MASSACHUSETTS INSTITUTE OF TECHNOLOGY

APOLLO

GUIDANCE, NAVIGATION AND CONTROL

Approved: Eldon C Hall Date: 5/26/67
ELDON C. HALL, DIRECTOR, DIG. DEV.
APOLLO GUIDANCE AND NAVIGATION PROGRAM

Approved: David G. Hoag Date: 6 Jun 67
DAVID G. HOAG, ASSOCIATE DIRECTOR
INSTRUMENTATION LABORATORY

Approved: Ralph R. Ragan Date: 6 June 67
RALPH R. RAGAN, DEPUTY DIRECTOR
INSTRUMENTATION LABORATORY

E-1944

THE IMPACT OF THE FLIGHT SPECIFICATIONS ON SEMICONDUCTOR FAILURE RATES

by
Jayne Partridge
and
L. David Hanley, Jr.

June 1967

MIT INSTRUMENTATION LABORATORY
CAMBRIDGE 39, MASSACHUSETTS

COPY # 97

ACKNOWLEDGEMENT

This report was prepared under DSR Project 55-23850, sponsored by the Manned Spacecraft Center of the National Aeronautics and Space Administration through Contract NAS 9-4065.

ABSTRACT

The procurement, screen and burn-in, and field history of the semiconductor parts in the Apollo Guidance Computer (AGC) is given. Both field failures and variability of performance through screen and burn-in are directly related to changes occurring in the parts manufacturer's facilities. The problems of developing and sustaining high reliability are discussed.

by:

Jayne Partridge
and
L. David Hanley Jr.

June 1967

THE IMPACT OF THE FLIGHT SPECIFICATIONS ON SEMICONDUCTOR FAILURE RATES

INTRODUCTION

The requirements of space technology have forced a new emphasis on parts reliability. For most applications failure rates of 0.1 to $0.01\%/10^3$ hours have been considered sufficient. Today with the advent of space technology, increased reliability has become necessary and a new term, "high reliability", has evolved. This term has come to represent a failure rate of $0.001\%/10^3$ hours at 90% confidence. To demonstrate this failure rate, it is necessary that no failures be generated in 2.3×10^8 component hours of test. Therefore, 10,000 parts must operate for 31.5 months without a single failure, or for one allowable failure the sample of 10,000 parts must operate for 53 months. Proving this failure rate is too expensive and too time consuming. The failure rate is usually demonstrated only after-the-fact in operating equipment. The more conventional methods of solving this dilemma have resulted in the battle of statistics which loses sight of real parts.

The Apollo program also had to face this dilemma. Large numbers of parts were subjected to extensive life testing and extensive computer field history was gathered. Over 330×10^6 operating part hours have been accumulated in operating computers, where some of the computers have been operating for over two years. Integrated circuits have been operated up to 31,000 absolute hours at maximum rating. As many as 400,000 integrated circuits have been purchased, tested, and stressed. Generated failures have been analyzed. The accumulated data will demonstrate that inadvertent changes occur with time. Comparison of the integrated circuits with the simpler transistors and diodes will further illustrate the subtleties of the elusive high reliability goal. A compilation and analysis of all the Apollo data forces strong conclusions concerning semiconductor procurement and reliability.

The method by which the awesome task of assuring high reliability was approached for the semiconductors in the Apollo Guidance Computer (AGC) has been previously presented[1]. This report will show that it is possible to reach the goal if the previously described requirements are met, and will further emphasize the problem of inadvertent changes which must be combatted before high reliability can truly be achieved.

INADVERTENT CHANGES

The semiconductor manufacturing industry is a classic example of "state of the art" processing where it can be shown that, with the exception of poor design, all detected failure modes can be eliminated by improved and controlled manufacturing processes. A dynamic, "state of the art" process will be subject to inadvertent process changes. Although inadvertent changes contribute greatly to increased failure rates they are often overlooked as major contributors to the decay of quality and reliability.

Inadvertent process changes are caused by the following:

1. Insufficient technical knowledge or understanding such that all the variables of the process are not known or their contributions and ramifications are not thoroughly understood.

2. Incomplete documentation where all of the details of the process, if understood are not completely or precisely specified in the manufacturing or process control documents.

3. Inadequate in-house control because the process control techniques do not guarantee the requirements of the process control documentation.

4. Inadequate control of suppliers where the procurement, inspection, and acceptance of incoming material is unsatisfactory.

5. Personnel. People are not specified as part of the process but they notably influence the interpretation and execution of procedures. Therefore, a personnel change in itself generates inadvertent changes.

6. "Improvement" changes. The manufacturer sincerely believes the instituted change will improve the quality and reliability of the product. However if the manufacturer has not thoroughly studied the effects of the "improvement" change on the process, he usually finds he has "traded a headache for an upset stomach", where he knows how to treat the headache but doesn't know how to treat the upset stomach.

Although the above list may appear to be self-evident, interesting subtleties are implied --- particularly when the goal is high reliability.

At this point it might be informative to illustrate some of the problems that were triggered by inadvertent process changes. A dramatic example occurred when an increase in integrated circuit failures manifested itself in metalization opens at oxide steps. Upon inquiry, the vendor claimed that no changes had been instituted. However, on closer examination, the vendor discovered that his oxide steps were steeper. Further study revealed that the photo-resist supplier had improved the photo resist such that greater line resolution, a constant objective of the industry, was possible. But, thinning of metalization crossing oxide steps was aggravated by the heating under power, resulting in opens at these oxide steps which had to be corrected by increasing the metallization thickness. Unfortunately, the change in the photo resist introduced a long-time-dependent failure mode which the vendor could not detect before shipping the product. Insufficient understanding by the vendor of his process and inadequate control by the vendor of his supplier were the basic causes of this sad experience.

Incomplete or imprecise specification of even the simplest of processes has been a source of much grief. An example of such a process is the wash cycles following the etching of silicon slices. At first, process specifications only indicated that slices be washed after etching. At the time, the production supervisor understood the requirements and instituted an effective wash cycle even though the process was not completely covered by documentation. When the particular production supervisor and some of his workers left for other jobs (a common occurrence in the semiconductor industry) and when new personnel with different backgrounds and experience took over, the entire handling procedure changed. It soon became obvious that the minimum wash time had to be specified. At a later date, in the natural course of events, production volume was increased. More slices were being processed, utilizing personnel of lower job classification. Because of the increase in volume, further experience showed that the number of slices washed at one time and the frequency of wash water change were important. Such process changes are not trivial when striving for high reliability.

Inadequate control of the manufacturing process presents a continuous source of problems. For example, integrated circuits are hermetically sealed to protect them from adverse environments. Although the environment was specified, some manufacturers were not controlling the environment during the sealing process and atmospheres inducing metallization corrosion were sealed in. The long-time-dependent failures caused by aluminum corrosion were directly attributable to this sealed-in water and oxygen.

