TEK INFORMATION FOR TECHNOLOGISTS

MAR/APR 1985

TECHNOLOGY report

AN EFFECTIVE TECHNIQUE FOR PROCESS CONTROL

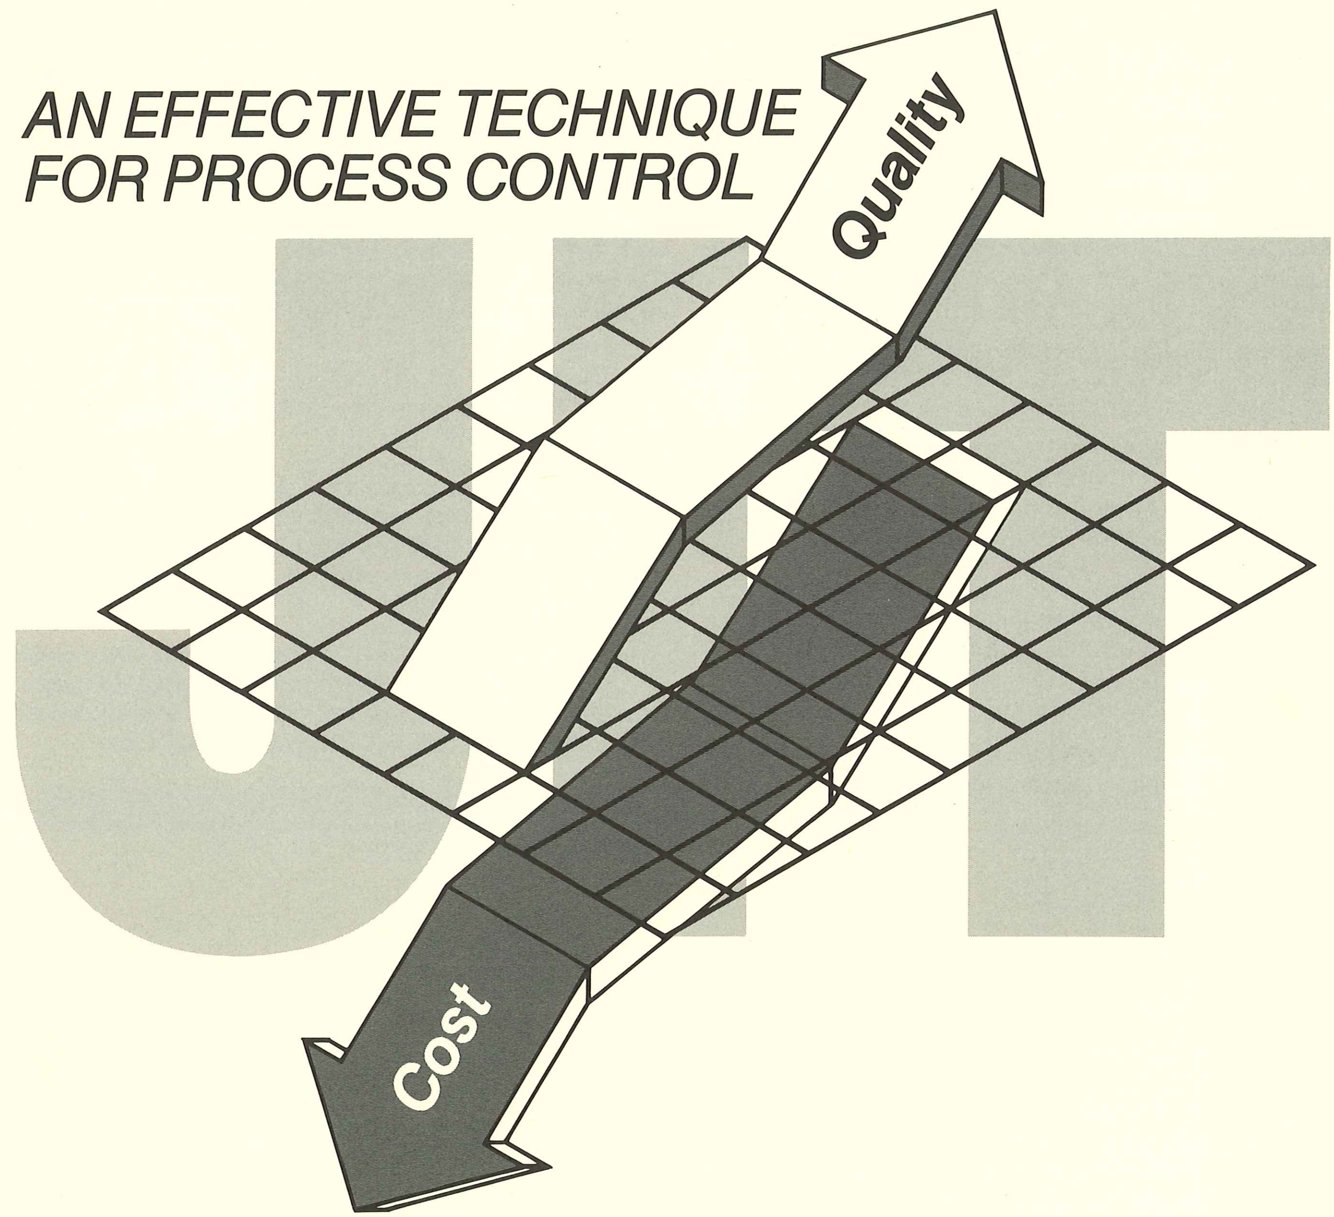

Tektronix®
COMMITTED TO EXCELLENCE

CONTENTS

Volume 7, No. 2, March/April 1985. Managing editor: Art Andersen, ext. MR-8934, d.s. 53-077. Cover: Darla Olmscheid; Graphic illustrator: Nancy Pearen. Composition editor: Sharlet Foster. Published for the benefit of the Tektronix engineering and scientific community.

Why TR?

Technology Report 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).

HELP AVAILABLE FOR PAPERS, ARTICLES, AND PRESENTATIONS

If you're preparing a paper for publication or presentation outside 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 technological 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 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 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. □

JUST-IN-TIME: AN ESSENTIAL TECHNIQUE FOR PROCESS CONTROL

Gayle Albin-Brooks is now a marketing manager in Hybrid Manufacturing, part of HCO. At the time this article was written, she was a production and test manager in the same organization. Gayle joined Tek in 1980. She has a BA in speech communications and English from Portland State University.

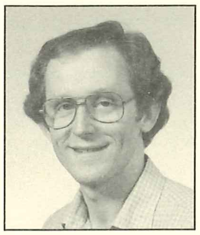

Michael K. Lewis is a process engineer in Hybrid Manufacturing, part of HCO. Michael joined Tek in 1980. He received his BS in ceramic engineering in 1980 from Rutgers University, New Jersey.

When Just-In-Time (JIT) techniques were integrated with process controls in HCO, yields of a complex thick-film multilayer hybrid circuit rose dramatically. Although budget and physical constraints blocked implementing a pure JIT line, by applying JIT principles a batch-oriented traditional line was transformed into an "almost-pure" JIT line. The authors believe that JIT techniques can be successfully modified and implemented in any manufacturing process at Tek, with similiarly dramatic results.

Early in 1983 our manufacturing team initiated a pilot project that was to apply JIT manufacturing concepts to process control. The project tackled a trouble plagued thick-film multilayer hybrid. We had been manufacturing this circuit on a traditional assembly line.

The hybrid consists of thirteen discrete devices and a thin-film resistor/capacitor (RC) network. These are epoxied down onto a thick-film multilayer substrate. The substrate is then wirebonded inside of a 24-pin ceramic header. Our yields were poor – only three out of ten units would pass final test.

With hard work and some tough compromises, JIT lived up to its promise. After six months, yields reached an unprecedented 98%. Today we consistently hit 99% and it is not uncommon to have some perfect lots. To achieve this our team did not practice pure JIT – if there is such a thing outside of the textbook. We worked with existing facilities, budgets, and an ongoing process. We adapted the textbook practices to our unique environment, rather than initiating an all new production line.

Although an 'ideal' JIT system would originate in and include the materials planning organization, this project was initiated and implemented in production only. This demonstrates that JIT principles can be applied in selective areas of an organization with good results. By tailoring JIT principles to their manufacturing operation, we believe that any production team, regardless of product complexity, can optimize their process-control systems and dramatically increase yield.

In this article we will discuss the differences between traditional and JIT manufacturing techniques, how we overcame the obstacles to implementing a JIT system, and the results of JIT in our application. We will also show how JIT benefits normal process control mechanisms and speeds up problem resolution.

Manufacturing: Traditional and JIT

In traditional semiconductor manufacturing, production systems are structured by function or process. This is done to minimize set-up times and maximize production efficiencies. In this function-oriented scheme, product batches are large – thousands of identical items are produced. Traditionalists justify this approach because of the long times needed to set up major process steps. They believe that producing large batches increases the efficient use of expensive production equipment. But this 'efficiency' actually burdens the system and obscures problems, because many parts must be handled simultaneously. Inventory climbs, lead times stretch out, and yields decline as a direct result.

In traditional systems, quality feedback is slow; slow feedback makes it difficult for operators to assume responsibility for quality. Individual processes are frequently optimized for volume at the expense of overall yield; since the emphasis is on quantity, process controls for quality are ineffective. Everyone's great fears are stopping the assembly line and missing production quotas. Therefore problems frequently remain, going unidentified and unresolved. While intended to use labor and equipment efficiently, the traditional approach, taken to its extreme, actually does the opposite and costs rise.

In contrast to traditional methods, 'pure' JIT is marked by dedicated product lines, small lots, and a willingness to practice line stop. In actual practice, more than one product will be run on a JIT line, but those products will be a family having a common technology and similar features. The equipment is dedicated to the family, rather than to just one product. Other marks of JIT are low inventory, short lead times, speedy feedback, and more quality-effective process controls. Figures 1 and 2 show the differences between the two basic flows.

