COMPANY CONFIDENTIAL

Published March, 1979

FORUM REPORT 12

RELIABILITY

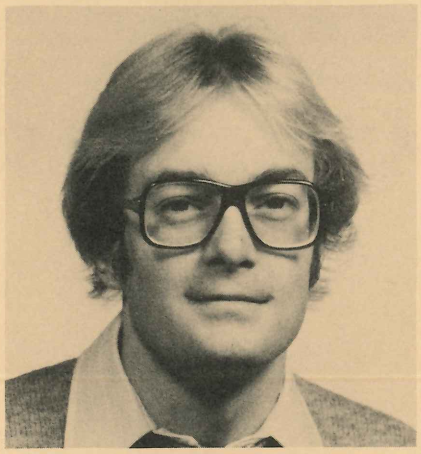

Tim Flegal

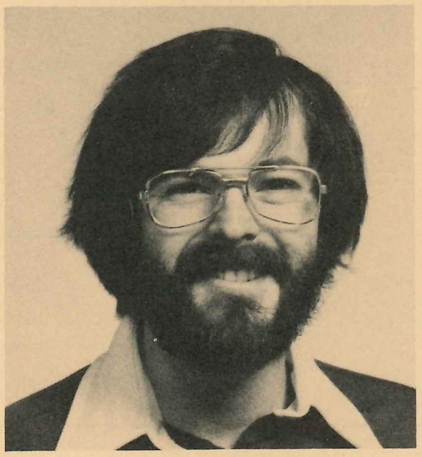

Mike McMahon

Bill Walker (executive vice president), in November 1976, formed the Engineering Activities Council to provide engineers with a forum in which to present directly, to multiple levels of management, what engineers themselves consider important in technology.

Forum 12, "Reliability Engineering," was held in January 1979 in the Technical Center (building 50) auditorium. Speaking at the forum were: Bill Pederson (Spectrum Analyzer Engineering), Bill Snell (Engineering Services), Bob Wallace (LID Operations), and Dan Wright (FDI Reliability).

Forum chairpersons were Mike McMahon (FDI Engineering) and Tim Flegal (TM500 Engineering).

RELIABILITY BASICS

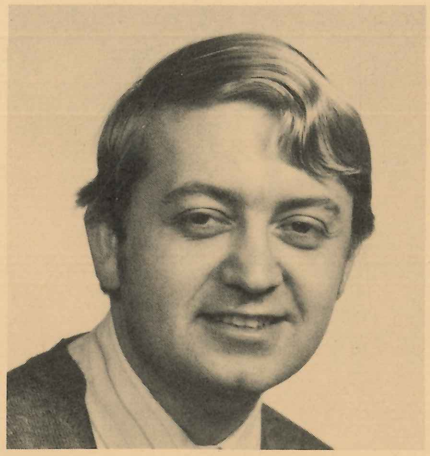

Dan Wright, FDI Reliability, ext. 7025 (Beaverton).

Although reliability engineering isn't a new discipline, only in the last few years has Tektronix recognized reliability as a full-time professional position. To make sure we are thinking of the same concepts, let's review some reliability basics.

Strictly defined, **reliability** is "the probability that a product will perform its intended function for a specified period under specified conditions." Informally this means the length of time a product performs without failing.

FAILURE TYPES

When we talk about failures, we are most concerned with customer field failures. Those failures have several causes. Customer abuse, especially inappropriate applications of equipment, is a major source of failures. Quality failures result from manufacturing defects such as pinched wires, loose hardware, and poor soldering.

The failures we are concerned with in this forum are **reliability failures.** These are failures that occur because of poor-quality parts or poor engineering application of good parts.

FAILURE MODE ANALYSIS

While reliability engineers are concerned about what failures occur and how many occur, their primary long-range concern is identifying *why* failures occur. **Failure mode analysis** is the process of "determining the reason a system failed by examining the failed parts to identify the cause of performance variations."

MEAN TIME BETWEEN FAILURES

A basic term used in many reliability discussions is **mean time between failure** (MTBF). MTBF is population size times the operating hours, divided by the number of failures. **Population size** is the number of units under test. **Operating hours** is the number of hours a product has actually been operating.

With very large systems (such as Tek's Cyber system), users think of "reliability" in terms of availability. **Availability,** is the probability that a product can operate at any given time. Availability, often called **up-time ratio,** is MTBF divided by the sum of MTBF and mean down-time. For example, a 95% up-time ratio means there is a 95% chance that a system will function properly at any one time.

PART FAILURE RATE

Another commonly-used term is **part failure rate,** which is the number of parts replaced (per unit of time) due to the failure of that circuit element (some parts fail because other parts in the circuit fail.) Like all engineers, reliability engineers use models to explain the phenomena they work with. A simplified example of a part failure-rate model is:

$$\lambda = (\pi_Q)(\pi_T)(\pi_S)(\lambda_b)$$

The π factors represent quality (Q), temperature (T), and stress (S). Lambda is the failure rate modified by these π factors.

The failure rate of a system is basically the sum of the failure rates of the system's parts:

$$\lambda_T = \lambda_1 + \lambda_2 + \ldots$$

An example will demonstrate how a reliability engineer uses the model. Refer to figure 1.