People are never specified as part of a process but they do constitute the means by which a device is made. When striving for a high reliability product, motivation of the personnel building the product becomes an important factor. This has been dramatically shown when a separate production line was set up for the Apollo gate. The appreciation of the application generated greater care. The importance of the application gave status to the job. Both of these intangibles resulted in a measurable difference as will be shown later by comparing the failure rates of integrated circuits with the failure rates of diodes and transistors. In the latter case separate high reliability lines were not established.

Many changes which have been instituted for the sake of improving the product have led to long periods of instability and non-uniformity. Many examples can be given of "Improvements" which put programs depending on high reliability parts out of business. Even if the changes are for the better, inadvertent changes result, as the previous example of the photo resist illustrated. Changing plant locations, personnel, material, or processes for the sake of "improvement", might over a long time period, result in the production of a better part but during the transient period the process was not the same and the reliability of the parts produced is not at the same level as the parts that preceeded or followed the transient. In the meantime some well-meaning reliability engineer is trying to predict a failure rate based on past performance of the devices.

SEMICONDUCTOR PERFORMANCE AND FAILURE RATE DATA

Integrated Circuits

The hard, cold fact that vendors vary with time can best be demonstrated graphically. The numerical variations through screen and burn-in of the TO-47 single nor gate, one integrated circuit from one qualified vendor, is summarized in Figure 1. The details of the screen and burn-in procedure have been previously reported[1], but a flow diagram is given in Figure 2 for handy reference. In Figure 1, each plotted point represents a shipment lot of 2000 to 5000 nor gates; the predominant lot size was closer to 5000 parts. The percentage of catastrophic failure is plotted against a linear time scale with the total number of parts processed superimposed.

Before analyzing the variations of Figure 1 in detail, it might be interesting to point out some marked correlations of the data trends with occurrences at the vendors operation. For example:

> 1963 represents the tail end of the vendor's learning curve.
>
> There was no buying of the nor gate between June and October of 1964 thus necesitating a temporary discontinuation of the production line. Reinstatement of the line produced a new region of instability during the latter part of 1964.
>
> The two isolated points shown in mid 1965 indicate shipments from a noncontinuous production line.
>
> After June of 1965, part of the vendor's production line was moved to another geographic location. The history of this move is reflected in the data after October of 1965.

The screen and burn-in data, failure analysis, and field experience of the lots prior to July of 1964 inspired the flight specifications described in Reference 1. This was because it was found that field failure rates and failure modes were directly related to the number of failures and failure modes generated during screen and burn-in. The acceptance levels were adjusted so that 80 % of the lots would easily meet them. An example of a lot which would not have met the flight acceptance criteria was lot 418 (a date code). The flight specifications were then applied to shipments after October of 1964. As more computers have been built using these nor gates, extensive field failure rate data has been generated. It now remains to be seen if the field failures were accurately predicted by the flight acceptance criteria.

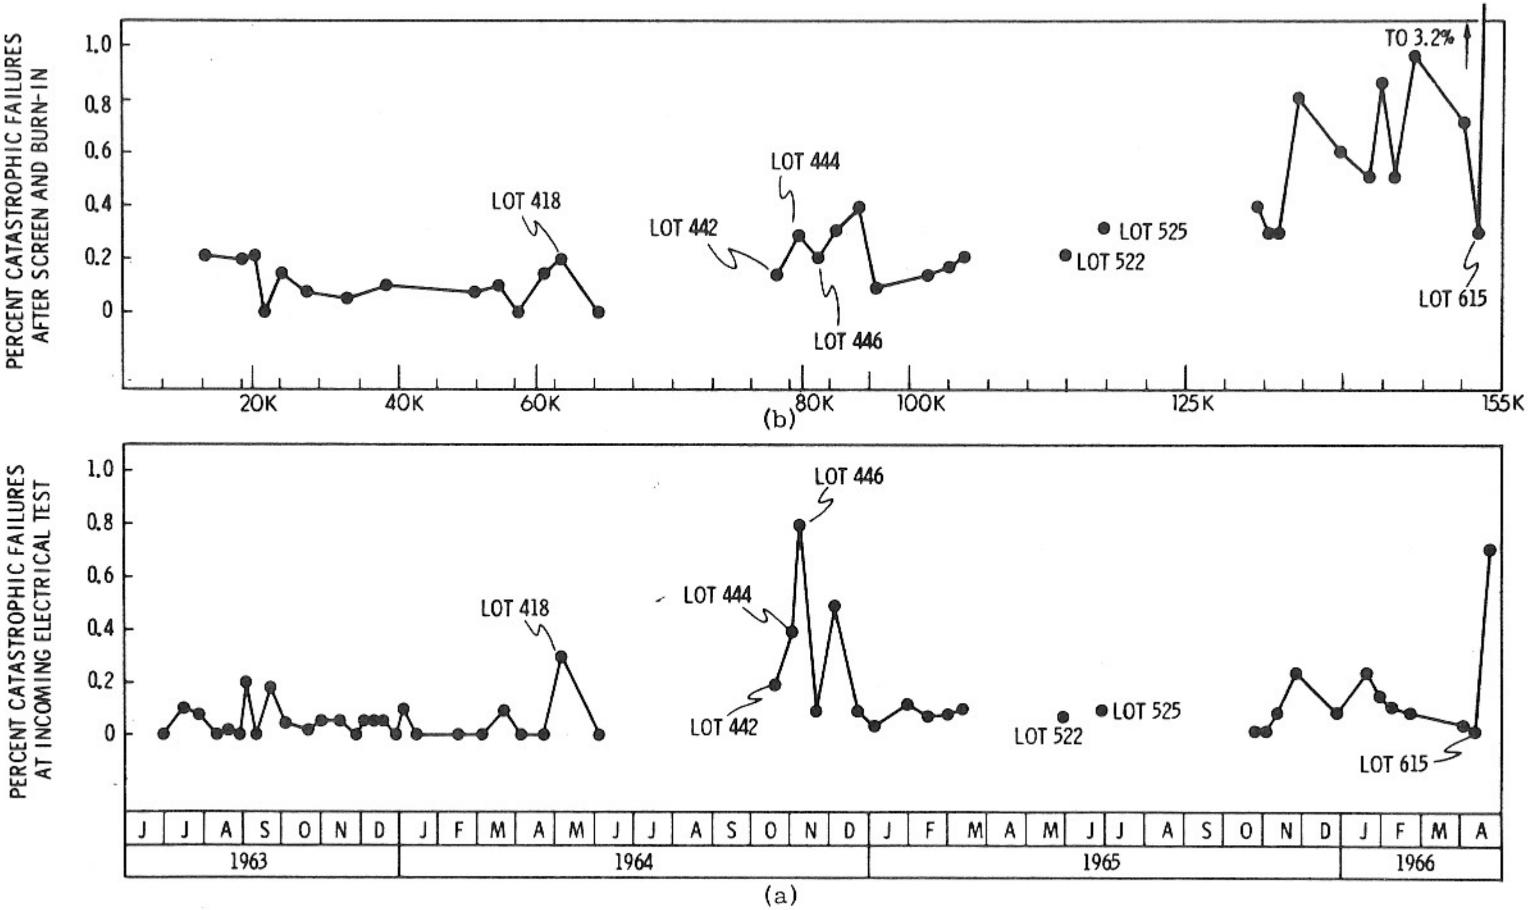

Fig. 1 A Vendor's Single Nor Gate Performance Through Screen and Burn-In Versus Time.