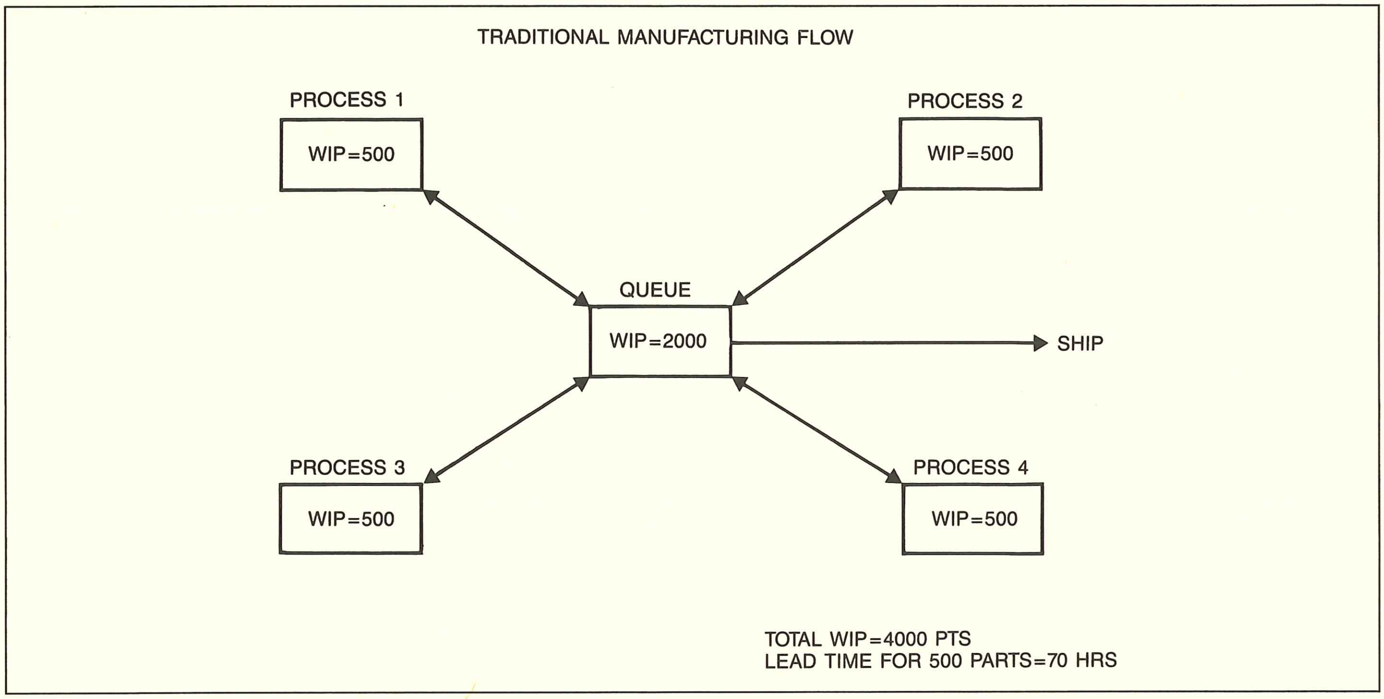

Figure 1. In traditional manufacturing, work in process (WIP) moves in and out of a queue. While supposedly using labor and equipment efficiently, this method is often quite inefficient as part quality is hard to track.

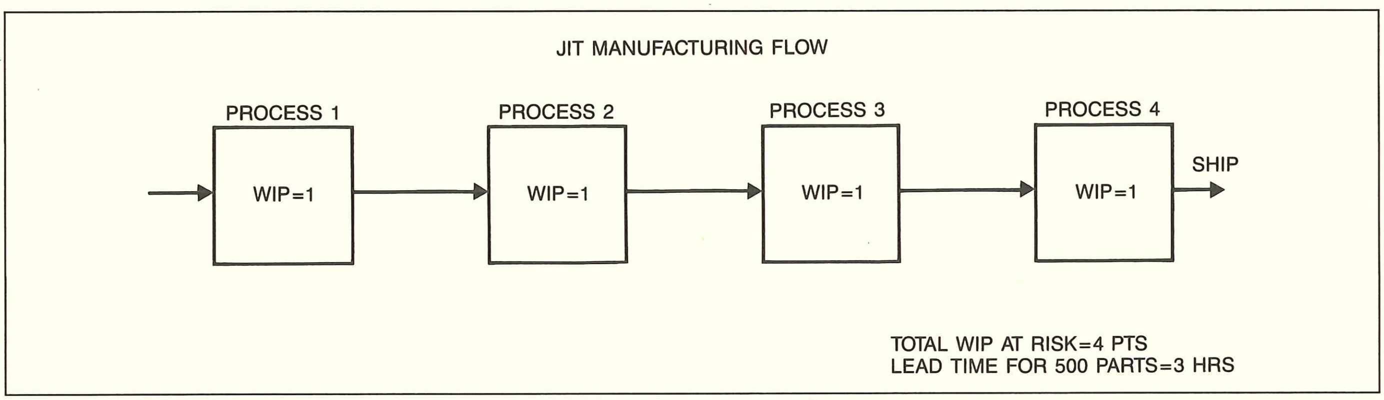

Figure 2. Pure JIT flow is linear. There is no queue and WIP is limited to one. Most bad parts are discovered promptly by the operator who made it. Processes can be corrected with equal promptness before yield deteriorates.

The high volume of work-in-process (WIP) in the traditional model obscures problems. It's almost impossible to track and coordinate the thousands of parts in a system such as the one shown in figure 1. While operators can easily handle few parts at a time and assume responsible for quality, they cannot and do not feel the same commitment to large batches. Just doing the paperwork needed to track 4000 parts becomes, by necessity, more pressing than quality.

In a pure JIT line, operators would think twice before rejecting a single part. With only a few parts in the loop, it's easy to pinpoint the cause in the sequence of events leading up to a problem.

In the traditional model long lead times prevent timely feedback and the prompt correction of problems. The traditional manufacturing model in figure 1 produces 4000 parts before operators receive feedback from test on the quality of the first 500 parts. All 4000 parts could be bad – and need screening, rework, or rejection.

In the traditional model, the pressure to ship, ship, ship is intense. The idea of shutting down the line to correct a problem meets great resistance. Instead, the emphasis is on screening for quality rather than perfecting process controls to build the parts correctly on the first time through the line. Often problems go unresolved and become chronic. In the traditional system, cures for chronic ailments seem to disrupt the system rather than help it.

In the JIT model, however, feedback is quick, a matter of hours. When a quality problem is discovered, the line is shut down until the problem is resolved. Since the amount of WIP is small – in the example in figure 2, only 4 parts – quality

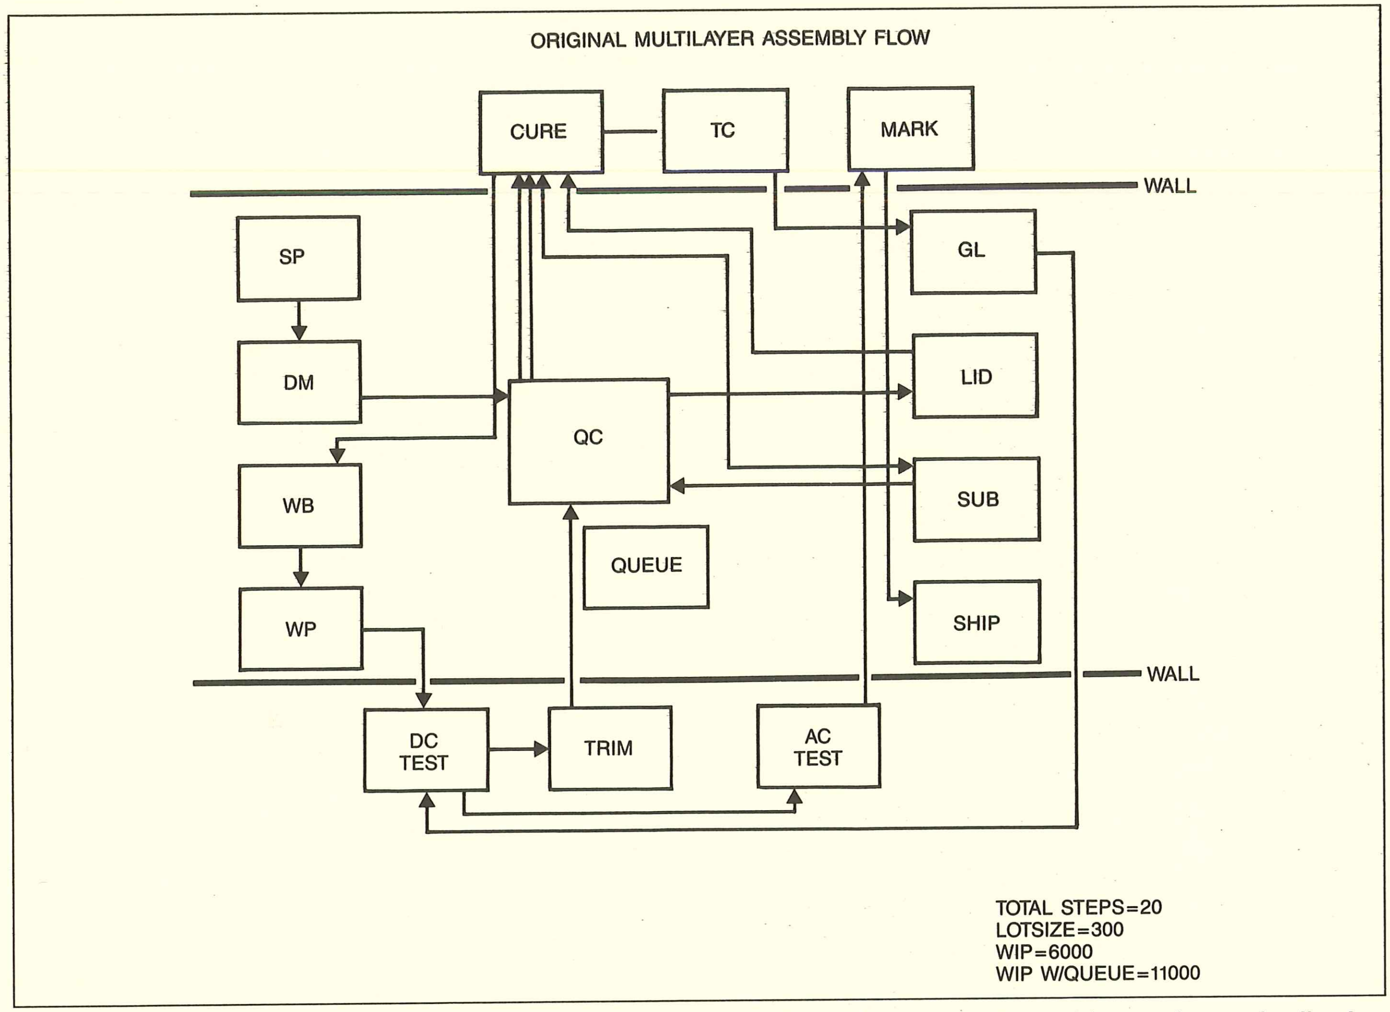

Figure 3. The hybrid assembly flow before JIT principles were applied. This complexity and its ponderous feedback isolates operators from the quality imperitive. The WIP and queue level were enormous – 11,000 parts whose quality was ambiguous.

problems are easy to identify and correct as they occur. The emphasis *is on fixing the process instead of screening the product.* This increases the yield since the process control system is increasingly fine-tuned so that only good parts will be produced. It is this interaction between JIT and process control that makes the system so dynamic, so responsive.

