing this mini-spec. Ideally, the name of the mini-spec process would be the name of the function. However, compilers and linkers severely restrict lengths and function names. Thus, the convention was adopted to name each function with a *p* following by the process number (periods being replaced by underscores).

parameters is part of the names of all data flows used by the mini-spec. Again, the name of the data flow should become the name of the parameter. However, hyphens must be converted to underscores in order to disambiguate parameter names from arithmetic expressions.

A sample of the output from *dfdtolist* is shown below.

```
/*
*************************************************
*
* DFD        - 9.1 - Look up dd entries (lookdd cmd)
* MINI SPEC  - 9.1.1
* TITLE      - Get entry from DD
* AUTHOR     - rainerw
* DATE       - 6/26/85
*************************************************
*/

dataflow     dd_id    i    char  *
dataflow     dd_entry_to_get i   char  *
dataflow     dd_internal  i    struct dd
dataflow     requested_dd_entry   o    char  *
dataflow     gdd_id   o    char  *
minispec ms9.1.1 p9_1_1 gdd_id requested_dd_entry    dd_internal    dd_entry_to_get dd_id
     /*
     *************************************************
     *
     * Get entry from DD
     * rainerw
     * 6/26/85
     * 9.1.1
     *
     * Repeat {
     *    if left hand side of dd entry != dd_entry_to_get {
     *          skip this dd entry.
     *    } else {
     *          SET requested_dd_entry = right hand side of dd entry.
     *          SET lhs_level = entry_level.
     *          SET lhs_dd_entry = left hand side from a dd entry.
     *          CLEAR dd_entry_to_get.
     *          CLEAR entry_level.
     *          Return.
     *    }
     * } until the entire dd_internal has been read one time.
     *
     * if (dd_silent_flag = FALSE) {
     *    print "Name <dd_entry_to_get> not found in DD <gdd_id>".
     * }
     *
     * CLEAR dd_entry_to_get.
     * CLEAR entry_level.
     * Return.
     *
     *************************************************
     */
```

Step 5. *Build programs* – This step consists of an *awk* program that produces the main program and the module stubs from the output of the previous step. The main program contains the data declarations for all data flows, the data declarations for all flags associated with each data flow, and the main loop that calls each mini-spec process in turn. The main program generated for the example is shown below. The include-file references were automatically generated for this application.

```
/*
* Main loop
*/
#include     "flag.h"
#include     "io.h"
#include     "error.h"
#include     "globals.h"

main (argc, argv)
     int   argc;
     char  *argv[];
{
     FLAG   loop_flag;

     do {

     loop_flag = 0;
     if (IS_SET(Fdd_internal)) goto skip001;
     if (IS_CLEAR(Fgdd_id)) goto skip001;
     p2_1 ();
     loop_flag = 1;
skip001:
     if (IS_SET(Fgdd_id)) goto skip002;
     if (IS_SET(Frequested_dd_entry)) goto skip002;
     if (IS_CLEAR(Fdd_internal)) goto skip002;
```

```
     if (IS_CLEAR(Fdd_entry_to_get)) goto skip002;
     if (IS_CLEAR(Fdd_id)) goto skip002;
     p9_1_1 ();
     loop_flag = 1;
skip002:
     if (IS_CLEAR(Frequested_dd_entry)) goto skip003;
     if (IS_SET(Fdd_definition)) goto skip003;
     if (IS_SET(Fdd_return_code)) goto skip003;
     if (IS_CLEAR(Fdd_silent_flag)) goto skip003;
     p9_1_2 ();
     loop_flag = 1;
skip003:
     if (IS_CLEAR(Fdd_definition)) goto skip004;
     p9_1_3 ();
     loop_flag = 1;
skip004:
     if (IS_CLEAR(Flook_dd_request)) goto skip005;
     if (IS_SET(Fdd_silent_flag)) goto skip005;
     if (IS_SET(Fdd_level)) goto skip005;
     if (IS_SET(Fdd_id)) goto skip005;
     p9_1_4 ();
     loop_flag = 1;
skip005:
     if (IS_SET(Fdd_entry_to_get)) goto skip006;
     if (IS_CLEAR(Fdd_level)) goto skip006;
     p9_1_5 ();
     loop_flag = 1;
skip006:
     ;
     } while (loop_flag);
}
```