Using what we have learned, the failure rate for this circuit is the sum of the failure rates for the individual parts. If we assume that the failure rates for the individual parts are:

$$\lambda_{R's} = .1 \text{ failures/1000 hrs.}$$

$$\lambda_{C's} = .05$$

$$\lambda_Q = .5$$

(These failure rates are hundreds of times higher than real-life failures.) Therefore,

$$\lambda_T = 4(.1) + 2(.05) + 1(.5) = 1.0 \text{ failure/1000 hrs.}$$

$$\text{MTBF} = \frac{1}{\lambda_T} = 1000 \text{ hrs.}$$

PRECONDITIONING

Another term you hear frequently in reliability discussions is **preconditioning.** This means removing bad parts from an electrically or thermally-stressed parts lot.□

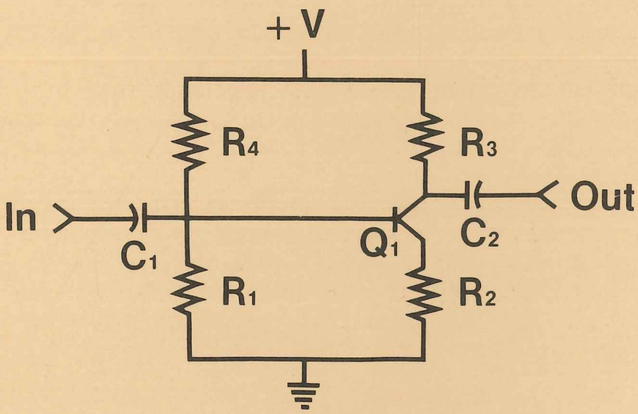

PREDICTED FAILURE RATE

$$\lambda_T = 4(.1) + 2(.05) + 1(.5)$$
$$= 1.0 / 1000 \text{ hrs.}$$

$$\text{MTBF} = \frac{1}{\lambda} = 1000 \text{ hrs.}$$

Figure 1. Computing the predicted mean time between failures for a simple circuit involves adding the failure rates of the individual parts (in this case, four resistors, 2 capacitors, and 1 transistor). The example failure rates (.1, .05, and .5) are hundreds of times higher than real-life failure rates.

HOBBY FAIR III ANNOUNCED FOR JULY

The Engineering Activities Council has announced that it will sponsor Hobby Fair III in July, 1979. Hobby Fairs are exhibits of engineering and G-job projects. Hobby Fair I (1977) and Hobby Fair II (1978) included only microprocessor projects, but Hobby Fair III is open to engineering projects in all disciplines.

Because exhibit space is limited, the Engineering Activities Council will select participants from among applicants.

If you would like to exhibit a project at Hobby Fair III, fill out the coupon below and mail it to Elske Cordell, 50-389. Applications must be in by June 1, 1979.

T&M Publicity has published reports describing Hobby Fair I and II. For copies, call ext. 6795 and ask for Forum Report 6 and Forum Report 11. □

Your name ____________________

ext. __________, D.S. __________

Your project's name ____________________

____

Brief description of your project ____________________

____

____

If you have questions about displaying your project, call Elske Cordell, ext. 6098 (Beaverton).

Component Failures, Causes And Controls

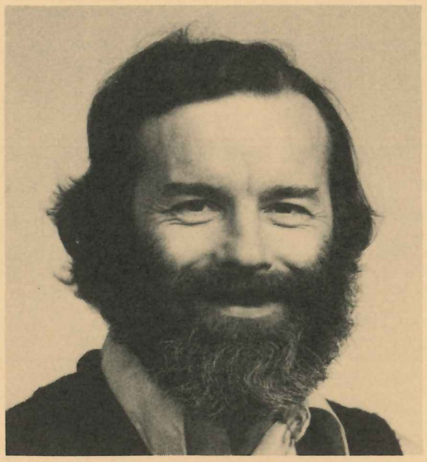

Bob Wallace, Lab Instruments Division, Operations Staff, ext. 7982 (Beaverton).

In discussing component failures, their causes and control, we'll talk almost exclusively about semiconductors, because they are the major cause of instrument failures. We will also examine a significant opportunity to control failure rates in future LSI-based products.

A well-designed product's quality and reliability are determined by the manufacturing process reliability and component reliability. Environmental variables alter a product's longevity, but a good design can protect a product from environmental variations as well as from most forms of abuse.

TYPICAL FAILURES

Figure 1 shows the failure distribution for lab scopes. Transistors account for 20% of all failures, and diodes, 15%. Tektronix IC's and workmanship are smaller sources of failure. The rest of the failures are due to resistors, capacitors, inductors, crts and other parts.

Figure 2 shows the failure distribution for logic analyzers. Here most failures occur in purchased IC's. Transistors appear to be in second place, but the figures are misleading because a design problem, not a component fault, caused the high transistor failure rate. Correcting the design problem greatly reduced the transistor failure rate.

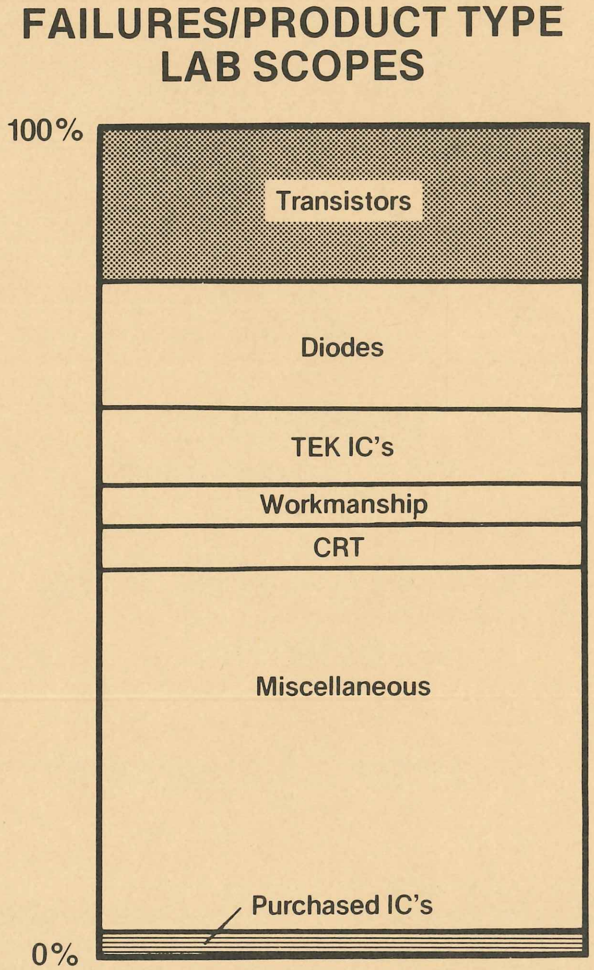

Figure 1. Failure distribution for Lab Scope components.