Our Pilot JIT Project

To launch our pilot project, we first needed to remove the obstacles to to a pure JIT manufacturing flow, as many as possible. We had to do this if we were to achieve our key objectives; that is, increasing yield, reducing WIP, and shrinking lead times. In analyzing our existing process flow for the multilayer thick-film hybrid, we discovered nine major obstacles (see figure 3 for the existing process flow):

1. Undedicated equipment.
2. A physical as well as psychological wall separated test from assembly operations.
3. Operators were neither versatile nor accountable.
4. Quality feedback was slow.
5. Too much paperwork; too many parts on the floor.
6. The MRP system would not accommodate JIT scheduling.
7. Two-hour cure cycles.
8. Too many QC steps.
9. A bottleneck at the wirepull operation.

To set up a JIT process, our first step was identifying the equipment needed for a dedicated product line. But to set up this line, we had to remove the physical wall that separated test from assembly. The wall was not only physical; it was psychological too. To the operators in assembly and to the people in test, what happened on the 'other side' had little relevance.

Since we couldn't actually remove the offending wall, a compromise was worked out: For certain periods, the test group dedicated their equipment exclusively to our multilayer. Next, to eliminate the psychological barrier, we developed a small team of people from both the assembly and test groups to work together on the JIT project team.

At this time, operators were skilled in no more than three of the five major processes used to make the hybrid. To learn all five, they began cross training in assembly and test. They were also trained in the basics of JIT manufacturing.

To reinforce the team concept, yields were tracked through both the assembly and the test areas. Operators at all process steps were given accountability for yields at their steps, as well as for the overall-yield target. Other team members, including managers and engineers, were also made accountable for the target yield.

Slow feedback had made pinpointing and solving process problems difficult. With the fast feedback of the JIT system, we quickly spot a troubled process.

Another crucial need was speeding up the feedback on those hybrids that failed final testing. Before JIT, failures were analyzed off-line by technicians and engineers; this delayed feedback. Instead of waiting for analysis, operators were trained to recognize the 'footprints' characteristic of the most common failures. Operators were also trained to correct problems on their own. With the feedback and the training, most problems now are resolved as they occur, or on the same day. Operators, by being able to analyze and correct failures on their own, gained insight into their processes – and pride in their skills.

Excess paperwork had been absorbing a lot of operator time, often focusing them on accurate bookkeeping rather than making good parts. The paperwork was simplified. Twenty reporting steps were reduced to three. Specifications and process flows were combined onto a single routing that travelled with each lot of parts. The results:

1. Operators did far less paperwork, concentrating on making quality parts.
2. Teamwork was reinforced – no one fought over counts and rejects.
3. We reinforced the JIT concept of waste reduction by eliminating paperwork that was both unproductive and costly.

The next obstacle to be removed was too much work in process. High WIP was dictated by the Tektronix Materials Resource Planning (MRP) system. Since we were not free to change the MRP system, we worked out a compromise with the scheduling group: The MRP system would continue to set weekly quantities. Manufacturing would then manually split these weekly quantities into daily buckets. This allowed us to process the smaller quantities necessary for JIT manufacturing.

A stubborn obstacle to JIT flow remained – the epoxy cure. The cure took two hours and prevented us from achieving pure one-at-a-time processing. It was impossible to eliminate the cure step entirely, but could a one-hour cure be substituted? Yes, engineering tests confirmed the reliability of parts subjected to the shorter cure.

The shortened cure cycle, itself, now became the critical obstacle to pure JIT processing. We established one-hour cycles at each of the process sequences instead of pure one at a time processing. The number of parts that could be processed in one hour at the slowest sequence – ten – became our JIT lot size. Although not 'pure' JIT, ten-part batching was a quantum jump upward from traditional manufacturing where we had 6000 or more parts as WIP.

After reducing lot size, we turned to the problem of visual inspection (VI). Too much emphasis was being placed on inspecting in quality rather than doing the job right the first time. Inspection was a large-batch procedure in between each process step. VI stations were 'holding centers' for a lot of potentially bad parts. The centers delayed feedback to operators, preventing them from assuming full responsibility for their work.

We eliminated the VI holding centers and trained operators to do all visual inspection themselves. This gave them responsibility for passing on only-good parts to the next step. Emphasis was put on the operators producing 100% good parts rather than just meeting efficiency standards. Operators were responsible for their quality and the narrow job of QC inspector was eliminated. Any part that did fail the specified quality criteria was reworked by the original operator. This increased their motivation to do the job right the first time.

Nondestruct wirepull is a wasteful process, a classic example of inspecting quality into a product. We were doing it to every part because the wirebonds on the multilayer substrates kept failing. Ironically, it turned out that wirepull itself was causing rejects. We found that just touching a part can damage it. Since one-hundred-percent nondestruct wirepull requires handling every wire on every part, the likelihood of damage is astronomical. It was the handling and the handling tool, not the stress of the pull, that caused failures. The wirepull tool would often scratch the die or cause two wires to short.

By receiving timely feedback from our JIT line, the manufacturer of the thin film, another part of HCO, was able to improve its process and supply 100% good parts.

We started a major effort to improve yields by tightening up process control at the automatic wirebonder. Detailed records were kept, at the wirebonder, of the location and bond parameters of each failure. Every four hours one part was subjected to one-hundred-percent destruct wirepull. This data revealed two problems in the thick-film multilayer substrate. First, since wirebond pads were printed in different gold layers some pads were exposed to the thick-film firing cycle more times than others, creating a difference in bondability. Second, sections of the substrate were being contaminated with slag from the laser scribe.

To solve the first problem, the multilayer was redesigned, placing all bond pads on its top layer. The second problem was fixed by coating the substrates with a protective layer before

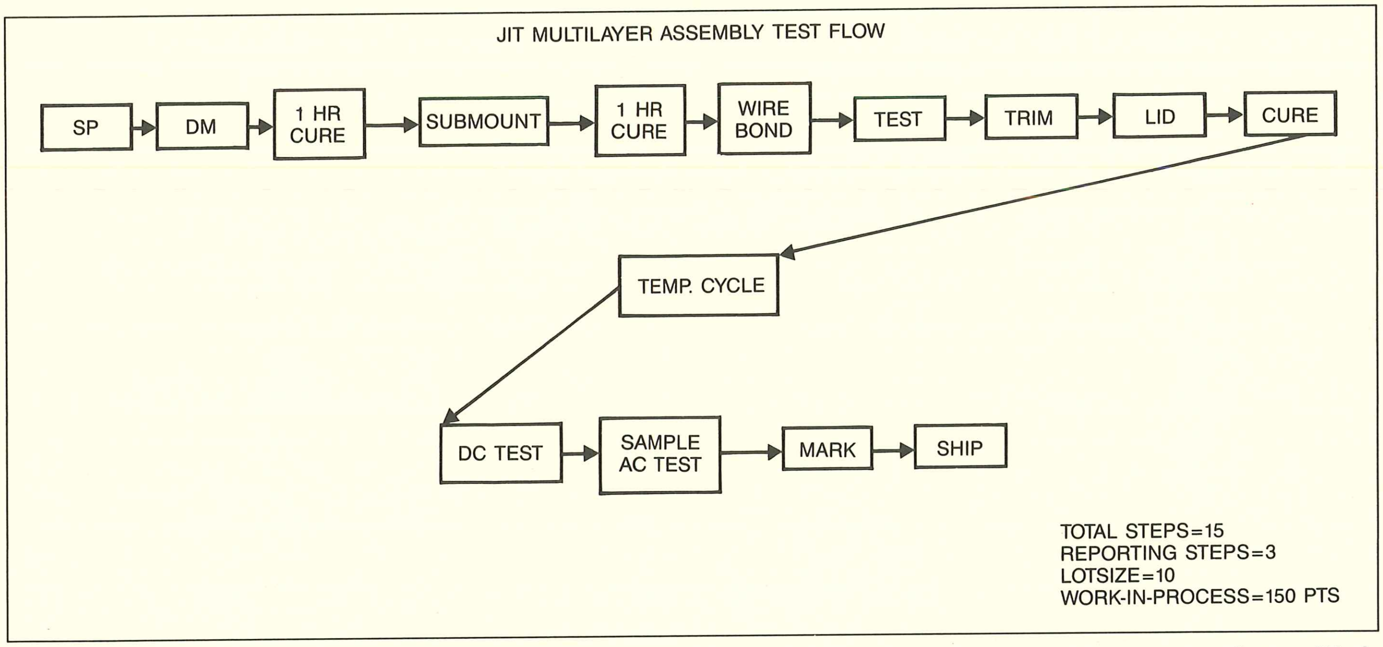

Figure 4. The JIT assembly flow established for the hybrid is linear and simple. Yields are 98 percent or better. Work in process is down to 150, not the ideal one per process step – some existing constraints could not be removed.

laser scribing. This coating was washed off afterwards with water. With these two corrections, bond yields stabilized at 99.8%, that is less than two failures per thousand bonds.

It was now possible to set up process control charts for our automatic wirebonder. We used the data from our destruct tests to build mean and standard-deviation charts. These charts are tools used in equipment qualification.

As a result of these investigations and changes, we eliminated nondestruct wirepull, increasing overall yield by about 10%.

JIT flow made wirebond control easier. On our old assembly line, attempts to improve wirebonding had failed, swamped by the effort needed to track all of the bond failures. Before JIT, getting out the day's production quota was the imperative – not collecting the data needed to improve yield.

Once we had removed the nine major obstacles to JIT, our process flow took the form shown in figure 4. Our team had achieved an 85% yield simply by tying our processes together in a modified JIT system. This showed us that there was an significant relationship between JIT processing and effective process controls.