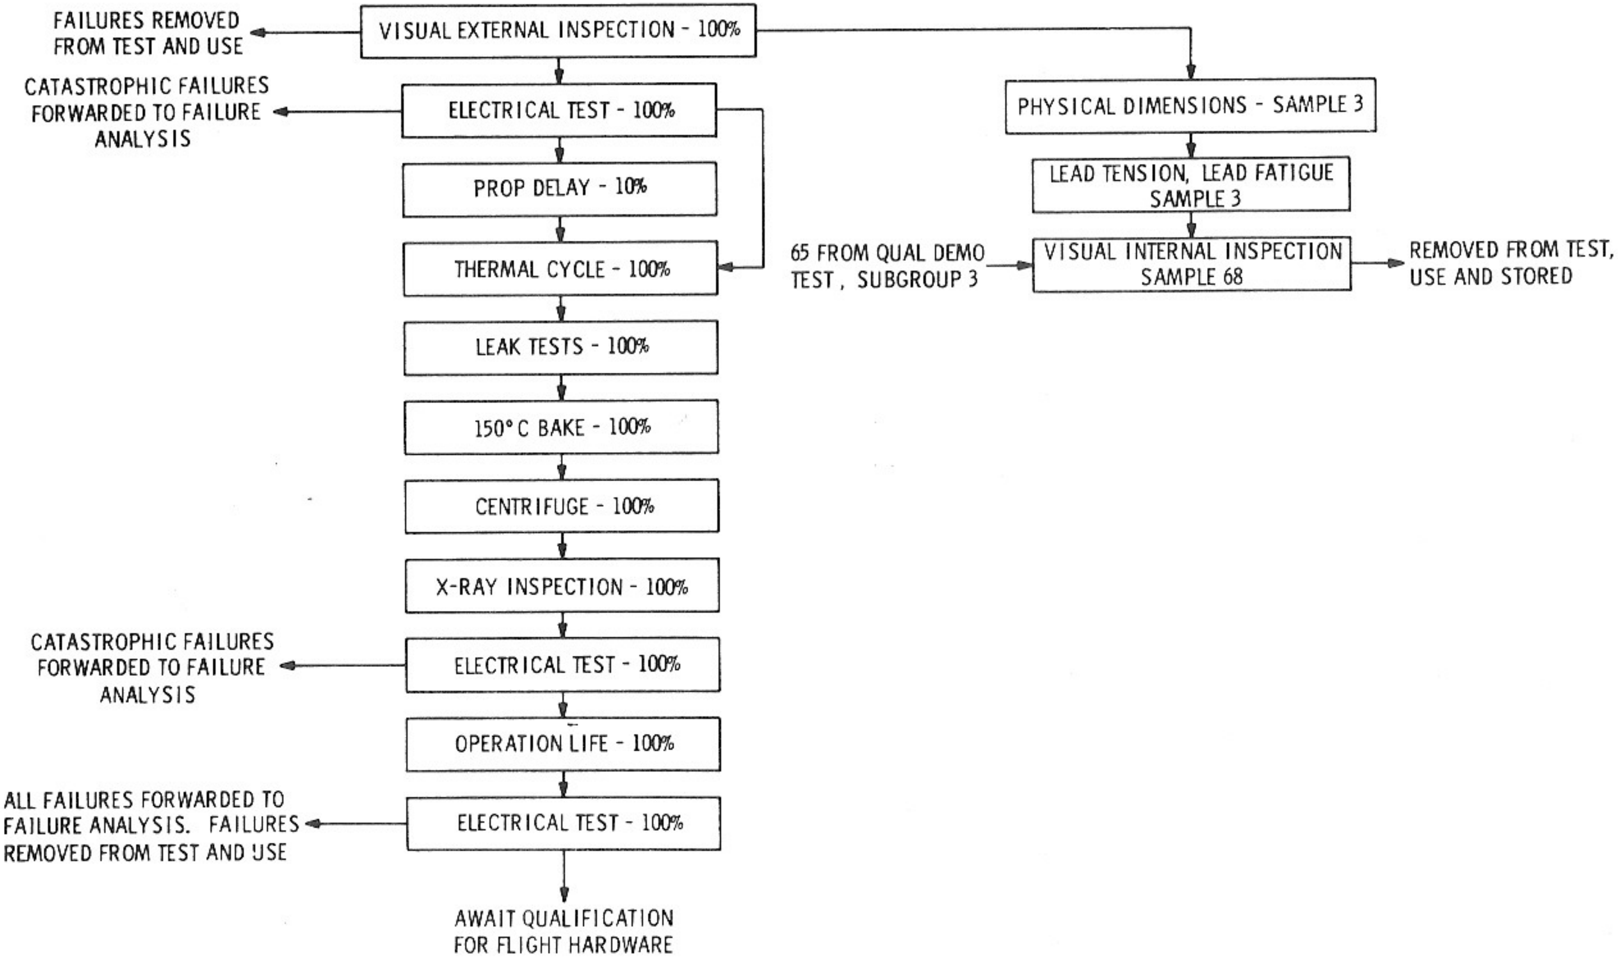

Fig. 2 Screen and Burn-In Flow Diagram for the Single Nor Gate.

Seven field failures have been generated and the lots of which these failures are members are shown in Figure 1. The description of the failure and the flight status of the lot from which they came are as follows:

Lot 418

The erratic operation was caused by a metal sliver or peeling of sufficient length extending normally from the internal surface of the can thus shorting the leads to ground. This lot would have failed the requirements of the flight specifications, electrically as well as visually. The visual samples exhibited loose gold particles and loose pieces of cap material. Unfortunately, loose pieces of cap material and particles were attributed to the opening of the package.