As shown in figure 3, the failure distribution for 8002A Microprocessor Lab components is dramatically different from lab scope and logic analyzer distributions. Designed by Millenium Systems, Inc., the 8002A uses a different component technology than the other two product lines. As a result, most 8002A component failures are due to purchased IC's and purchased subassemblies such as floppy disk drives.

CAUSES

Examining the causes of transistor and diode failures (refer to figure 4) reveals that apparent electrical overstress is the major cause of component failures.

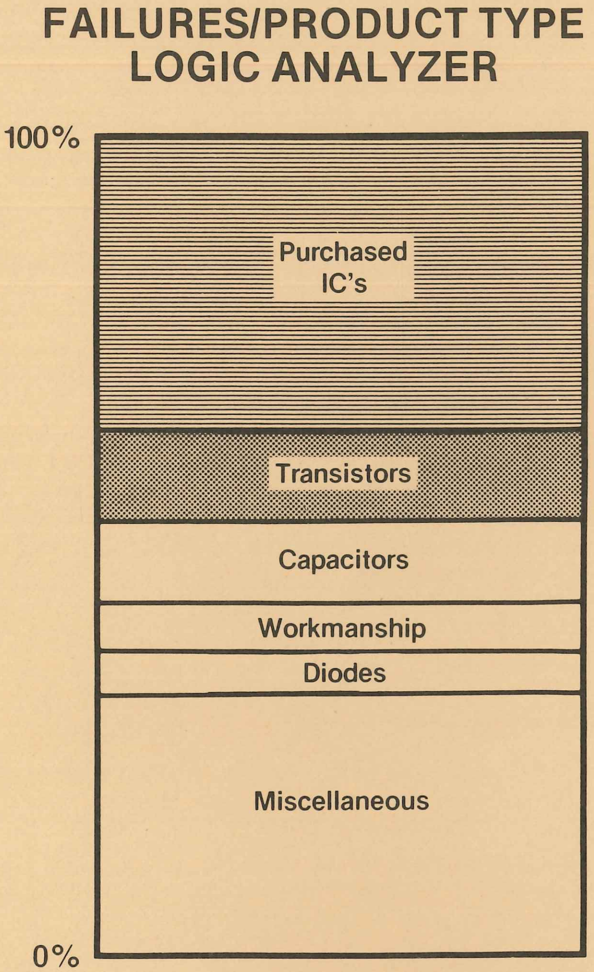

Figure 2. Failure distribution for Logic Analyzer components.

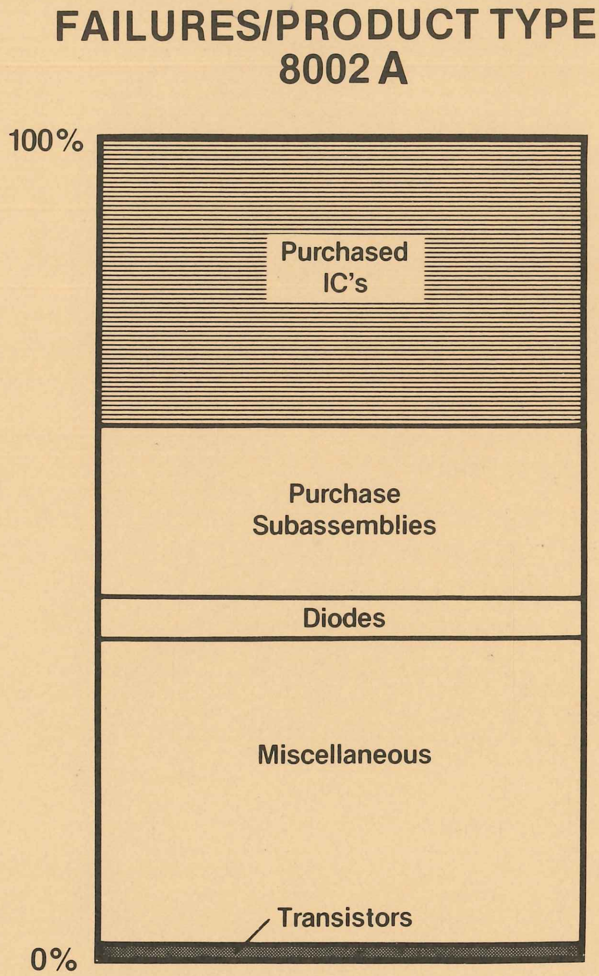

Figure 3. Failure distribution for 8002A Microprocessor Lab components.

The electrical overstress is "apparent" because these failures really come from component capability poorly matched to the applied stress (we may not know if the component was too weak or the stress too high). Failure analysis of overstressed parts often reveals little because little is left to indicate original part capability (the parts burn up).

"Retest-OK" components are components that service technicians incorrectly identify as failures, but which (1) are parts incompatible with other parts, (2) are parts rejected during "shotgun" troubleshooting, or (3) are failures that result from socket problems rather than actual component failures. (In "shotgun" troubleshooting, a servicer changes all components that are paired or grouped even if only one failed.)

Component vendors are at fault in the smaller failure categories. These faults almost always result from poor component packaging and processing.

As shown in figure 5, the failure-cause distribution for linear and small-scale integrated digital IC's is very similar to the distribution for transistors and diodes shown in figure 4.

However, if we look at more recent technology such as LSI circuits and memory IC's, the failure distribution is quite different from that of other components. Refer to figure 6. Many LSI failures are caused by surface and oxide defects rather than by improper applications or poor design. Large-scale IC's usually have complex metallization patterns, frequently including multilayer metallization. This intricate and delicate surface topology accounts for most inherent LSI failures. Fortunately we don't have to live with this high level of LSI failures. We do have alternatives.

SCREENING TECHNIQUES

We can screen-out surface and oxide defects with two commonly-used techniques. The first approach is high-temperature parametric testing which screens-out about 50% of the failures which would occur without the testing. This technique adds only about six cents to the cost of each part. High-temperature parametric testing is a very popular screening technique in Wilsonville.