One-at-a-time processing is the ultimate process control because each time a part is made no more parts are made until the first part proves to be perfect. Even with our modified JIT system we do almost as well.

With our JIT system started, we quickly identified these remaining obstacles to 100% yield:

1. Thick-film shorts between the layers in the multilayer substrate.
2. Thin-film shorts in the RC network.
3. Tin-contaminated inductors.
4. Incorrectly mounted substrates in the ceramic header.

These problems had been obscured by the thousands of parts in process in our old manufacturing system. Slow feedback had made pinpointing and solving process problems difficult. With the fast feedback of the JIT system, we quickly spot a troubled process.

Half of all failures were thick-film shorts in the multilayer substrate. These substrates were not being electrically checked before assembly. Our thick-film fabrication area was not getting feedback on their quality. It took only a week to set up a test using existing laser trim equipment to identify bad substrates.

The substrate had 52 independent conductor runs in three layers. Since the laser equipment could only scan 17 at a time and it would be too expensive to run the test more than once, we had a problem. And, we couldn't just go out and spend $50,000 on another test system. But, we could combine three or four runs – that could not short to each other – into a group. In this way all 52 runs were formed into 17 simultaneously testable groups.

Feedback from this test was sent immediately to thick-film fabrication. With it, they could evaluate process changes intended to prevent the printing of shorted substrates. Initially, 25% of the substrates produced were screened out. After two batches, rejects were only 4%. This shorting problem was resolved quickly because our test results sped quality feedback to the fab area.

In looking into the electrical shorts in the thin-film RC network, we found two contributing factors: One in the manufacturing process for the thin-film and the other in our handling of the networks during assembly. By receiving timely feedback from our JIT line, the manufacturer of the thin film, another part of HCO. was able to improve its process and supply 100% good parts.

The RC network has six laser-trimmable interdigitated capacitors. These are made up of one-mil-wide fingers and spaces, mostly covered with a protective polyimide layer. Our handling techniques were damaging the part of the fingers that were in an unprotected area, left open for laser trimming. To solve the problem, operators were asked to use Teflon rather than metal tools to handle the networks during sub mount. The timely feedback from the JIT flow increased the operators' willingness to use the Teflon tools. This resulted in a yield increase. Teflon tools, however, could not be used to repair wirebonds. To reduce failures caused by metal tweezers, the polyimide protection was extended. This reduced the chances of damage from the tweezers.

Open inductors were a most elusive failure. Our hybrid has six 510-nH inductors epoxy mounted to the thick-film multilayer. Many eventually became electrically open. We could find no pattern to the failures. So, consistent with JIT philosophy, we shut down the entire assembly line until the cause could be identified. Within a week the problem was traced to the coating on the inductor terminations. The vendor was tinning the inductors for solder mount. The tin oxidized through the epoxy we used to mount the inductors and created an open circuit. Even when the inductors were ordered untinned, the problem did not go away. The vendor's equipment was contaminated with tin from earlier production lots and our inductors ended up with tin residue. After helping the vendor to improve its process, we began to receive uncontaminated parts.

As their JIT skills developed, the operators became intensely committed to identifying problems. For example, they identified the need for a pin-1 location marker inside rather than on the bottom of the ceramic header where it was invisible to the operator during sub mount. An inside – visible – flag was added to the header. With this simple change, no more packages were assembled incorrectly.

The last assembly step is sealing the packaged hybrid with a B-stage epoxy lid. The sealed package was being subjected to a gross-leak test before shipping. We questioned if the test itself was necessary. A reliability test was started on a sample that parts failed gross leak. These 'failed' parts proved to be reliable. We removed the gross-leak test from the flow and substituted a destruct lot qualification of the lids to guarantee good product. Eliminating gross leak further shortened our feedback loop, helping us further optimize our process controls.

Conclusions

Six months after implementing our pilot JIT manufacturing system, we had achieved 98% yields. Shortly thereafter, JIT was successfully introduced to a high-volume assembly line and to a low-volume custom assembly area. In both cases, the pure JIT model was modified to match the unique environment of each product line. The results were comparable to those realized in our pilot project. Yields increased from 75% to 96% or better, feedback cycles shortened, and leadtimes reduced 75%. As our production teams continue to gain skill in using JIT in process control, they are further refining their processes, driving for 100% quality and the lowest product cost.

Once the basics of JIT manufacturing are grasped, it becomes clear that the principles of one-at-a-time processing are, in themselves, critical components of a process control system. Shortened feedback cycles, low inventory, and line stop all maximize process control. Indeed, we have concluded that JIT is the ultimate process control. It reveals how a process is performing and enables everyone – engineer, manager, or operator – to exert maximum control over their work.

In this article, we discussed methods to remove obstacles to JIT and how we used JIT to control our processes and improve our yields. We also briefly discussed the successful application of the JIT technique in other product lines. From our experiences with the dissimilar nature of these product lines, we conclude that JIT techniques can be successfully modified and implemented in any manufacturing process at Tek – with the same dramatic results.

For More Information

For more information, call Gayle Albin-Brooks 627-6803 (13-850) or Michael K. Lewis 627-3967 (13-035). ☐

GOT A PART PROBLEM? PERHAPS A METALLURGICAL EVALUATION IS THE RIGHT STEP

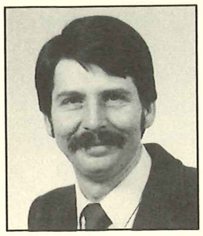

John DeHaven is a metallurgical engineer in the ACC Metallurgy Lab, part of Applied Chemical Components (ACC). He has had 10 years experience at Omark, Hyster, and other companies. He taught several metallurgy classes at the college level before joining Tek in 1984. John's BS in metallurgical engineering is from Oregon State University. He has done graduate work in metallurgy and failure machanics.

Knowing just how a vendor's manufacturing process is causing trouble enables you to rapidly negotiate solutions – because your position is very strong. This article describes one case involving thousands of expensive parts and how ACC's Metallurgical Lab was instrumental in helping a division resolve a critical problem.

Initial start up for the new 4692 Color Printer was being delayed by ink-jet problems. Plugged ink-jet orifices were causing erratic spray patterns and intermittent operation, sometimes completely blocking ink flow. The contaminating particles continued to be generated, despite repeated flushing with filtered de-ionized water.

The ACC Metallurgy Lab was asked to determine the cause of contamination and just how the vendor was manufacturing the heads. A look with a scanning electron microscope showed that the inside bore had been finished by honing or grinding. The tool pressure used during this operation had been great enough to separate grain boundaries, releasing tiny particles (20–100 μ) to drift through the head assembly. These particles, when worked by continued tool movement, became long, thin flakes of metal. It was these flakes that caused the problems.

The Lab then cross sectioned, mounted, and prepared the ink-jet head for optical metallographic examination. Metallographic preparation entails imbedding the sample to be examined in one of several types of plastic, and grinding and polishing to obtain a mirror smooth flat surface. This surface, when exposed to certain etching acids, reveals structures within the metal. These structures help to identify how the part was made (casting, extrusion, rolling) and its thermal history (heat treating, annealing, welding). This examination allowed us to determine that the head was cast, most likely by investment casting.

While investment casting is normally the most cost-effective method for precisely casting many parts, in this instance the casting quality was somehow poor. We saw many open voids in the critical area just under the bore surface. This type of void is known as shrinkage porosity; it's usually caused by allowing the metal to solidify at the entry point of the mold cavity before the metal within the mold itself has solidified. These voids trapped the metal flakes but allowed only a few to wash out during each flushing cycle, leaving many behind to be released during later cycles.

With the evidence we had obtained from these and other tests, Maridana Whitlow of GPP's Process Engineering confronted the vendor in Japan. The vendor, recognizing that a serious performance problem existed and that their manufacturing technique was clearly at fault, accepted the return of 4800 heads, which at $157 each were worth about $750,000 according to Doug Stanley, manager, GPP Process Engineering.

The evaluation of the ink-jet head is an example of how our metallurgical lab can help resolve your metal-working problems. We offer a complete range of services including:

- Mechanical testing including tensile/compression/peel strength of a variety of materials.
- Optical and electron microscopy including coating thickness measurement, material quality verification, and failure analysis.
- Batelle-chamber accelerated corrosion tests.
- ECB plating thickness, multilayer registration, solderability and terminal pull testing.
- Design consultation including materials and process selection as well as quality testing recommendations.
- Vendor qualification.

For More Information

For more information or for a consultation on your particular application, please call John DeHaven 627-0259 (38-314). □

THE MODULAR INTERFACE SYSTEM, HANDLES UP TO 249 NODES THROUGH A TWO-WIRE BUS

David Williams is an engineering project manager in the Modular Interface Systems group, part of the Engineering Computing Systems Division. David joined Tek in 1977. Earlier, he had worked at the Data Systems Division of Litton Industries and the Commercial Electronics Systems Division of RCA. He received his BSEE from the University of Portland.

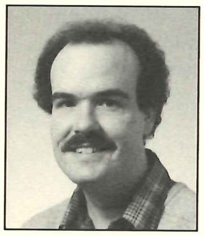

John Balding is a software/hardware engineer in the Modular Interface Systems group, part of the Engineering Computing Systems Division. John joined Tek in 1981. He received his BSEE from the University of California at Davis.

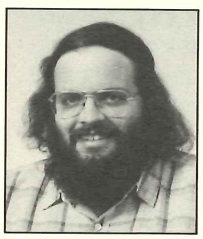

Jim Sullivan is a software/hardware engineer in the Modular Interface Systems group, part of the Engineering Computing Systems Division. Jim joined Tek in 1979. Earlier, he worked at the Boeing Aerospace Division. Jim received his BSEE and BSMTH from the University of Portland.

This article reports on the Modular Interface System (MIS). This system, developed by ECS, allows up to 249 input devices to be operated from a single port on a workstation, terminal, or other similar host.

Terminals, workstations, and programmable instruments usually come with a keyboard for communicating with the host. When other input devices are added, such as joysticks or trackballs, the usual practice is to add a port to the system for each device. These ports are designed for the specific input device. If input devices are added after the host product is built, the host must be modified. Because such modifications can be expensive and take the host out of service, few customers are willing to reconfigure a system after it has been purchased.

In contrast, the Modular Interface System (MIS), developed by ECS, allows a customer to "safely" purchase a system configured to meet immediate requirements. When those requirements change, that customer can easily change the input-device complement to fit the new requirements – without modifying the original.