The include file *flag.h* contains the global data declarations for each data flow and its associated queue-state flag. All queue states are automatically initialized to empty (FALSE).

```
FLAG    Fdd_definition = FALSE;
FLAG    Fdd_entry_to_get = FALSE;
FLAG    Fdd_id = FALSE;
FLAG    Fdd_internal = FALSE;
FLAG    Fdd_level = FALSE;
FLAG    Fdd_return_code = FALSE;
FLAG    Fdd_silent_flag = FALSE;
FLAG    Fgdd_id = FALSE;
FLAG    Flook_dd_request = FALSE;
FLAG    Frequested_dd_entry = FALSE;
char    *dd_definition = {NULL};
char    *dd_entry_to_get = {NULL};
char    *dd_id = {NULL};
char    *gdd_id = {NULL};
char    *look_dd_request = {NULL};
char    *requested_dd_entry = {NULL};
int     dd_level = {NULL};
int     dd_return_code = {NULL};
int     dd_silent_flag = {NULL};
struct dd   dd_internal = {NULL};
```

Each module stub contains the correct external data declarations for the data flows used by a module, and the rudimentary C statements to make the file compileable. Even though all of the data flows are global variables in this implementation, each module can only access those data flows that are directly attached to its process bubble since other data flows are not explicitly declared. The following is an example of a *mini-spec* stub.

```
#include     "io.h"
#include     "error.h"
#include     "globals.h"

/*
*************************************************
*
* DFD        - 9.1 - Look up dd entries (lookdd cmd)
* MINI SPEC  - 9.1.1
* TITLE      - Get entry from DD
* AUTHOR     - rainerw
* DATE       - 6/26/85
*************************************************
*/

/*
* FLAGS
*/
extern FLAG   Fdd_id;
extern FLAG   Fdd_entry_to_get;
extern FLAG   Fdd_internal;
extern FLAG   Frequested_dd_entry;
extern FLAG   Fgdd_id;
```

```
/*
 * GLOBALS
 */
extern char    *dd_id; /* i */
extern char    *dd_entry_to_get; /* i */
extern struct dd    dd_internal; /* i */
extern char    *requested_dd_entry; /* o */
extern char    *gdd_id; /* o */

p9_1_1()
{
    /*
    •••••••••••••••••••••••••••••••••••••••••••••••••••••••••••••••
    •
    • Get entry from DD
    • rainerw
    • 6/26/85
    • 9.1.1
    •
    • Repeat {
    •     if left hand side of dd entry != dd_entry_to_get {
    •         skip this dd entry.
    •     } else {
    •         SET requested_dd_entry = right hand side of dd entry.
    •         SET lhs_level = entry_level.
    •         SET lhs_dd_entry = left hand side from a dd entry.
    •         CLEAR dd_entry_to_get.
    •         CLEAR entry_level.
    •         Return.
    •     }
    • } until the entire dd_internal has been read one time.
    •
    • if (dd_silent_flag = FALSE) {
    •     print "Name <dd_entry_to_get> not found in DD <gdd_id>".
    • }
    •
    • CLEAR dd_entry_to_get.
    • CLEAR entry_level.
    • Return.
    •
    •••••••••••••••••••••••••••••••••••••••••••••••••••••••••••••••
    */

    BEGIN
    STATE(STATE0)
    END
}
```

Step 6. *Make module stubs* – This step consists of the following command sequence:

    mkproc filename

The program *mkproc* splits up *filename* – the module stubs produced by the previous step – into separate files. One file is created for each module. Having each module stub in a separate file permits better management of the system components and lets the user take advantage of UNIX utilities such as *make.*

Step 7. *Code modules* – In this translation step the mini-spec is coded from the algorithm described by the *mini-spec* body. The body of the *mini-spec* has been put into each module file as a comment to aid translation.