Continued on page 6

Continued from page 5

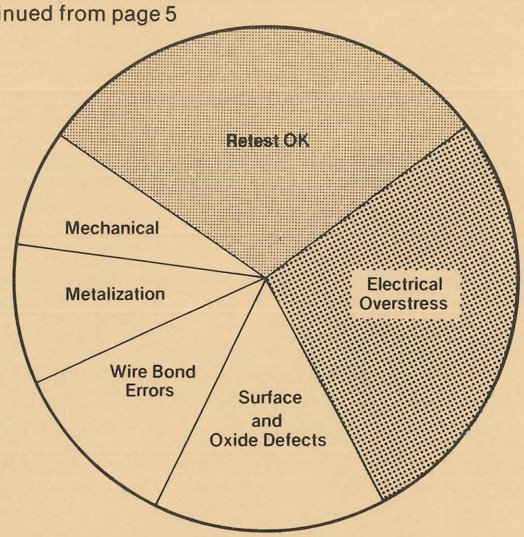

Figure 4. Causes of transistor and diode failures.

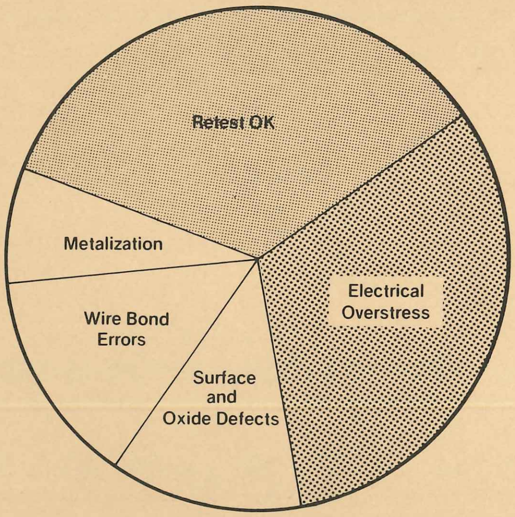

Figure 5. Causes of linear and small-scale integrated circuit failures.

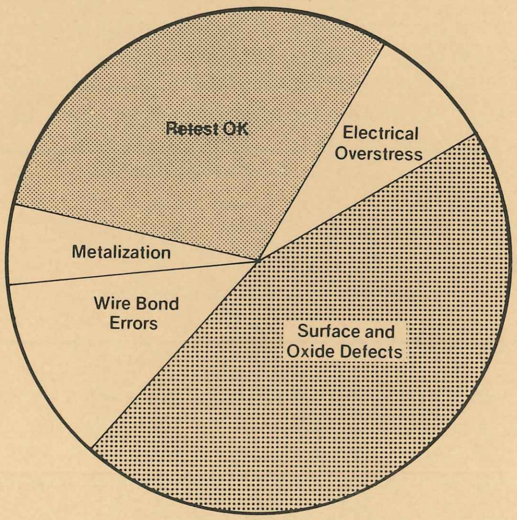

Figure 6. Causes of large-scale integrated circuit and memory IC failures.

A second oxide and surface defect screening technique is high-temperature accelerated burn-in, a process of powering-up a part in an aging rack, followed by parametric testing. Burn-in is 95% effective as a screening technique, leaving only about 5% of the failures that would occur with untested parts. Burn-in adds about 15 cents to the cost of each tested part.

Figures 7 through 10 further describe burn-in. Figure 7 shows a generalized time/temperature acceleration factor plot for semiconductors (the plot applies to most other components as well).

Figure 7 references a typical 40° C junction temperature (X1 acceleration). Suppose we first burned in a set of semiconductors at 90°C. For every hour of burn-in at 90°C, the parts "age" as much as they would in 100 hours at 40°C. How does this relate to the longevity of a typical bag of parts? In figure 8, the nearly-horizontal line represents the failure rate of the main population of a group of parts. This failure rate is very low and constant. If we could buy all our parts from this population, we would have excellent product reliability.

FAILURE DISTRIBUTIONS

With each parts shipment, we receive parts that have early failures which we call "infantile failures." (Refer to figure 9.) These parts often pass vendor testing and our manufacturing process. However, due to marginal flaws, these parts fail soon after customers receive the products. They are definitely bad parts.

We also receive a "freak" parts population with each shipment. See figure 10. These parts fail for reasons that may or may not be similar to infantile failures. The factor that sets freak populations apart from infantile failures is time. A freak population's failures cluster around one point in time.

Any given parts lot may contain more than one freak population, but each freak population's failure distribution centers around one point within the main population's lifetime. In a normal application environment, a freak population usually appears in the first thousand hours of use.

As shown in figure 11, combining failure rates for the infantile, freak and normal parts populations produces a typical composite curve of failure rate versus time. If we used accelerated burn-in so that the effects of the first 2000 hours of normal life are achieved in 200 hours or less then we could quickly "kill off" the unwanted populations.

The result would be a more reliable population of parts. Burn-in eliminates most failure-prone parts. From another perspective, with burn-in we move the first year of use beyond the life span of most infantile and freak populations.

In real life, each lot of parts is different. Some lots have very few weak parts while others may have a freak or infantile population for the majority.

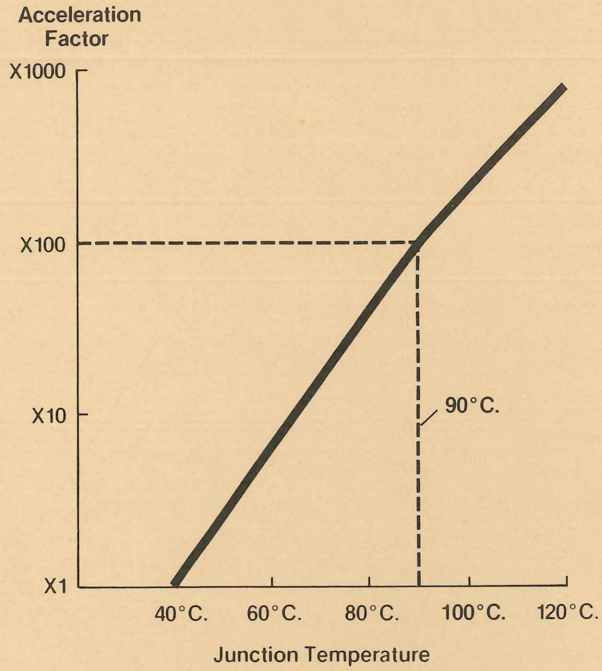

Figure 7. A generalized aging acceleration factor versus junction temperature for semiconductors.