Another advantage of the MIS approach is that input devices can be interchanged among other Tek products that also employ the MIS system. For example, a Japanese keyboard for a product with an MIS can be used with any other MIS-employing system used in Japan. This, of course, is far better than having unique keyboards for each product going into Japan or other countries or regions.

The interface system allows up to 249 MIS nodes to be connected to a single two-wire bus. This bus is routed to the various device nodes through a shielded cable (see figure 1). This cable also carries 12 Vdc and return. An Intel RUPI (Remote Universal Peripheral Interface) microcontroller is located at each of these device nodes and at the host. The RUPI at the host acts as the controller, regulating all the traffic on the bus. We usually refer to the bus as MUIB (Modular User Interface Bus).

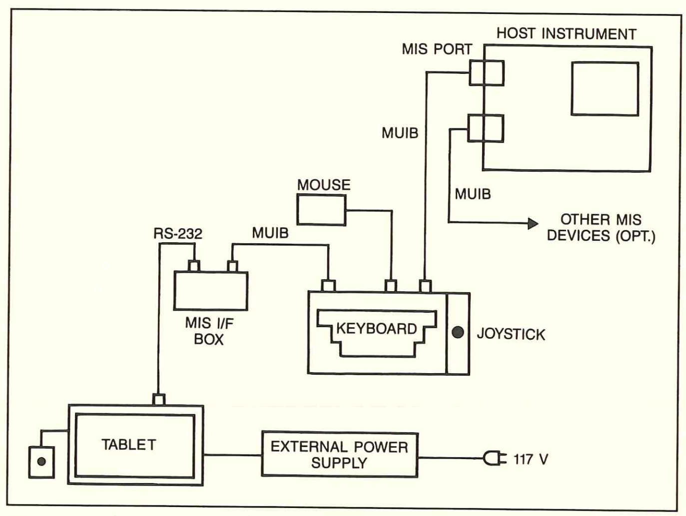

Figure 1. One type of port in the host that's compatible with all devices, this is the MIS concept. Rather than a mess of parallel cables to a bunch of varied host-ports, up to 249 input devices can be operated on a single two-wire bus. A microcontroller in the host regulates the bus traffic. Each device (slave node) has its own microprocessor. RS232 devices may be interfaced to the MUIB through an MIS Interface Box. Each graphic-input device is housed in a physically attachable enclosure, enabling joysticks, auxilliary keypads, and so forth to form a compact, almost one-piece unit with the keyboard.

The MIS configuration can also support output devices such as low-speed printers and signaling devices. In ECS products, the host provides 12 Vdc at 2 amps to power the various devices on the MUIB. This voltage, of course, can be regulated within the device to the level needed. If 2 amps is not adequate, the device will have to be powered separately. The MUIB controller located in the host will support up to 16 devices of any one type on the bus at the same time. And the controller can communicate to each of these devices individually.

The Intel RUPI microcontroller (see figure 2) is the heart of the MIS. This chip features an 8-bit 8051 microcontroller CPU connected to an intelligent high-performance SDLC serial-communications controller through a dual-port RAM. The communications controller, called the serial interface unit (SIU), manages the interface to a high-speed serial link (312.5 kbits/sec). The SIU off loads communication tasks from the 8051 CPU, thereby freeing the CPU to concentrate on real-time control tasks.

The RUPI chip has 4K×8 ROM and 192×8 RAM, 32 programmable I/O lines, and two 16-bit timer/counters. There are three RUPI types: The 8044 has masked ROM; The 8344 has no onboard ROM; The 8744 has a user-programmable/erasable EPROM. ECS will use all three types. The SIU of the RUPI interfaces to the MUIB through a TI 75176A differential bus-transceiver (RS-422) chip (see figure 4). This chip provides a differential pair operating at +5 Vdc and ground. A 10-MHz crystal is used with the RUPI, giving the bus a data-transfer rate of 312.5 kbits/sec.

The MIS System has two states of operation: *initialization/configuration* and *data manipulation/control.*

In the initialization/configuration state, the MIS power-up diagnostics are executed; next, each power-up device informs the host of its presence on the MUIB. This announcement is a 5-byte code we call the "long id." The first two bytes indicate the device type; each MIS device has a unique type (such as keyboard, keypad, function pad). The next byte is the device's serial number. The remaining two bytes are reserved for expansion.

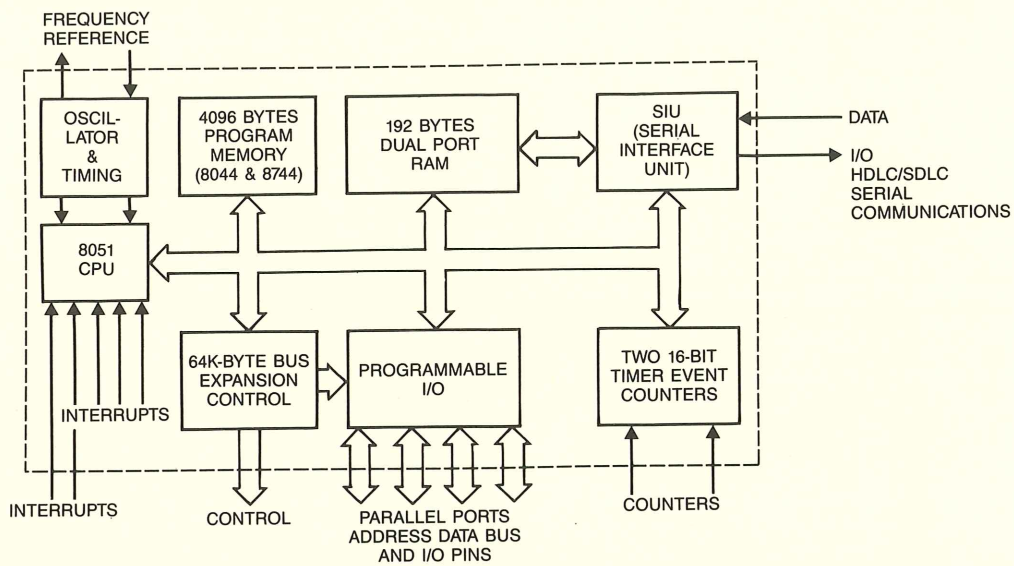

Figure 2. The heart of MIS is one of Intel's three 8X44 Series RUPI chips and its onboard 8051 CPU. This 8-bit microcontroller is connected within the chip to a a dual-port RAM (see figure 4). The RAM, in turn, is connected to the SIU. By managing the interface to the serial-data link, the SIU frees the on-board CPU to concentrate on real-time control. The three RUPI chips differ in the presence or absence of program memory. See figure 3 for the devices that with the RUPI chip form the "ECS" RUPI Controller.

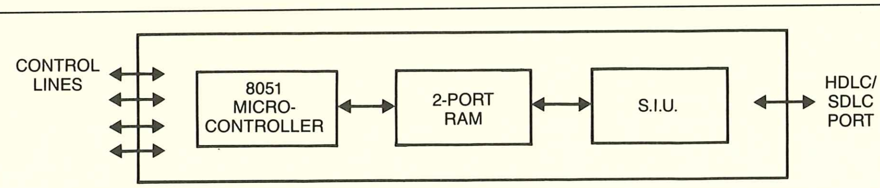

Figure 3. The architecture of the dual controller on the 8X44 Series RUPI chip.

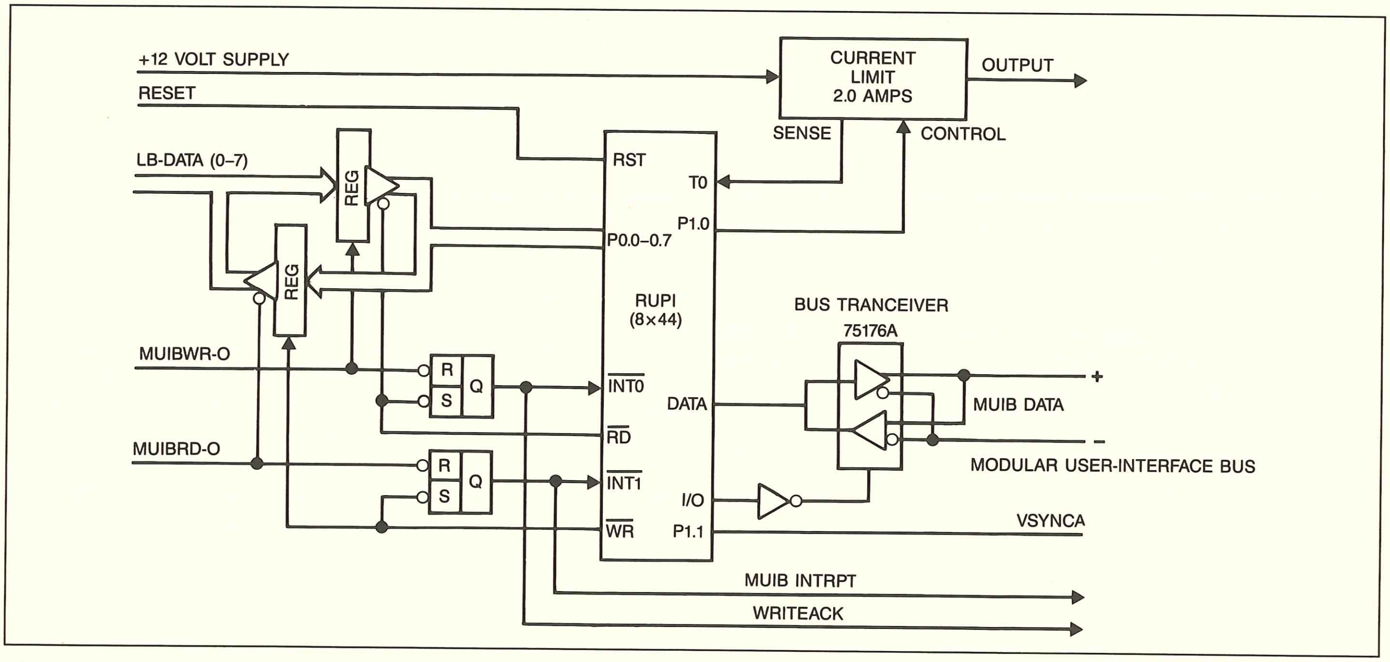

Figure 4. The "ECS" Controller. Its functions center on the Intel 8X44 RUPI chip. The serial interface unit on the RUPI chip (see figure 2) interfaces to the MUIB through a TI 75176/75176A differential bus-transceiver.