In some instances the module for a *mini-spec* may have multiple internal states. The states are a means of introducing control inside the module. Multiple states are needed when a module is to control the sequence of execution of other modules. Macros are used to define states and state transitions. This allows the source code to remain readable.

Step 8. *Add "Or" conditions* – This step is optional. Some modules must execute even with some inputs not set. When this is so, the conditional invocation in the main program for such modules must be modified by adding a logical "or" to the list of conditions preceding the module's call.

Step 9. *Add argument processing* – This step is optional. It is required when the main program must obtain user-supplied parameters from the invoking command line. In this case, the designer must supply the code required to process the command line.

Advantages of This Technique

Generating code from the DFDs ensures that the specification is very close to the final code in the implemented product. If the specification is correct, the implementation will be correct.

The use of compile-time macros for module entry, module exit, module-state control, and queue-state control makes it easy to add or subtract debug hooks into or from various parts of the system. The macros simply need be changed to include the desired debug print statements.

The laborious conversion of DFDs to structure charts is skipped. This saves considerable time. It also preserves the original information about the system. Usually DFDs are discarded after structure charts are drawn. This does not happen here.

Disadvantages of This Technique

Reading code without the original DFDs is difficult. To understand the code, you must have the specification because the hierarchical nature of DFDs is invisible in the code itself. It is possible to reconstruct a flat single-level DFD from the main program, but the result is messy (much like a flat, detailed structure chart).

As with all new techniques, people have to learn how to use it. Maintainers of a product developed with this technique must also understand SA.

Another disadvantage is that not all steps are automatic. Thus, changes are still made to the code rather than in the specification. If all steps were automated, only the specification would need to be changed and the code would simply get regenerated.

Future Work

Many paths can be followed from here to extend the advantages and to reduce the disadvantages of this scheme. Eliminating the manual steps from figure 1 seems like an obvious next step. High-level data flows could be removed from the file *types* without too much trouble. Data declarations, if contained in the data dictionary, could be automatically extracted. If mini-specs were written with more structure, the translation of mini-specs to code would be easier. "Or" conditions could be placed directly into the DFD with a graphics editor permitting correct mini-spec invocation conditions to be generated the first time.

Work has also been started on animating a DFD to monitor the execution of a program.

Summary

There are significant implications for using this approach to develop programs for computers with multiple central processors. It would be possible to have each module execute on a separate hardware processor. In this way, CPU intensive programs could execute much faster.

For More Information

For more information call Rainer Wieland, 629-1029 (92-525). □

References

1. DeMarco, Tom, *Structured Analysis and System Specification,* Prentice-Hall, Englewood Cliffs, New Jersey, 1979.

2. Page-Jones, Meilir, *The Practical Guide to Structured Systems Design,* pp. 182-183, Yourdon Press, New York, 1980.

3. Babb, Robert G. II, "Parallel Processing with Large-Grain Data Flow Techniques," *Computer,* vol. 17, no. 7, pp. 55-61, 1984.

4. Peterson, J.L., *Petri Net Theory and the Modeling of Systems,* pp. 18-21, Prentice-Hall, Englewood Cliffs, New Jersey, 1981.

5. Pugh, J.R., "Actors Set the Stage for Software Advances," *Computer Design,* vol. 23, no. 10, pp. 185-189, Sept. 1984.

The Functions in *SA Tools*

SA Tools support structured analysis (SA) by automating the routine tasks of specifying a system. The Tools let designers create, analyze, modify, and display a specification of the system to be developed.

SA methodology defines a system using three types of documents. First, there is a hierarchy of data-flow diagrams (DFDs). A DFD graphically shows the flow of data and the processes that transform the data. Starting with a DFD showing the overall system, each process is further structured by a more detailed DFD or described by a mini-specification, the second type of document. A mini-specification describes (in structured-English, for example) the transformation of data at the lowest level considered. This hierarchy of DFDs and mini-specifications reflects the structure of the system being developed.