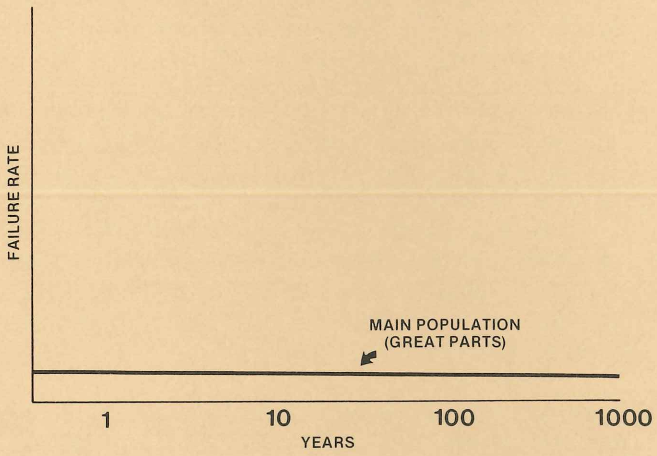

Figure 8. A plot of the failure rate for the main population of a hypothetical parts lot.

AN OPPORTUNITY

At Tektronix, we have an opportunity to lead our competitors in product performance by controlling our failure rates through careful circuit design and judicious use of screened parts. To achieve acceptable reliability, future designs must incorporate component burn-in or high-temperature parametric testing, techniques commonly used in the electronics industry.

ACKNOWLEDGEMENT

Ron Schwarz (Component Reliability Engineering) provided much information about component failure modes and screening. □

INFANTILE FAILURES

(really bad parts)

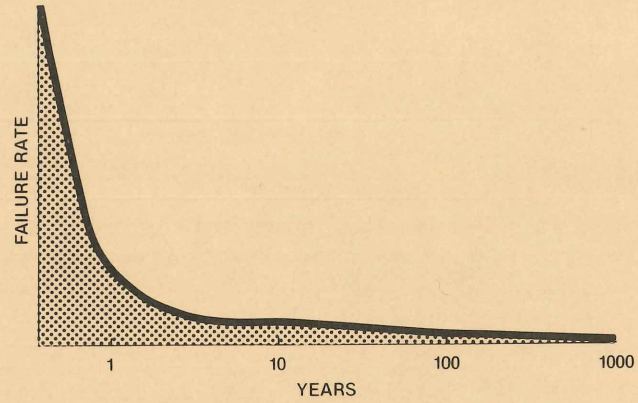

Figure 9. Infantile failures are parts that passed vendor testing and survived the manufacturing process, but failed shortly after our customers received their products.

FREAK DISTRIBUTION

(not-so-bad parts)

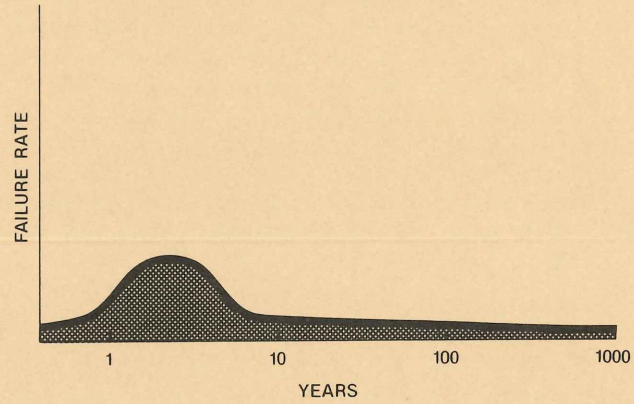

Figure 10. "Freak" failure populations cluster around a given point in time.

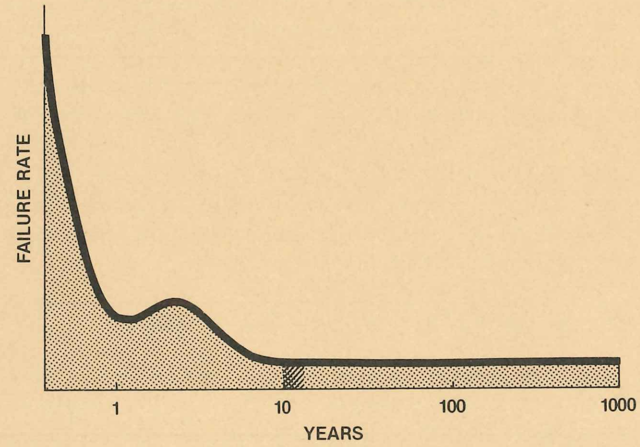

Figure 11. This curve, which incorporates the infantile and freak populations in figures 9 and 10, is typical of failure rates.

DESIGNING FOR RELIABILITY

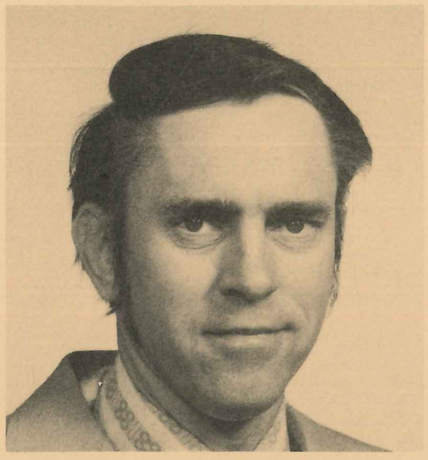

Bill Pederson, Spectrum Analyzer Engineering, ext. 7025 (Beaverton).