Lot 442

An open was generated in the aluminum interconnect at an oxide step where the metallization was excessively thin. The failure mode, "disappearing aluminum", was most probably caused by the formation of Al_2O_3 + Si triggered by excessive local heating [2,3]. This lot failed the requirements of the flight specification. Opens in the aluminum at oxide steps was one of the more predominant failure modes causing lot rejection.

Lot 444

A large, loose nickel inclusion caused a power-to-ground short. The source of the nickel was most probably the cap. The lot failed both the electrical and visual requirements of the specifications. Loose conducting particles were detected visually in some of the screen and burn-in electrical failures.

Lot 446

Shorting from input to ground occurred as a function of pressure applied to the cap. Examination revealed two metal peelings inside the top of the can. This lot failed the flight specifications. One of the reasons for lot rejection was the occurrence of intermittent electrical failures. The cause of intermittency was not determined, but based on past symptoms, the analyst hypothesized conducting particles as the cause.

Lot 522

A gold input lead shorted to the edge of the chip, thus causing intermittent shorting to ground. It appeared that a blunt tool could have sufficiently depressed the gold wire prior to capping to cause intermittent shorting. This fault is difficult to detect.

The lot passed the electrical requirements of the flight specifications but failed the visual requirements by only one unit. The failure mode described was not detected during the screen and burn-in procedure. It is interesting to note this was one of the lots exposed to 100 % X-Ray after a Y_1 centrifuge for lifted chip and excess lead length. The depressed lead could have been shadowed by the other leads and posts. Because the lot came so close to passing and because the field failure mode had not been detected in this lot, this failure will be counted as though the lot passed flight specifications.

Lot 525

An intermittent open due to an open thermocompression bond caused by insufficient pressure and/or temperature applied during the bonding operation. This particular lot failed the flight processing specifications because of an excessive number of open bonds caused by underbonding.

Lot 615

Failure was due to an open thermocompression bond caused by insufficient pressure and/or temperature applied during the bonding process. This particular lot failed the flight processing specifications because of an excess number of open bonds caused by underbonding..

All of the lots in Figure 1b up to November 1965 have been built into flight computers. After November 1965, only two of the better lots were selected for flight computers and Lot 615 was one of these.

From the field data generated, up to October of 1966, it should be obvious that the techniques of the flight processing specifications had predicted six of the seven field failures.

The single nor gate in the TO-47 package was the integrated circuit used for the logic of the Block I Apollo Guidance Computer. The Block II Computer incorporated a dual nor gate in a flat package. The screen and burn-in summary for the one vendor who supplied the dual gate is given in Figure 3. The flow diagram for the screen and burn-in procedure to which the dual gate was exposed is given in Figure 4. Here again, up to October of 1965 during which time the first 10,000 dual gates were shipped, the vendor had gone through a learning period. Thereafter, with the close cooperation of customer and vendor coupled with the aid of continuous buying, the vendor has maintained an exceptionally low reject level.

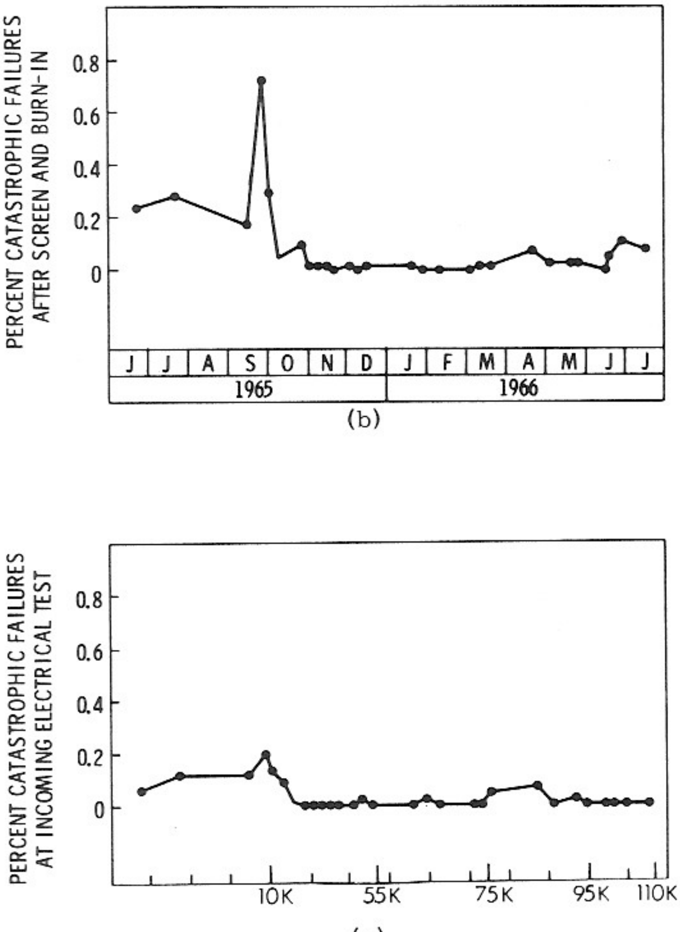

(b)

(a)

Fig. 3 A Vendor's Dual Nor Gate Performance Through Screen and Burn-In Versus Time.

Although the lots comprising the first 10,000 shipment did not pass the flight processing specification requirements, the parts were used in non-flight applications. To date, one flat-pack failure, because of a part fault, has been generated and it was a member of the very first lot shipped. The failure was due to a high series resistance at the aluminum-to-silicon contact. Shrinking geometries necessitated by demands of increased speed, coupled with insufficient etching of SiO_2 at the small emitter and base contacts were the major manufacturing contributants to increasing series resistance. The lot did not pass the flight specifications because it exhibited more than an allowable number of this time-dependent failure mode. Corrective action instituted by the vendor at that time has eliminated the problem.

At this point it might be interesting to compare the failure rates of the various populations of single nor gates shipped by the vendor of Figure 1 and the dual gates shipped by the vendor of Figure 3. Table I summarizes this data.