The third type of document in an SA specification is the data dictionary. It contains definitions of all the data items used in the DFDs and mini-specifications.

Now let's look at the functions in SA Tools.

In using the *editing function,* a special-purpose interactive graphics editor lets designers create and modify DFDs on-line. The editor commands can create, label, move, and delete each item on a DFD.

Designers can also tailor the appearance of DFDs and use the graphics capabilities specific to the terminal. The *visit* commands let the user move through the SA specification hierarchy without leaving the DFD editor. The *shell* escape command provides the capability to temporarily leave the editor to do other tasks.

The data dictionary and mini-specifications can be edited with any text editor, so designers can use the editors they are most familiar with.

The *evaluation function* verifies the consistency and complete-ness of the SA specification. The *evalsa* command analyzes the SA documents both individually and with respect to one another. It identifies inconsistencies between documents and omissions in a document for adherence to SA guidelines. For example, the *evalsa* command checks the consistency of data flows between a DFD and its parent, checks that the data dictionary defines all data used in DFDs and other data, and checks DFDs for unconnected and read- or write-only data items.

The *correction function* preserves consistency in the SA specification as it is developed. The *fixsa* command automatically maintains consistency of process names and numbers in the specification. This command creates new documents with the proper name and number, and corrects documents with incorrect names and numbers. It also maintains consistency between data dictionary entries and the data names in DFDs by adding or deleting data dictionary entries to match new or changed DFDs.

The *display function* is used to show the SA specification on-line or to produce a hard copy. The *showsa* command displays formatted SA documents on the terminal.

Designers can also produce hard copies of data-flow diagrams quickly on Tektronix copiers and plotters. Copies can be full or half size, and in color or black and white.

SA Tools provide *auxiliary functions.* These are used for other modifications of the specification. These commands perform tasks such as sorting the data dictionary or listing the processes and data used in the SA documents.

SA Tools can be combined with one another and with standard TNIX or UNIX commands to perform the more-complex specification tasks. For example, the *listpnn* command can be used with the UNIX *sort* and *uniq*] commands to produce an overview of the hierarchy of DFDs and mini-specifications in the SA specification. □

Display Filters Offer Variable Viewability
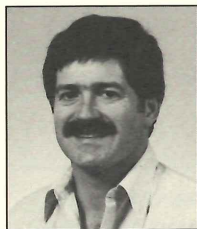
Bob Hubbard is a senior scientist in the Imaging Research Laboratory, part of Tek Labs. Bob joined Tek in 1978. Earlier, he worked for Texas Instruments. Bob's PhD in organic chemistry is from Texas A & M University.

Contrast-enhancing filters have been employed for many years. Many are simple and effective. However, these simple filters are fixed. Now a filter is available whose density can be varied over a continuous linear range. The characteristics of the filter itself can be changed to meet a range of applications.

Variable density filters are Tek developed liquid-crystal devices intended for controlling the *viewability* of CRT displays.

With these voltage-controlled devices, contrast levels can be dynamically varied and incident light rejected without degrading resolution. The amount of light transmitted from the CRT can be adjusted to match the CRT's characteristics and the ambient conditions. Because contrast and transmission are voltage controlled, these characteristics are readily adjustable by manual or automatic techniques.

To match a variety of applications, the Liquid Crystal Shutter SPU provides its customers with a choice of two basic filters:

- The Variable Density Filter (VDF)—The VDF is a simple variable density filter available in a range of variable densities, in either a gray (neutral) or a color.

- The Variable Circular Polarizer (VCP)—The VCP combines a variable density filter with a quarter-wave plate to substantially reduce viewing problems caused by ambient light.

Transmission Characteristics

Variable density filters are liquid-crystal cells whose optical density is variable over about a thirty-five percent transmission range. The cell is darkest with no voltage applied and less dark with about 10 volts applied (see figure 1). Contrast is changed by adjusting the voltage.