With our products becoming more complex, we can no longer simply accept whatever product reliability we get. Instead, we must design it into our products. There are several ways we can improve product reliability: (1) make reliability a specification, (2) create a closed-loop system that gives designers reliability feedback, (3) use accelerated-life testing to shake-out reliability problems, and (4) return at least some field failures to product designers and reliability engineers to fix and analyze.

SPECIFICATION

A reliability specification gives design engineers reliability goals, and allows designers to trade off reliability against costs and performance.

The 7L18 Spectrum Analyzer has a reliability specification. Because the first 7L18 production units did not meet the specification, Manufacturing postponed customers shipments. While there was some delay, we corrected some serious reliability problems. For example, after we replaced an "A" series CMOS part with a "B" series CMOS part, there were no field failures associated with the part. However, we had to trade off this increased reliability against the effects of shipment delays: added engineering costs, increased work-in-progress and canceled orders.

THE RELIABILITY ENGINEER'S ROLE

Every new project should have a reliability engineer working with design engineers. A reliability engineer can help define the reliability specification, and (if involved early in the project) can develop a reliability program for the product. A formal reliability program provides several benefits. First, the program gives designers reliability goals. Second, a reliability program enables designers to single out problems in past products, perform more reliability tests, and receive feedback early.

Having a reliability engineer involved early in a design project also helps the project move smoothly. For example, in packaging, a reliability engineer can define air-flow requirements and high-voltage supply spacing requirements to eliminate arcing. A reliability engineer can also make sure a product performs under conditions for which the product is designed. For example, even though both are portable instruments, a 492 Spectrum Analyzer requires considerably less mechanical durability than a 1510 Ground Fault Locator (which may be dropped into the back of a service truck).

CLOSED-LOOP SYSTEM

Another factor that can enhance product reliability is a closed-loop design system (refer to figure 1). In this system, we design an instrument, build a few, test for reliability, analyze all failures, redesign to prevent those failures, and start the loop again. With our current reliability testing procedures, the new product introduction phase system is too short to allow closing the loop. Waiting for test results takes too long. With an open-loop system, an instrument is usually in pilot (or even production) when the designer receives meaningful data from reliability tests. At this point, redesigning for reliability isn't always possible, and certainly isn't as cost-effective.

There are only two ways to close the loop. Lengthening the phase system is unacceptable (it's too long already). That leaves one option: accelerating the reliability testing.

ACCELERATED-LIFE TESTS

In Component Reliability (under Engineering Services), we performed high-temperature accelerated life tests of several component types. We also analyzed all component failures, both from accelerated-life tests and from field failures.

Since high temperature caused almost all field failures, I decided to perform high-temperature accelerated-life testing. We tested the 492 Spectrum Analyzer (a complex instrument) at 75°C, and the PA1 Cable Comparator (a simple instrument) at 125°C. In a third project, the ambient temperature was 170°C. That temperature was excessive. The tested units had

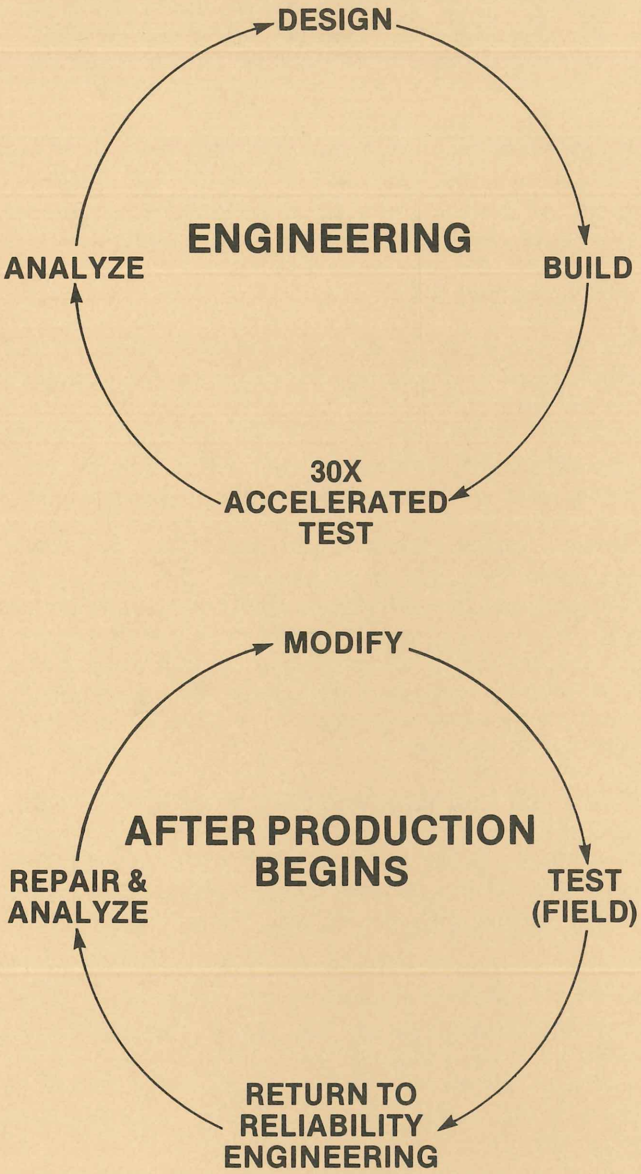

Figure 1. A closed-loop design system gives product designers reliability feedback before a product enters production. Even after production begins, a closed-loop system can help reliability engineers analyze failures and correct product design.

problems (the solder and wires melted) that presumably will not occur under normal operating conditions.

The Reliability Laboratory (in Engineering Services) performs reliability tests at lower temperatures than we used. The acceleration factor is lower with these tests (ranging from a factor of three at 40°C to six at 50°C) than with our accelerated tests (ranging from a factor of 30 at 75°C to 1024 at 125°C). However, Reliability Lab's tests are essential for uncovering failures not caused by high temperatures (humidity and vibration are two common causes of failure).

In our A-phase accelerated-life testing of the 492 Spectrum Analyzer, we found six failures (some due to design problems, and some to component failures). After each failure, we stopped the test, analyzed the failure, and redesigned the failed circuit. Then we restarted the test, evaluated the new design, and looked for new problems.