The serial number is used to distinguish between devices of the same type. In the ECS implementation, up to four jumpers in the device may be used to change this serial number. The placement of the jumpers is not critical, except that no two devices of the same type should have identical serial numbers. The serial number could also be implemented in a small serial EEPROM. The keyboard used by ECS – built to our specs by KeyTronic – can provide eight serial numbers because it only uses three jumpers.

To avoid bus collisions, the MIS controller in the host employs polling. As the controller receives the long id from each device, it issues a short id of one byte to each active device. The number used as a short id is not significant in itself; the device uses the short id as its address for the current session on the bus. Up to 249 short ids are available. All communications with devices after initialization are prefixed with this short id.

The data manipulation/control state is entered after configuration is complete. When the system is in this state, it can control device functions or transmit and receive data to or from a device on the MUIB.

Data for ECS's "standard" devices come in two formats (types). The first type of data is *device-specific.* This can be key-up or key-down data, or it can be encoded (ASCII) data. The encoding would come from a table stored in the device's RUPI ROM. The second type is *graphic-input* (GIN) data; this typically consists of quadrature counts. The mouse in our MIS supports relative delta counts. The counters for the ECS mouse use a 16-bit modulus. One counter is used for each axis.

The MIS controller uses a poll table that enables it to control up to six active devices at a time. A RUPI slave node can have more than one pollable device. Our MIS keyboard, for example, has two pollable devices at the same node: the keyboard itself and the mouse. Each takes one slot in the poll table.

There are two ways of polling devices. In the *periodic poll,* devices are polled at a rate determined by the number of devices in the poll table and the bus baud rate. The controller loops through the table, polling each entry.

The other entry type is the *sync poll.* This type operates similarly to the periodic poll, except before the controller traverses its poll table, the controller waits for a transition on an external input pin of the RUPI. This allows slave-node polling to be externally synced. ECS uses sync-mode polling to avoid cursor jitter; the system is synced to the vertical retrace of the display.

To enable RUPI slave nodes to have many host-addressable functions, the MIS employs the auxilary-address concept. Each slave RUPI is assigned an 8-bit SDLC address. If there is more than one device at a node, each device is assigned an auxilary address, for example:

Device	SDLC Address	Aux Address
Keyboard	"01"	"01"
Mouse	"01"	"60"

These auxilary addresses are type specific. After a slave RUPI receives its SDLC station address (short id) during configuration, all data packets to and from the addressed slave contain an appended auxiliary address in the data field of the packet. This allows the host and slave RUPIs to determine which device at the slave node is communicating. Not all addresses are pollable by the controller: Some addresses are receive-only. On the ECS keyboard, for example, the sound generator, key-repeat map, key-click map, and LED register have receive-only addresses.

Graphic-Input Devices Attach to Keyboard

The ECS keyboard features an unique enclosure. At each end are locking features for attaching other similar enclosures. Adding devices to the MIS such as an additional keypad, thumbwheels, joydisk or other graphic-input device is simply a matter of attaching application-specific enclosures to the right or the left side of the keyboard. Four-contact connectors at each side of these enclosures expand the MUIB by passing the signals and dc voltage between the connected units.

When an application calls for just one unit, such as a keypad, it may be operated alone on the bus.

All the attachable devices can be operated as standalone units, not physically locked to the other devices' enclosures. Standalone operation is implemented using eight-pin telephone connectors for connecting to the MUIB cable; these connectors are located on the back of each enclosure. When operating unattached, the locking features and the four-pin connectors are covered by snap-on end covers.

RS-232 and Other Applications

A box is being designed now for interfacing RS-232 devices to the MUIB. This unit, known as the MIS Interface Box, will be powered off the MUIB and will accept an RS-232-type device, such as a tablet or a quadrature-type mouse.

As new applications arise, new devices can be designed to operate on the bus. Only the application software needs to be upgraded to accept a new device.

For More Information

For more information, call David Williams, 685-2560 (61-217). □

ADDING LOCAL AREA NETWORK TECHNOLOGY TO TEK PRODUCTS

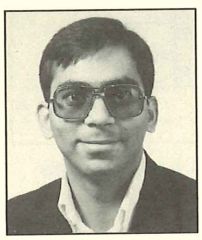

Atul Bhatnagar is manager of the Network Development Group of the Computer Science Center. He joined Tek in 1982 after receiving his MSCS from the University of New Mexico. He received his BSEE from BITS, Pilani, India in 1978 and was employed as a software engineer by DCM Data Products, India until 1981.

Now there is an economical way to add local area network (LAN) functions to a Tek product. By incorporating the Network Protocol Module (NPM), you will give your product both a front-end communications processor and industry-standard networking software. This easily retargetable software is independent of the host's software and hardware. The NPM was developed by the Network Development Group of the Computer Science Center.

As hardware costs plummet, more and more information processing is being distributed among specialized systems. These specialized systems are handling increasingly complex problems. This trend has made information exchange and resource sharing among distributed computing nodes very difficult. Recognizing these challenges, LAN technology is being aimed at fast and reliable exchange of information among interconnected computers. Today, commercial LAN technology will handle data transfer at rates from one to ten M-bit/sec.[1]

What does this mean to Tektronix?

It means most high-performance Tek products will need networking capability.

Providing LAN Functionality to Tek Products

Since Tek products have different hardware and software environments, networking software should be independent of both the host operating system and host hardware. This independence can be readily achieved by incorporating a standalone front-end communication processor. We designed the Network Protocol Module (NPM) with these concepts in mind. It contains networking protocols required for reliable end-to-end data communications. The NPM isolates host-specific hardware and software, thus minimizing the time needed to retarget the NPM to a Tek product.

The NPM software supports multiple industry-standard data-communication protocols and multiple connections. The NPM hardware consists of a Motorola 68000 microprocessor, 128 K-bytes of RAM, 128 K-bytes of PROM, two RS-232 ports, and a DMA interface to the local host.[2] Because the bus-dependent part of the NPM can be altered to suit a specific product's needs, the systems integrator has a complete data communications module to work with.

ISO Reference Model

The International Standards Organization (ISO) has divided networking functions into seven layers, each layer providing a well-defined interface to its adjacent layers and a specific networking service.[4] Each layer has two interfaces as shown in figure 1. The ISO seven layer model is shown in figure 2.

Physical layer

The physical layer transmits data over a communication channel by providing data signaling in the form appropiate to the physical medium.

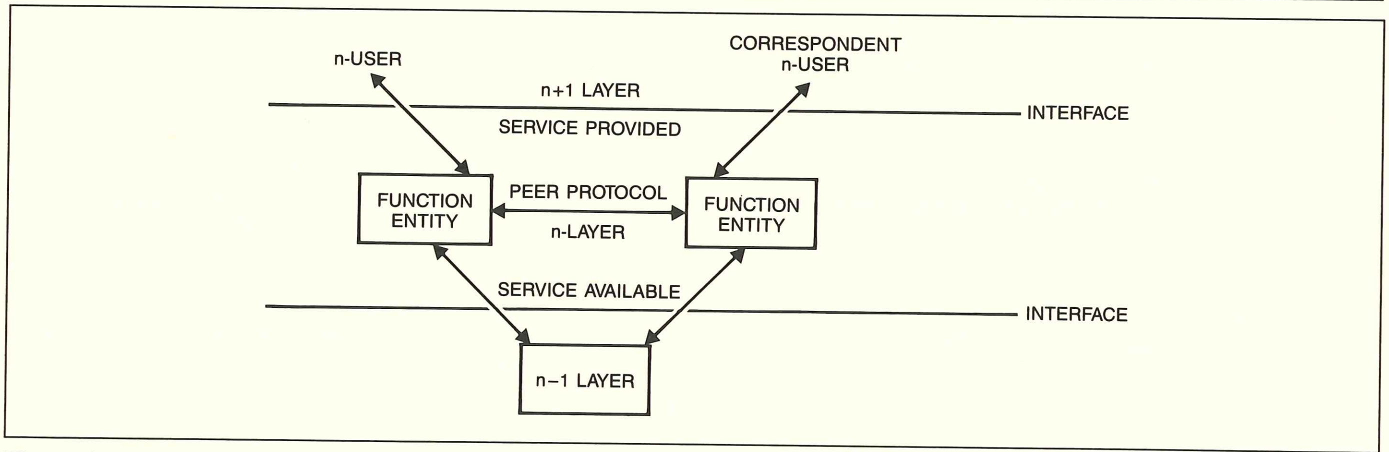

Figure 1. Network layer and its two interfaces. Each of the seven layers in the ISO model has two interfaces.

Data link layer
This layer supports data exchange between devices on the same medium. It also provides packet framing, checksumming, collision detecting, retransmitting, and station addressing.

Network layer
The network layer supports data communication over different networks by routing, segmenting, and reassembling packets as apporpriate to the network used.

Transport layer
The transport layer provides reliable end-to-end data transmission regardless of the number of networks between the end systems. This is done by providing sequencing of data, checksumming, windowing, and end-to-end data acknowledging.

Session layer
The session layer is a user interface into the network that enables a user to establish a connection with a process on a remote machine. The session layer manages the dialog in an orderly manner. It does this by providing tokens, synchronization points, and dialog units.

Presentation layer
The presentation layer performs user-directed transformations on the data. Some typical transformations are data encryption, data compression, and conversion of character codes.

Application layer
The application layer provides application tools such as file-transfer programs, virtual-terminal programs, and electronic mail.

Which Protocol to Use?

Presently, most systems in Tek and in industry follow the Transmission Control Protocol/Internet Protocol (TCP/IP) networking standards. These standards were established by the Defense Advanced Research Project Agency (DARPA).

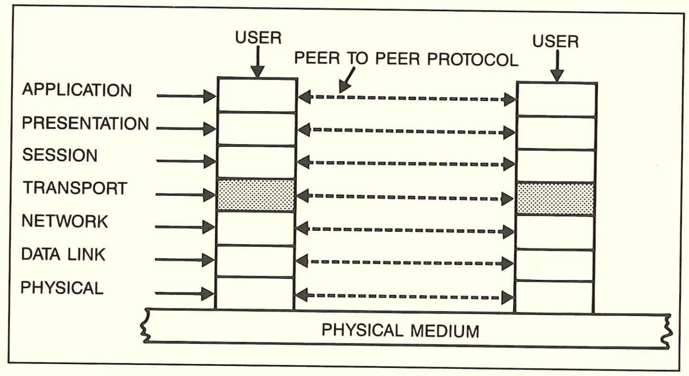

Figure 2. The ISO model has seven layers between the physical medium and the user.