The light transmitted through such filters can be a level of white (neutral gray) or blue, red, green, and other colors to suit the application.
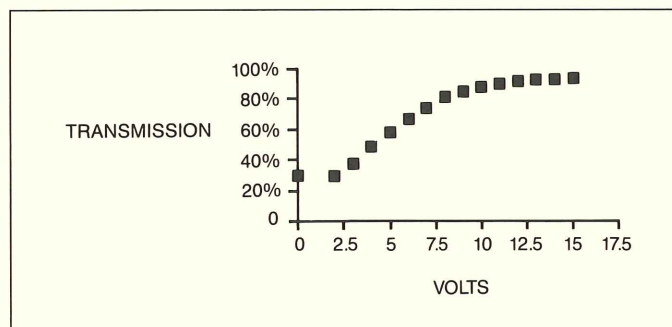

Figure 1. A variable density filter (VDF) changes density with applied voltage. This change of density changes the amount of light transmitted.

In manufacturing these filters, density can be adjusted by varying the concentration of the dye in the liquid crystal or by varying the cell thickness. The effect of varying the dye concentration can be seen in figure 2 in which T_{off} represents the cell with no voltage applied and T_{on} represents the cell with full voltage applied. By varying the voltage applied, the transmission can be varied within the thirty-five percent range.
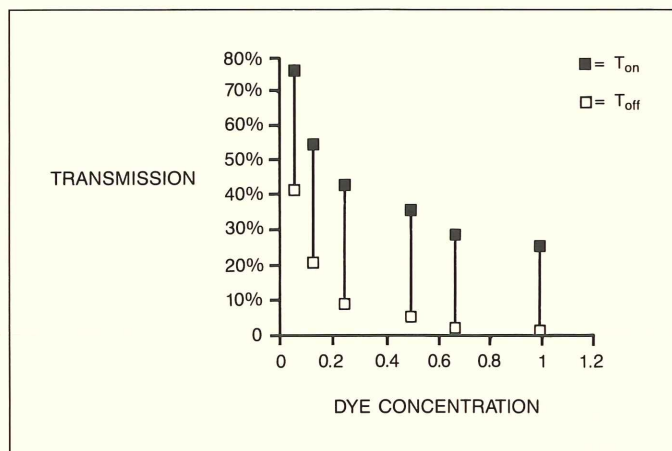

Figure 2. At low dye concentrations, the range of voltage-controlled change is about thirty-five percent (represented by the vertical excursions between "on" and off"). With higher dye concentrations, less light gets through and the range of density change is somewhat reduced.

Increasing cell thickness also reduces light transmission. And as thickness increases, the range of transmission variability decreases, particularly above 10 microns.

The optimum devices are cells of 5 to 10 micron thickness having dye concentrations of less than 50 percent. Applications dictate the final parameters. The 10 micron spacing allows a large margin for processing error, but 5 micron spacing is more compatible with the LCCS production line.

To produce a highly transmissive neutral (gray) filter, it is possible to make a thin low-dye-concentration cell with a transmission range of about 50% to 80% in the off and on states respectively. For a low-transmission device it would be necessary to use a VDF with a more concentrated dye mixture (see figure 2).

For a variable circular polarizer (VCP) the lowest transmission in the off state would be around 25% with an upper limit of around 85% in the on state since the quarter-wave plate does not attenuate the transmission of the variable contrast filter.

The Cell and Associated Electronics

The cell of a variable density filter is a sandwich of three pieces of glass with two 10-micron layers of dyed liquid-crystal material. The inside surfaces are coated with a transparent conductor (ITO) and a thin polyimide alignment layer. The alignment layers are parallel to each other in each compartment and the liquid crystal follows this alignment. The alignment in the two layers is orthogonal.

The dichroic dye in the liquid crystal follows the alignment of the liquid crystal and is much more absorptive in that axis than in the other directions. This configuration will linearly polarize light. When voltage is applied, the liquid crystal (and therefore the dye) realigns to follow the field perpendicular to the substrates. The dye is not very absorptive in this axis so light transmission is higher.