After 40 hours of testing at 75°C, six failures had appeared. Assuming an acceleration factor of 30, our accelerated tests were the equivalent of 1200 hours at 25 degrees. After the sixth failure, we continued the test for another 40 hours with no more failures. By contrast, Reliability Lab's low-acceleration temperature tests revealed one failure.

USING FIELD FAILURES

The closed-loop system should continue after Engineering Release. In this post-ER loop, we build instruments and, in effect, field experience tests them. Also in this loop, the factory fixes the failed instruments and returns them to customers.

Because few instruments are available for reliability testing, we must rely on field failures for some failure data. The goal of field service people is fixing failures as quickly as possible. That quick turnaround often doesn't allow time to identify causes.

From a reliability engineer's viewpoint, there are several reasons for factory repair of at least some field failures. Factory repair enables a reliability engineer to analyze problems and suggest redesigns. Another advantage of factory repair is that production engineers can also provide very quick turnaround (they are most familiar with servicing the instrument).

Reliability engineers often can not rely on field failure reports. About half of such reports don't specify which part failed. Also making the reliability engineers job difficult is the fact that about 30% of the "failed" parts meet all their specifications. There are two reasons for this: (1) the part isn't properly specified, (2) the part returned was not the part that failed (if many parts were changed, the wrong part may have been returned, or — in the case of integrated circuits and transistors — the sockets may be faulty). Field servicers, shooting for the quickest possible turnaround, often don't have time to analyze such "failures."

RESULTS

In Spectrum Analyzer Engineering, we used the post-Engineering-Release closed-loop system on the 7L18 Spectrum Analyzer project. The system worked very well. We discovered failures whose causes might never have been identified (and therefore corrected) without the closed loop. Our management was favourably impressed with our accelerated-life test results and closed-loop system, and plan to continue them. □

RELIABILITY AT TEKTRONIX

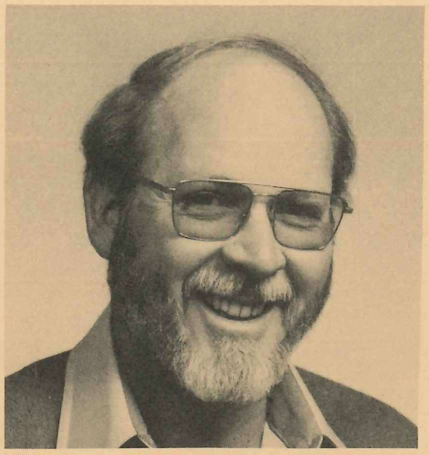

Bill Snell, Engineering Services (Operations), ext. 7897 (Beaverton).

About three years ago, reliability of Tektronix products became a major cause of complaints from the field. The result has been growth of formal reliability organizations within the company and adoption of new reliability programs by various product groups.

Unlike many companies which organize reliability functions into one corporate structure, Tektronix has taken a divisional approach in which the division operating manager is responsible for reliability of the division's products. Our current structure has both advantages and disadvantages.

ADVANTAGES

First, having reliability engineers close to responsible product engineering groups is a major advantage because they can better understand the products. A uniform reliability program may not benefit all products. Some reliability programs require different features than others. For example, eliminating integrated circuit sockets might be appropriate for one instrument but not for another. In any case, the trade-off's between reliability and costs aren't the same for all products. Reliability engineers need to be fully aware of these differences.

Second, as part of the division engineering "team," reliability engineers can better identify with product group goals, thereby more effectively integrating reliability into the design process.

Third, reliability engineers' familiarity with a product enables them to react quickly when problems occur and corrective action is needed.

Finally, whether reliability engineers are trying to influence new product designs or implement corrective action on products already in production, acceptance as one of the group is an important asset.

DISADVANTAGES

The divisional approach to reliability offers several disadvantages. First, after three years of operating independently, the divisions naturally do not present a common image to an outside observer. Several customers who recently visited Tektronix commented unfavorably on our lack of reliability coordination.

Second, quality control and reliability are almost completely independent at Tektronix. Neither function can be truly cost effective without close cooperation with the other. Even the best reliability program can't succeed unless the products are made with high-quality workmanship and materials. On the other hand, a poor reliability design can frustrate the best efforts to build a high-quality product.

Third, many projects (such as acquisition of large capital resources) require several organizations' cooperation. Some projects are difficult to implement when the various reliability organizations don't share objectives and priorities (which is most of the time).

Fourth, support groups cannot effectively budget for reliability needs. Reliability engineers can't affect support budgets and are not part of the justification process for additional resources such as money to buy additional component failure analysis equipment to provide the support that reliability engineers have requested.

Fifth, with responsibility for reliability widely dispersed throughout the company, measuring overall effectiveness is difficult. In fact, if one were to ask "what are we getting for our money?," determining what has been spent would be very difficult.

PROGRESS

Tektronix has made some progress in reliability efforts. We now have a Tektronix S-3260 Semiconductor Test System in Incoming Inspection (Central Manufacturing). The S-3260 tests about 150 part types (about one million parts per year). Equipment ordered in the last six months will significantly upgrade Incoming Inspection's ability to inspect parts.