The International Standards Organization (ISO) is working with the European Computer Manufacturers Association (ECMA) and the National Bureau of Standards (NBS) to define internationally acceptable networking protocols. When they are established, almost all major computer vendors and users will adopt the ISO networking protocols[5,6] in future products.

Because NPM supports both the TCP/IP and ISO protocols, the transition from the TCP/IP to the ISO standards will be straightforward.

NPM Hardware Environment

As shown in figure 3, the key features of the NPM hardware are:

- Multibus-I based
- Dedicated 8-MHz Motorola 68000 microprocessor
- DMA interface to the local host
- Two RS-232 ports
- IEEE-802.3 or Ethernet, supported by an Intel 82586 LAN controller
- 128 K-bytes of RAM
- 128 K-bytes of PROM

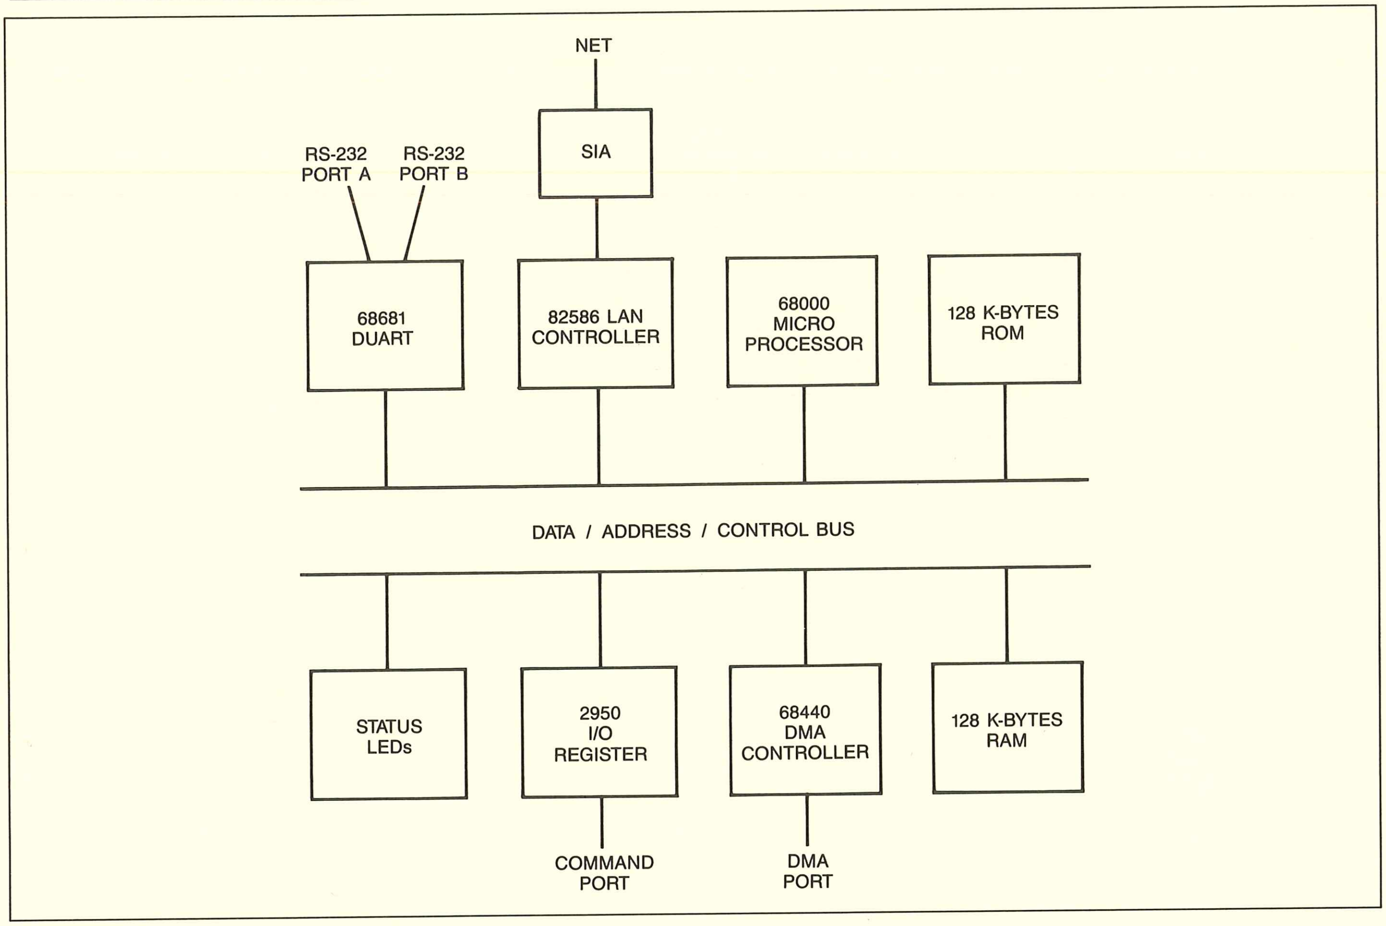

Figure 3. The NPM hardware environment.

NPM Software Environment

The NPM software environment (figure 4) is supported by a real-time operating system (RTOS), which provides multitasking, interprocess communication, memory management, and I/O driver integration. In this environment, each networking protocol layer is a separate process with an associated message exchange. This exchange enables processes to exchange information through interprocess messages.[2] Each process waits for an incoming service-request message at its message exchange.

To guarantee end-to-end data transfer, the NPM supports the concept of a pipeline. A typical pipeline consists of user-defined protocols for sending data over a particular connection. Hence, a protocol is bound to a pipeline at run-time by the user. This allows a user to specify the networking protocols at the time of establishing a connection.

Current NPM Processes

The *NPM-Host Interface* (NHI) transmits data to and from the local host. The local host and the NPM communicate via messages. Each message consists of a fixed-length header and a variable-length data portion. The fixed-length header contains the destination and the source of the message, the type of the request, a pipeline identifier, and the length of the accompanying data.[2] The NHI programs the hardware to transmit or receive data through DMA. When messages arrive at the NHI, they are dispatched to the appropriate process via interprocess messages.

The *ISO Transport Protocol* (ISO/TP) process provides transport layer services to the local host.[5] These services include:

- Sequencing of user-supplied data
- Checksumming end-to-end data
- Retransmitting host data
- Acknowledging received data
- Expediting data transmission

In short, the ISO transport protocol guarantees the reliable transmission and reception of data.

The *ISO Internet Protocol* (ISO/IP) process provides network-layer services to the network-layer user.[6] It routes packets of data over different networks; it also segments and reassembles packets.

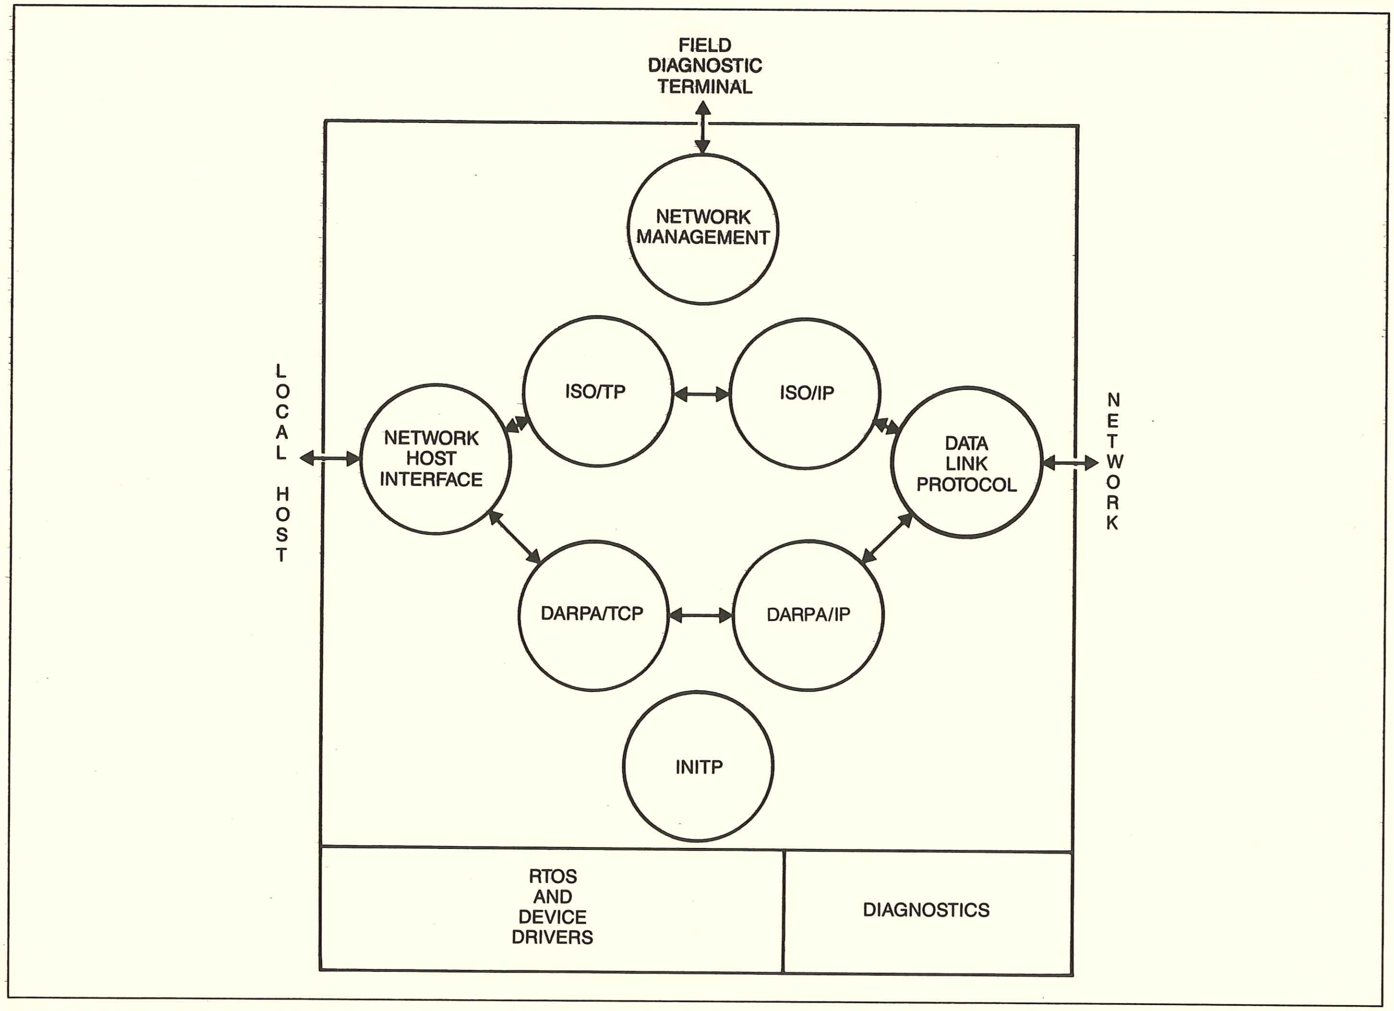

Figure 4. The NPM software environment. NPM supports both TCP/IP protocols and the upcoming ISO protocols. Since it is completely modular, NPM works in a wide range of environments. It is also easy to retarget to a different environment.