The drive electronics need only provide 0-10 volts at any convenient frequency above 30 Hz. The current is around a milliamp for a typical cell fully driven at 10 volts.

Variable circular polarizers are made by placing a quarter-wave retarder at 45 degrees to the alignment of a single-layered variable density filter.

In the VCP, after light has been plane-polarized the velocity of light is differentially retarded along two axes disposed 90 degrees to each other and each at 45 degrees to the axis of the VCP. If these waves are reflected back through the retarder, one ray lags 180 degrees. The plane of polarization of the slower wave has been rotated 90 degrees and the light will be cancelled when it re-enters the VCP on its way out. The overall result is the same as obtained with a standard circular polarizer: ambient light is efficiently absorbed, but light emmited from the display comes through with little attenuation; emmitted light is only polarized, not absorbed.

With a standard circular polarizer, the transmission range is limited and fixed. With the VCP, however, the transmission can be much higher and by applying voltage the light is circularly polarized to a lesser degree and the transmission is better, enabling higher brightness. The transmission is about the same for the VCF as shown in figure 1. The wave plate does not add density.

Transmission Selection

The transmission of the VDF cell (and therefore the VCP cell) varies with voltage. Any transmission level between the specified extremes is selectable and repeatable, with the appropriate voltage. The uniformity of the transmission across the cell is better than 1%.

The variability of these cells lends itself to automatic adjustment of light transmission to match ambient light levels. *By combining adjustability of transmission with adjustment of the brightness of the light source itself, a display device can be adapted to a wide variety of display applications.*

It is also possible to optimize areas of a display for certain display functions. These sections might employ different densities or different colors.

Environmental Stability

The VDF cells have been evaluated successfully against MIL-STD-810C and have passed this spec's vibration, shock, and altitude tests. The cells also passed 100% humidity at 90°C in the operating mode. The operating temperature range is −50°C to +100°C and the storage range is −65°C to +115°C. These cells are much more inexpensive and rugged than typical LC cells due to the materials used and lack of polarizing films.

For More Information

For more information call Bob Hubbard, 627-6688 (50-320) or Tom Haven, 627-4910 (48-300). □

Tek CAX Files Now on Your Terminal
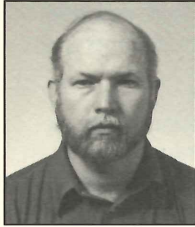
Jim Carden is the manager of Mechanical CAX in the CAX Center, part of Scientific Computer Center. Jim joined Tek in 1979 from Transamerica DeLaval, where he was CAD/CAM System Coordinator. Earlier, he was employed by American Products, a small aerospace company in New Jersey.

Ever wish you could get information about a Tek product on your terminal no matter where the information is stored? Maybe you need to see a parts drawing, look at a schematic diagram, or check on an engineering change made several years ago. The CAX File Management System (CFMS) will do all this. It manages many kinds of files, controlling the development, and the sharing use of data files through the Tek Engineering Network.

The file development controlled by CFMS involves file creation by any system, usually a CAD/CAM system. File control includes tracking by version or alternative, as well as the transferring and archiving of files. File sharing means that someone in Metals, for example, can use files created by an engineer in a division.

CFMS enables users to freely share information stored in data bases on various computer hosts companywide (see figure 1). The physical connection between data bases is across the Tektronix Engineering Network (hyper-channel).
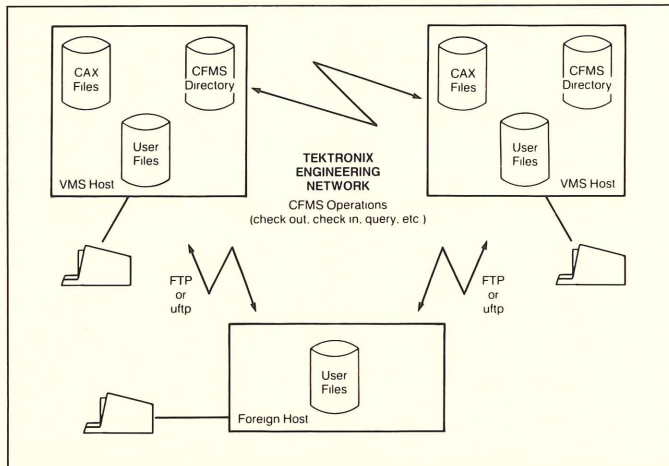