Several support groups have been formed since 1975. Among them are Test Engineering and Component Pre-Conditioning (Central Manufacturing), Component Reliability Engineering (Operations), and Failure Analysis (Engineering Services). Related

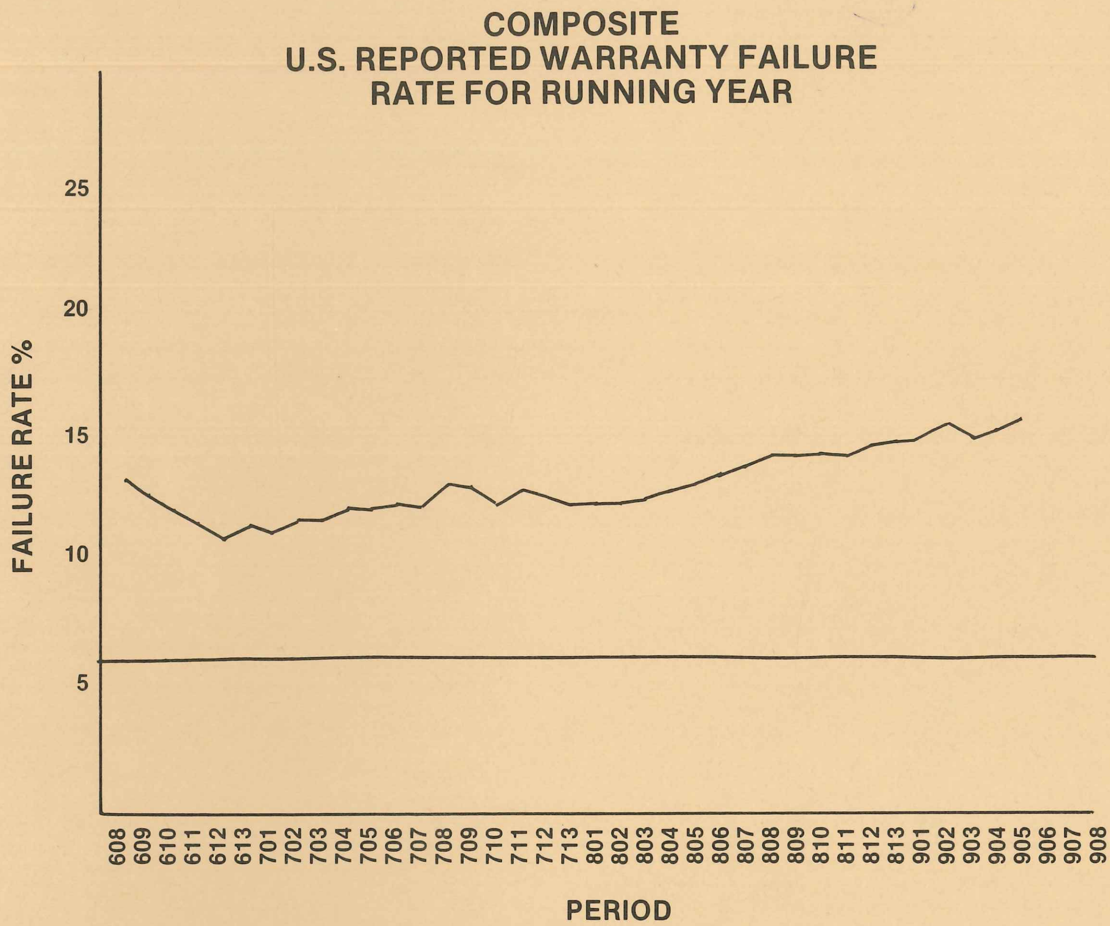

Figure 1. As shown in this moving annual average of field failures, failures in Tektronix instruments are rising. Several factors may be contributing to this increase. One such factor is outgoing product quality which has been dropping.

groups, such as Reliability Information Services (Engineering Services), have been strengthened.

The various divisions are more frequently employing designing-for-reliability programs (Bill Pederson discusses such a program elsewhere in this Forum Report).

In the last couple of years, Tektronix has greatly expanded field failure reporting. We now capture the total service record, and we are publishing several new reports.

A computerized component information system is well underway. Managed by Harriet Krauss, Component Information Systems will assemble into one data base everything that is known about a component. Users will be able to access the data through the Scientific Computer Center's Cyber system terminals.

PROBLEMS

Despite progress, problems remain. Average failure rates have not improved. As figure 1 indicates, failure rates are increasing. There is no simple explanation for

Continued on page 12

MAILING LIST COUPON

☐ Engineering News
☐ Forum Reports
☐ Software News
☐ ADD
☐ REMOVE
☐ CHANGE

Name: ____________________

Old Delivery Station: ____________________

New Delivery Station: ____________________

Payroll Code: ____________________
(Required for the data processing computer that maintains the mailing list)

MAIL COUPON TO: 19-313

Continued from page 11

this. However, there are indications that out-going product quality may be responsible for a large share of the slippage. Dan Harper (Quality Assurance) audited the manufacturing quality of finished products and found that the average defect rate is currently about 18%. For an equivalent period in FY700, the defect rate was about 10%. If Tektronix can return the defect rate to 10%, field failures should drop about 5%.

A second problem is that we have no vendor rating system to enable us to identify and screen out vendors who consistently ship poor quality parts.

Third, Tektronix has made little progress in adjusting parts inventory levels to enable us to reject bad lots. Often we have to use bad lots because replacements are not available.

No quality control of farm-out components is another problem. Farm-out parts are a major source of defective material on the production lines. Poor control of a circuit board cleaning process recently caused us to scrap several hundred finished boards that could not be re-worked. Several departments at Tektronix perform no incoming inspection of boards produced by outside contractors.

IMPROVEMENTS NEEDED

Design engineers and reliability engineers need clearly-defined company reliability objectives. An overall strategic plan is also badly needed to avoid dissipating resources and to allow the company to focus on issues affecting more than one division. A plan combining elements of each division's objectives, defining required resouces, and setting priorities would be a significant benefit. ☐

To add your name to the **Forum Report** mailing list, or change your delivery station, fill out the coupon inside.

Forum Reports *can not be WOW'd. However, unused* ***Forum Reports*** *may be sent along with other non-WOWable paper, in large quantity, to Material Evaluation at D.S. 71-474. Non-WOWable paper should be marked "SALVAGE."*

Managing editor: Burgess Laughlin, ext. 6795, D.S. 19-313. Compiled and edited by the Technical Publications Department for the benefit of the Tektronix engineering, software, and scientific community in Beaverton, Grass Valley, and Wilsonville.

Graphic design: Joan Metcalf. Typesetting: Jean Bunker.

COMPANY CONFIDENTIAL

(NOT AVAILABLE TO FIELD OFFICES)

FORUM REPORT
MAUREEN L KEY
60-553