DARPA's Transmission Control Protocol (DARPA/TCP) provides functions similar to those of the ISO/TP.[7]

DARPA's Internet Protocol (DARPA/IP) supports functions similar to those of the ISO/IP.[8]

The *Network Management Protocol* (NMP) process provides facilities for gathering network statistics, benchmarking, tuning parameters, and performing extensive diagnostics. These diagnostics include:

- Internal and external loopback
- LAN chip testing
- Time-domain reflectometry
- DMA link testing
- RAM testing
- PROM testing

The Network Management Protocol can also be used to display the internal state of the NPM.[2]

The *Data Link Protocol* (DLP) process supports the IEEE-802.3[3] protocol as well as Ethernet protocol.

The *Initialization process* (INITP) is responsible for starting multitasking in the NPM. The INITP supports all the functions needed to reliably download the protocol software into the NPM. The downloading can be done via an RS-232 port, DMA transfer, or over the LAN. The INITP also runs extensive diagnostics before downloading.

Why NPM?

Too often, the process of copying data from one layer to the next slows data throughput. In contrast, NPM is direct. NPM avoids layer-to-layer data copying by simply exchanging the ownership of the buffers holding the data.

The NPM supports both the TCP/IP and ISO networking protocols. This means that the cost of retargeting network functions will be minimized as networks move from using TCP/IP protocols to the upcoming ISO protocols.

The NPM environment will accommodate any customer-specific protocol.

The NPM software is independent of host operating systems. Hence, retargeting NPM software to a different environment is simple.

The NPM offers a completely modular user interface that works in a wide range of environments.

NPM supports extensive network-management software; this will help users monitor, maintain, and tune their network's performance. This software will also help service organizations diagnose network problems from any physical point of the network.

For More Information

For more information, contact Atul Bhatnagar at 627-6833, (50-709).

References

[1] Tanenbaum, A.S., "Network Protocols," Computing Surveys, Dec. 1981.

[2] NPM Architecture document, Network Development Group, Computer Science Center, Tektronix, Inc., Beaverton, OR 97077.

[3] IEEE Standard 802.2, "Logical Link Control"; IEEE Standard 802.3, "CSMA/CD Standard"; IEEE Standard 802.4, "Token Bus"; IEEE Standard 802.5, "Token Ring"; IEEE Standards Office, New York, 1984.

[4] Zimmerman, H., "OSI Reference Model," IEEE Trans. Commun, COM-28, April 1980.

[5] DIS 8073, "N1990 Connection-Oriented Transport Protocol," American National Standards Institute, Washington, DC.

[6] DIS 8348 DAD1 N3152, "Commectionless Data Transmission," American National Standards Institute, Washington, DC.

[7] MIL-STD-1778, "Transmission Control Protocol," Department of Defense, Washington, DC 20301.

[8] MIL-STD-1777, "Internet Protocol," Department of Defense, Washington, DC 20301.

BEEN ASKED TO ORGANIZE A SESSION OR TO TALK AT A CONFERENCE?

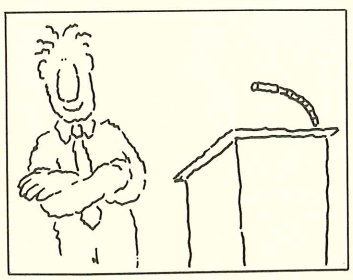

Professional conferences are an integral part of the information-transfer process. The success of such conferences depends on many things. The theme, the facilities, even the weather are important, but the speakers and the organizers are critical.

To help session organizers and speakers, Technology Communications 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 Professional Conference* will help you organize your talk and prepare effective slides to go with it. The rules in this booklet are few and simple, but they can help you look like a "pro" instead 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 Eleanor McElwee, d.s. 53-077 (642-8924). □

TIME SAYS, "A CLEAR MANUAL IS A THING OF BEAUTY."

***Time* Magazine, a few months back, published a piece about user manuals. TR, with *Time's* OK, reprints that piece here because *Time* gives the user's view rather then a dry analysis of what makes a good manual.**

How Does This #%*@! Thing Work?

No matter what happens, do not look at the manual

"You press the button, we do the rest." That marvelously simple slogan helped sell millions of Eastman Kodak cameras starting in 1888. Today, however, the owner of a new video cassette recorder or some other electronic wonder must turn to an instruction manual to get his machine working. But that is often when the trouble begins: the comsumer opens a booklet to find a compilation of jargon, gibberish and just plain confusion. "There is a major disease in this country called wall-stare," says Sanford Rosen, president of Communication Sciences, a Minneapolis consulting firm. "When people read a computer manual, they just want to put it down and stare at the wall for as long as possible."

Bad instructions are bad business as well as a torture to read. A maddening manual can cripple sales of products that might have been successful. Coleco lost $35 million in the fourth quarter last year partly because people flocked to return the initial version of its Adam computer, which the company offered for $600. In a statement to shareholders, Coleco blamed much of the consumer dissatisfaction on "Manuals which did not offer the first-time user adequate assistance." Observes Joseph Sugerman, president of J S & A, a mail-order house that specializes in high-tech merchandise: "Very often, items with the highest rate of return are those where customers are frustrated with the instructions." Coleco has reintroduced the Adam computer, complete with a new instruction manual.

Directions for hooking up and operating video cassette recorders can be particularly maddening. A frequent mystery is how to connect the machines to television sets and antennas. Owners must often pick their way through mazes of diagrams and technical terms like "One-touch type F connector" that seem to have been written for licensed electricians. Some manuals compound the confusion with illustrations that differ from the actual machine. Notes the 46-page booklet for a Panasonic OmniVision model: "Please be assured that this difference is not due to mistake but to ongoing product improvement."

Manuals for smaller, less expensive items can also be frustrating. Instructions for a Pulsar digital quartz watch ($59) go on for 13 pages before telling how to set the time. One Hewlett-Packard financial calculator ($110) comes with an operating booklet that runs to 246 pages of small type. The company supplements that with a 170-page training guide that sells for $15. "People have said we should do something like this for all our manuals," observes Janet Cryer, who wrote the guide.

Consumer electronics companies insist that customers are generally satisfied with the directions they get. "Over the past year the number of complaints we have received because of difficulties understanding our user's manuals would probably fit in one hand," says W.T. Collins, a vice president for consumer affairs at RCA. The firm's instructions used to be written by design engineers, but now they are prepared by technical personnel who train distributors in how to operate and service RCA products. Says Collins: "We realized that engineers have a tendency to make the content of a manual a bit too technical."

Various causes are behind impenetrable operating booklets. Some publications are slapped together quickly just as the product is about to be introduced. "Manuals are too often the last things that are done," says Communication Sciences' Rosen. The pressure is particularly intense in the fast-moving personal-computer industry. "A lot of the problem in that market is the haste to get the product out first," says Lois Schwartz, a New York City specialist in the preparation of instructions.

Some gadgetry from Japan and its Asian neighbors helps swell the confusion. Says Bob Budnek, a former Atlanta audio consultant: "The instructions are written in Japan, translated in Japan and printed in Japan, and sometimes the intention of making it clear to people in English does not come through." For example, the directions for one Japanese-made turntable cartridge advise, "Furthermore, cantilever would be damaged when the stylus guard is touched and detouched." Even simple points about simple products can get lost in translation. The instructions for Swimotor, a Hong Kong-made toy that pulls children through water, warn that "the user must every time pay attention especially to the time used with this machine."

A clear manual can be a thing of beauty and a joy forever. "Those that are well thought out make good reading," says Catalogue Merchant Joseph Sugerman. "They sound as if they were written by a teacher with plenty of patience who is aware of all the mistakes a consumer can make."

Many retailers are impressed with the manual for Apple's new Macintosh computer. Designed to be used with tapes and video displays, it guides Macintosh owners gently through a technological thicket. Says Chris Espinosa, 22, an eight-year Apple veteran who supervised the booklet's preparation: "A good manual is not a narrative; it is an outline or report. Nobody every reads a manual cover to cover – only mutants do that."

Fortunately, better manuals may be on the way. Leading technical schools like Rensselaer Polytechnic in Troy, N.Y., and Pittsburgh's Carnegie-Mellon have writing programs that teach students how to translate complex facts into clear directions. Enrollment in the classes is high, and instructors say that corporations have been snapping up their graduates.

But for now, at least, many consumers are likely to continue to find operating booklets more frustrating than enlightening. Indeed, some may feel like twisting the famous bromide "If all else fails, consult the manual" into a new admonition: "No matter what happens, do not look at the manual!"

By John Greenwald. Reported by Dorothy Ferenbaugh/New York and Carol Fletcher/Chicago

SYSTEM AVAILABLE TO CHARACTERIZE A/D CONVERTERS

A test system for characterizing the performance of high-speed A/D converters has been developed in the Electronics Systems Lab of ARL. They are making this system available for use by other Tektronix groups.

The system will test A/D converters having resolutions up to 8 bits and sampling rates up to 150 MHz. A sinewave curve fit test – with calculation of effective bits and S/N ratio – and a histogram test are available. Other tests can be implemented by the user, using SPS BASIC.

For more information or a copy of the user manual call Val Garuts (627-2514) or Erik Hultine (627-6162). □

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.

COMPANY CONFIDENTIAL

NOT AVAILABLE TO FIELD OFFICES

TECHNOLOGY REPORT

RICHARD E CORNWELL

mail label goes here

19-285

DO NOT FORWARD