TABLE I

NOR GATE FAILURE RATES AT 90% CONFIDENCE UP TO 1 NOVEMBER 1966

	Gate Operating Hours In Computers	# of Failures	Failure Rate %/ 10^3 Hours
Total Block I & II Nor Gates	363,000,000	7	0.0032
Block II Dual Nor Gates Lots which passed flight spec.	56,000,000	0	0.0040
Total Block I Single Gate	307,000,000	7	0.0038
Block I Single Nor Gates Lots which passed flight spec.	256,000,000	1	0.0015
Block I Single Nor Gate. Lots which did not pass flight spec.	51,000,000	6	0.02

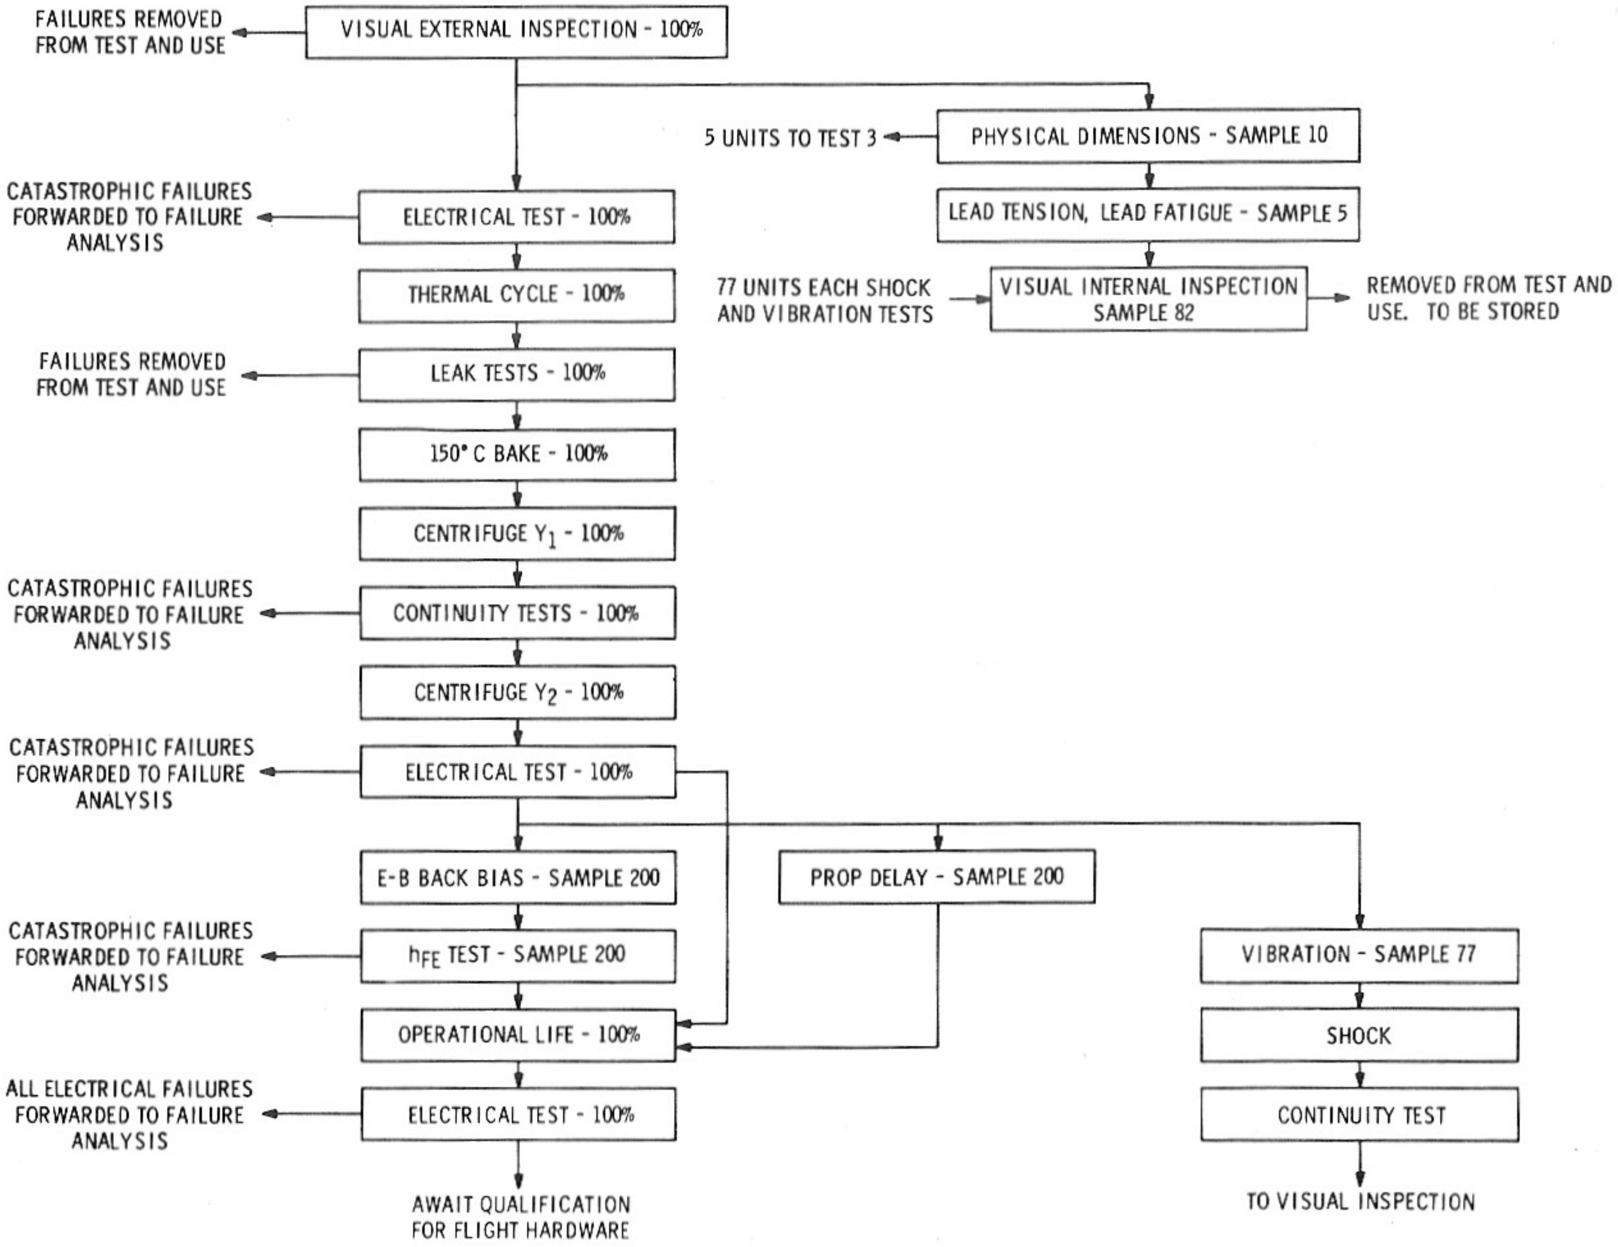
FAILURES REMOVED FROM TEST AND USE
VISUAL EXTERNAL INSPECTION - 100%
5 UNITS TO TEST 3
PHYSICAL DIMENSIONS - SAMPLE 10
LEAD TENSION, LEAD FATIGUE - SAMPLE 5
77 UNITS EACH SHOCK AND VIBRATION TESTS
VISUAL INTERNAL INSPECTION SAMPLE 82
REMOVED FROM TEST AND USE. TO BE STORED
CATASTROPHIC FAILURES FORWARDED TO FAILURE ANALYSIS
ELECTRICAL TEST - 100%
THERMAL CYCLE - 100%
FAILURES REMOVED FROM TEST AND USE
LEAK TESTS - 100%
150° C BAKE - 100%
CENTRIFUGE Y_1 - 100%
CATASTROPHIC FAILURES FORWARDED TO FAILURE ANALYSIS
CONTINUITY TESTS - 100%
CENTRIFUGE Y_2 - 100%
CATASTROPHIC FAILURES FORWARDED TO FAILURE ANALYSIS
ELECTRICAL TEST - 100%
E-B BACK BIAS - SAMPLE 200
PROP DELAY - SAMPLE 200
CATASTROPHIC FAILURES FORWARDED TO FAILURE ANALYSIS
h_{FE} TEST - SAMPLE 200
VIBRATION - SAMPLE 77
OPERATIONAL LIFE - 100%
SHOCK
ALL ELECTRICAL FAILURES FORWARDED TO FAILURE ANALYSIS
ELECTRICAL TEST - 100%
CONTINUITY TEST
AWAIT QUALIFICATION FOR FLIGHT HARDWARE
TO VISUAL INSPECTION