Figure 1. The CAX File Management System enables users with authorization to access and share databases companywide.

CFMS allows any number of files to be grouped and managed as a unit, thereby reducing file duplication, (see figure 2).

CFMS lets users search for a file by part number or project number. If the borrower has authorization, CFMS finds the file and retrieves it.
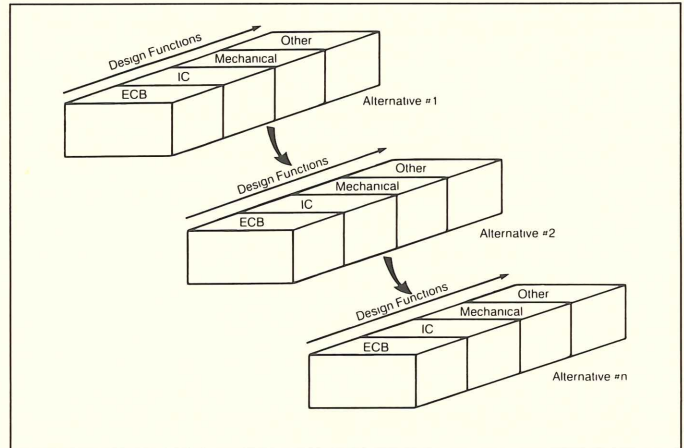

Figure 2. By grouping files as families of designs, files can be managed as units.

When you get the file you want, it will contain all changes. You can rely on this because CFMS stores all information about each revision level together. By default, the system retrieves the most current information, say, on a part—but it will access earlier versions when needed.

Although CFMS manages version control, it does not manage changes within the files themselves. Version control prevents conflicts and inconsistencies when several users use the same file. Only users with authorization can update files or change access permissions.

Figure 3 shows the CFMS architecture. The prompting, terminology, and functions vary according to the specific applications of users.
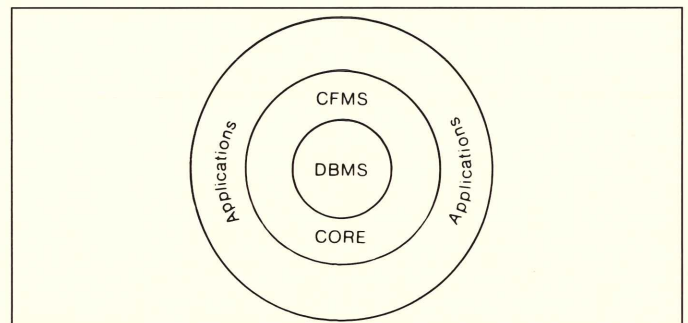

Figure 3. CAX File Management System architecture. (DBMS stands for Database Management System.)

CFMS applications are general. For example, CFMS manages files containing all types of project information needed by mechanical engineers, including bills of material, security, and administrative details. In other cases, it may handle similar details for hardware designers, software engineers, or even chemists.

If you would like to help develop applications tailored to your group's needs, call an application group manager: Electrical CAX—Jack Hurt 627-2599, Mechanical CAX—Jim Carden 627-1533. □

Seventeen Steps to Product Reliability

By Jon S. Potts, Manufacturing Engineering,
Special CRT Manufacturing, DDO (627-5113/48-320).

With the increased emphasis on rapid product introductions and on Just-In-Time manufacturing, we must not forget reliability.