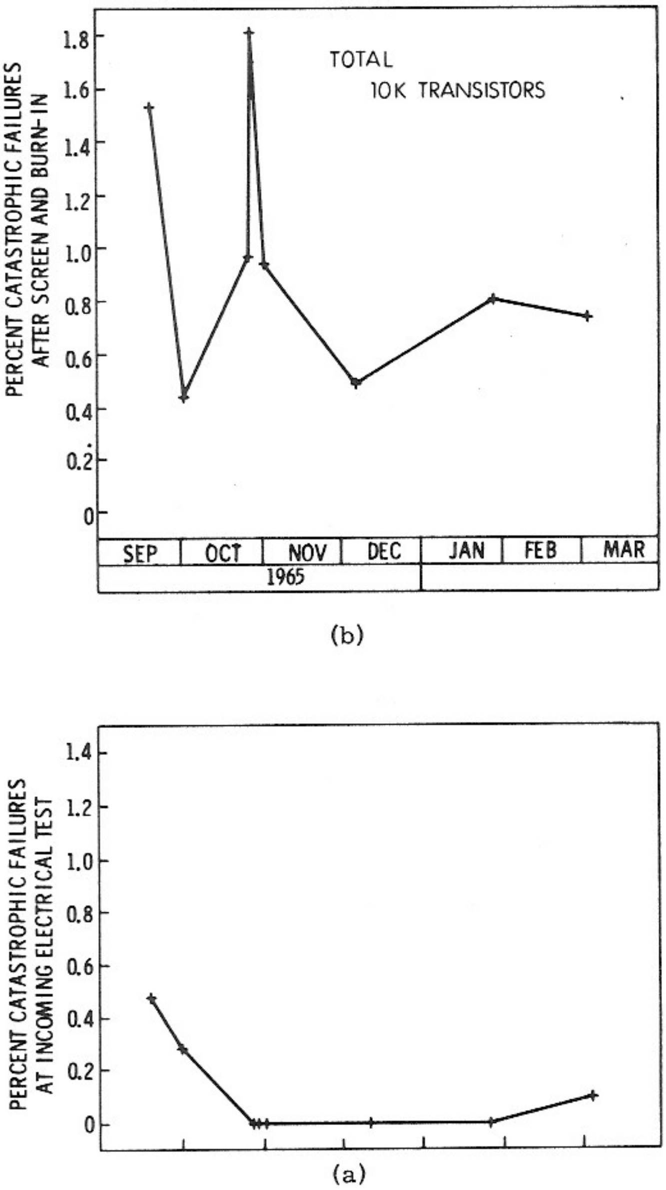

Fig. 5 A Vendor's Transistor Performance Through Screen and Burn-In Versus Time.

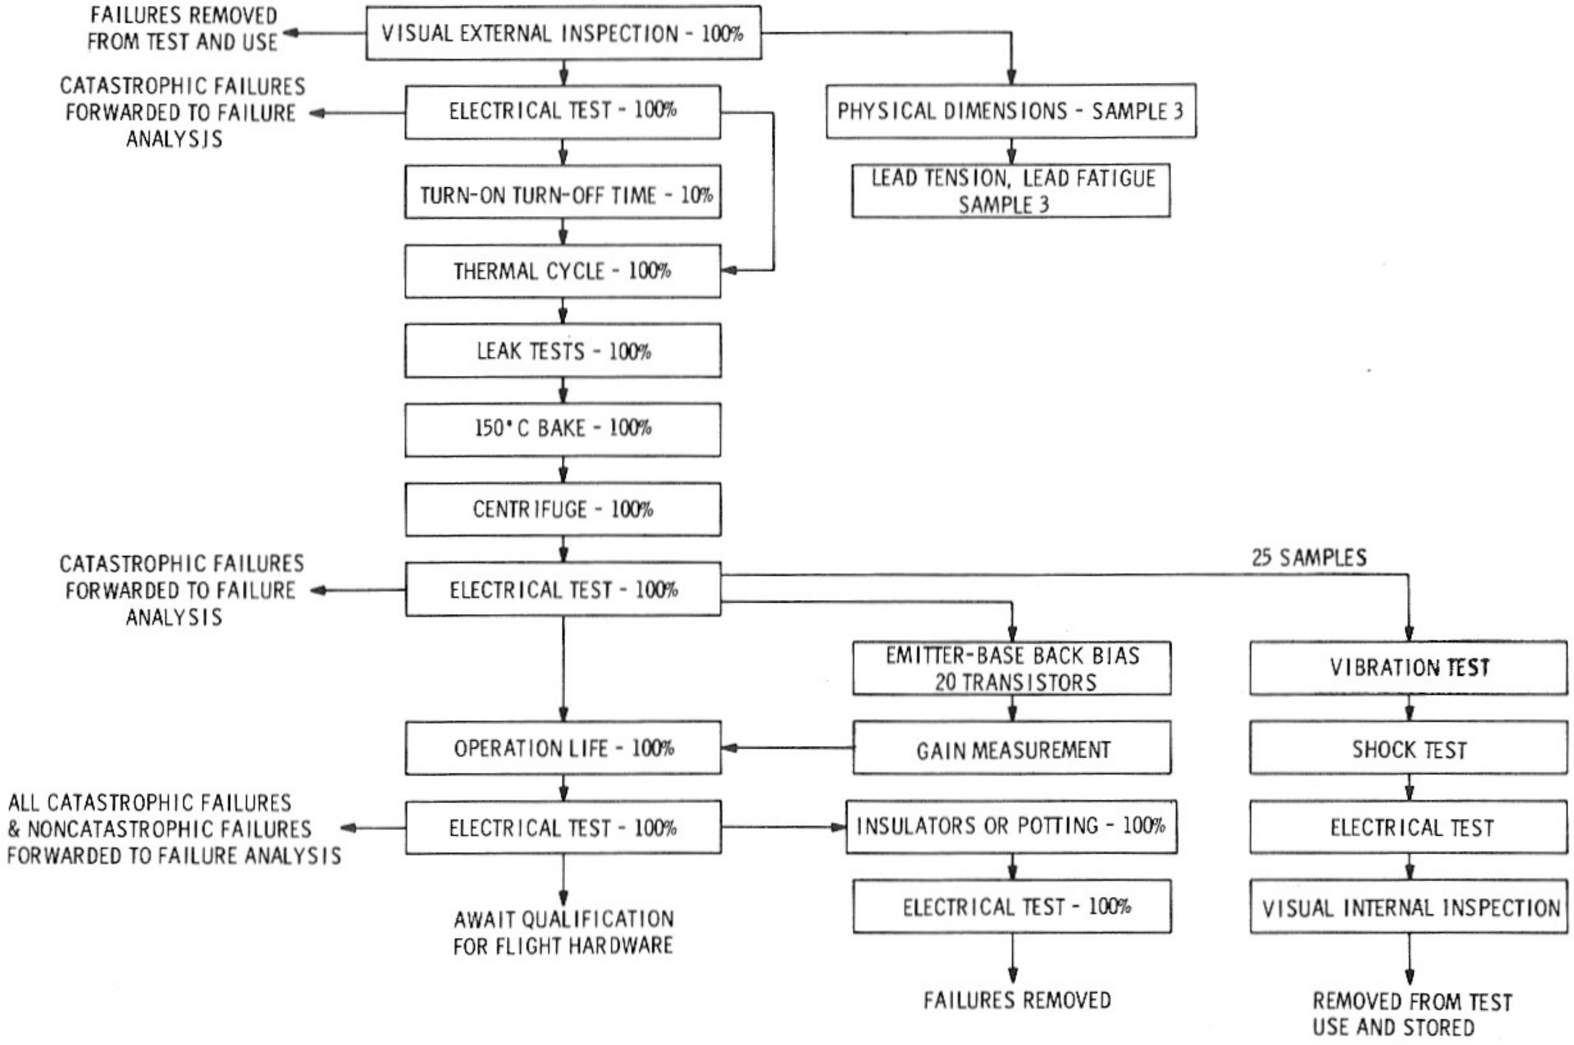

Fig. 6 Screen and Burn-In Flow Diagram for Transistors.

The single gate in the TO-47 package shipped from a single vendor is separated into two populations, those that were screened and burnt-in but failed the requirements of the flight process specifications and those that were screened and burnt-in and passed the requirements of the flight specifications. It is seen that in field operating computers there is a difference by a factor of more than an order of magnitude between lots that passed the flight specifications and those that did not.

The dual gates used in the AGC Block II design all passed the flight specifications and have accumulated as much field use time as the single gates which did not pass the flight specifications. The difference as shown in Table I in the number of failures generated and the resulting failure rates speaks for itself.

Transistors

As indicated in Figures 1 and 3 the single and dual nor gates were purchased in large volume. There was need of other semiconductor parts in the AGC, for example the transistors. It is interesting to compare the discrete transistor to the integrated circuit, which contains three transistors and four resistors.

The screen and burn-in results of medium power silicon planar transistors shipped from a single vendor is given in Figure 5. The flow diagram of the screen and burn-in procedure is given in Figure 6. The stress levels are indentical to the stress levels to which the integrated circuits were exposed with the exception of operating life. The transistors were, of course, operated at higher power levels, but still within the rating of the device. It is seen that the parts show a greater percentage fallout after incoming electrical test than either the single or dual nor gate. It has been shown[1,2] at least for integrated circuits, that the percent fallout during screen and burn-in is directly related to field failure rates. The transistors of Figure 5 which are being used in Block II design also show this relation. At this writing, these transistors have accumulated negligible operating element hours but five catastrophic failures have already occurred after screen and burn-in. These five failures were detected at the module level where the reliability clocks (which record accumulated time from final sale) have not yet been turned on. The failures were caused by a failure mechanism detected during the screen and burn-in procedure and that particular failure mechanism was the basis for the transistor lots failing the flight specifications. Because of small sample sizes, negligible operating field history has been accumulated to date, for the Block II transistor supplied by the vendor of Fig. 5.

Unfortunately the Block I transistors have to be treated differently from the Block II transistors, for two reasons. First, a flight specification was not imposed on the Block I transistors and secondly, the transistors were supplied by three

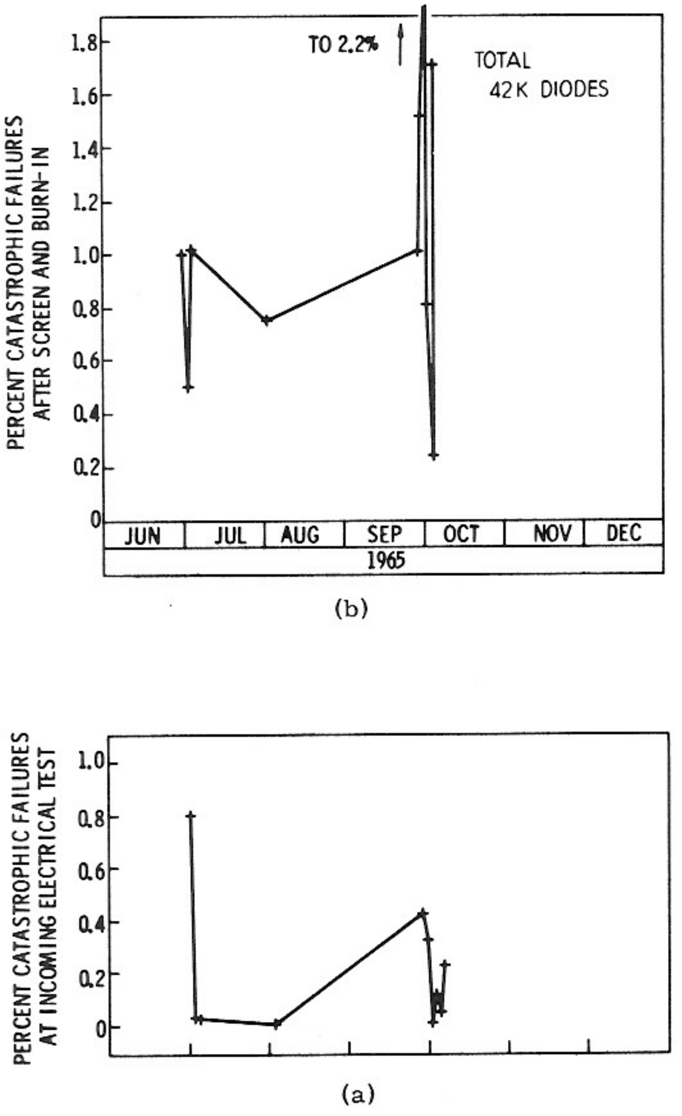

Fig. 7 A Vendor's Diode Performance Through Screen and Burn-In Versus Time.