Perhaps my *Seventeen Steps to Instrument Reliability* can play a small role, helping managers, project leaders, designers, and buyers to build reliability into products and then assure that those products stay reliable:

1. Buyers, managers, and designers – Determine the number of defectives acceptable for each component. Whether the rate be one defect per billion or one in ten, the number must be determined—and both you and the vendor must agree on that rate.

2. Buyers and designers – Qualify and accredit vendors following a structured process.

3. Buyers and designers – Procure and *expect* quality parts. Require vendors to document quality.

4. Designers – When appropriate, pre-stress, inspect, and test incoming parts.

5. Managers – Support the gathering and listing of recommended parts of known high quality and make the lists available to your designers.

6. Designers – Design for reliability, derating components consistent with value and need. Use recommended parts and multisource parts – even when buying from a single vendor.

7. Managers – Put the responsibility for reliability with the project manager first; the project manager should then delegate responsibility to the design engineer.

8. Managers – Set aggressive, but *realistic* goals with each designer for each circuit, module, or frame.

9. Managers – Provide information, guidelines, tools, resources and reliability data to designers – Q.A. feedback, for example.

10. Buyers and designers – Don't *assume* you will always get quality parts from the vendor (the odds are that you won't).

11. Buyers and designers – Don't assume vendors will tell you when they change their processes.

12. Buyers and designers – Don't assume that a vendor's manufacturing processes and testing are infallible.

13. Managers and designers – Don't assume you have quality design.

14. Managers – Don't assume your own manufacturing processes are producing quality or will continue to produce quality without a continued and concerted effort on your part.

15. Managers and designers – Routinely sample and stress test to verify the reliability of the design, components, and manufacturing-process.

16. Designers – Determine the true cause of failure (component, design, or environment) before attempting a fix.

17. Managers – In practice treat the importance of reliability as equal to that of schedule or cost. Don't just talk about it, show it with your actions. ☐

Technology Report
MAILING LIST COUPON

☐ ADD
☐ REMOVE

Not available to field offices or outside the U.S.
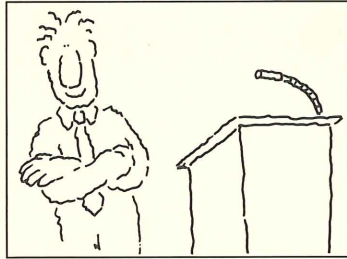
MAIL COUPON TO 53-077

Name:_____ D.S.:_____

Payroll Code: _____
(Required for the mailing list)

For change of delivery station, use a directory change form.

Been Asked to Organize a Session or to Talk at a Conference?

Professional conferences are an integral part of the infor-mation-transfer process. The success of such conferences depends on many things. The theme, the facilities, even the weather are important, but the speakers and the organizers are critical.

To help session organizers and speakers, Technology Com-munications Support has prepared two guides, each based on extensive experience in preparing and supporting professional communications.

If you've been invited to talk, *So You're Giving a Talk at a Pro-fessional Conference* will help you organize your talk and pre-pare effective slides to go with it. The rules in this booklet are few and simple, but they can help you look like a "pro" in-stead of an ill-prepared amateur.

If you've been invited to organize a session, you'll be expected to be not only the manager of a team but the coach, scheduler, and expediter as well. So *You're Organizing a Session for a Professional Conference* gives the basic rules and guidelines.

To get a copy of either or both guides, and to get professional support in preparing talks and slides, contact Technology Com-munications Support d.s. 53-077 (642-8920). □

COMPANY CONFIDENTIAL
NOT AVAILABLE TO FIELD OFFICES
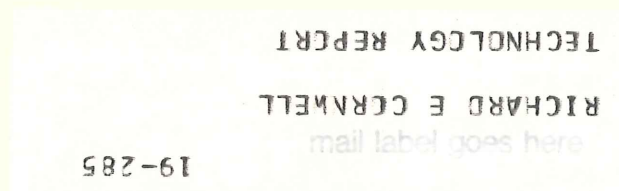
DO NOT FORWARD

Tektronix, Inc. is an equal opportunity employer.