different vendors. Although the transistors were exposed to a screen and burn-in procedure very similar to Block II transistors, complete screen and burn-in data by vendor is not available However, we can calculate the combined vendor Block I failure rate. Two Block I transistor field failures have occurred:

> One failure was caused by an open bond due to the gold-aluminum eutectic, "purple plague".
>
> One failure was caused by cracking under the bond thus creating a collector-to-emitter short.

Therefore with two failures in 3.5×10^7 element operating hours, Block I transistors exhibit a failure rate of $0.011\%/10^3$ hours at 90% confidence. Note that this is comparable to the single nor gates which were screened but did not pass the flight processing specifications.

Diodes

Block I diodes were exposed to a screen and burn-in procedure as shown in Figure 7. The summarized results of the screen and burn-in procedure are given in Figure 8. Again it is seen that the percent of catastrophic failures during screen and burn-in is high compared to the integrated circuits. The predominant failure modes detected during screen and burn-in were channelling, contact resistance problems and chip cracking. The two field failures which have been generated were due to excessive leakage currents and a cracked chip. The resulting failure rate at 90% confidence is $0.005\% / 10^3$ hours.

In summary, it should be emphasized that most high reliability procedures rely on screening only. However the flight processing specifications impose the additional requirement that entire lots be rejected if the number of failures by failure mode exceeds the reject levels. The rejection of entire lots forces corrective action.

As the data compiled on integrated circuits has shown, more than an order of magnitude difference in failure rates can be expected between lots which pass the flight specifications and lots which are screened but do not pass the flight specifications. Further, it can be concluded that the flight processing specifications have predicted the types of failure which have occurred after screen and burn-in.

One side advantage of the flight processing specs is that through the flight report, better records are kept. Without a manditory documentation procedure. extracting data can be difficult as exemplified by the transistors and diodes.

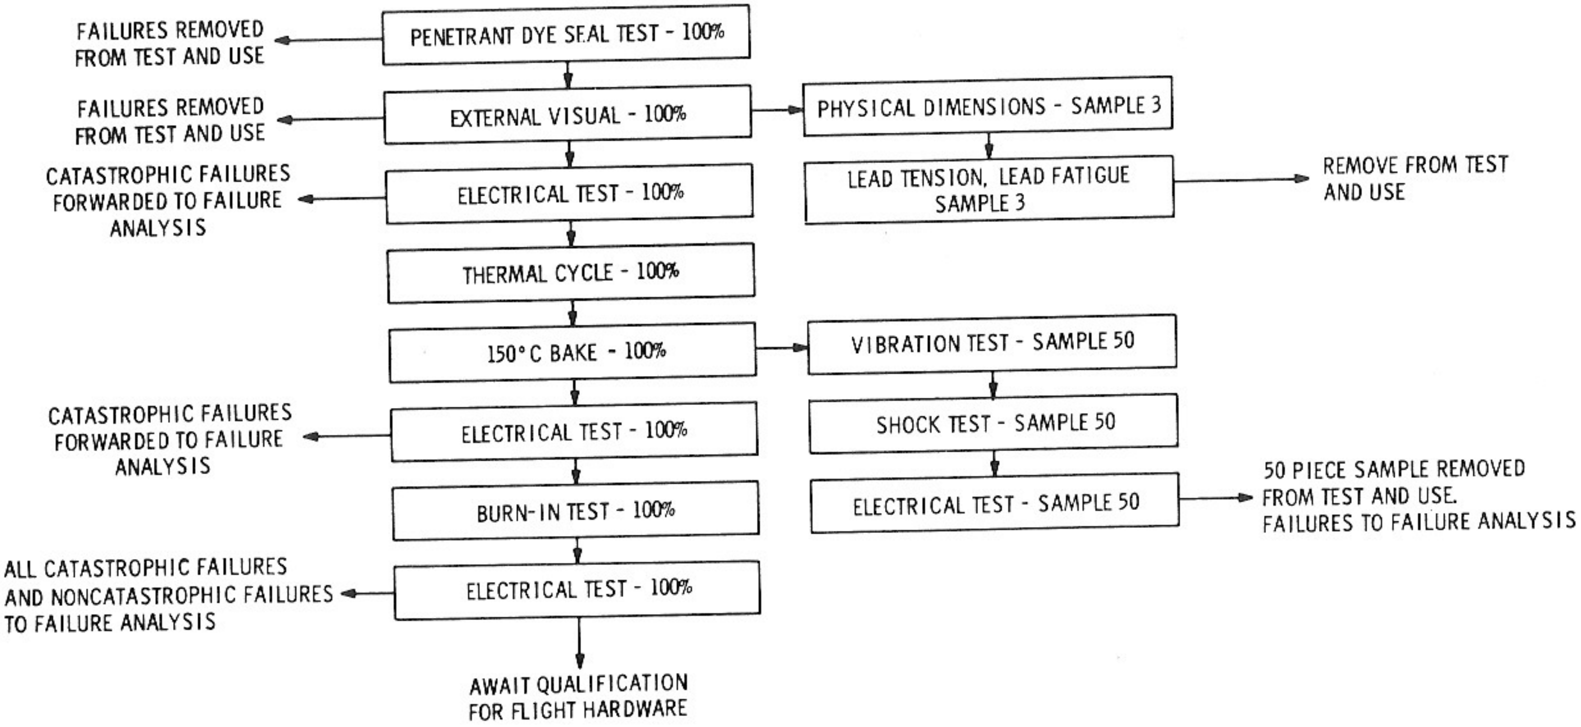

Fig. 8 Screen and Burn-In Flow Diagram for Diodes.

COMPARISION OF PART FAILURE RATES

A comparison of the semiconductor parts failure rates yields some surprising information which would not normally be anticipated. Because of complexity, it is usually assumed that the dual nor gate composed of six transistors, eight resistors, and interconnections should be less reliable than the single nor gate composed of three transistors and four resistors. Carrying the assumption further, discrete transistors should be more reliable than integrated circuits and the diode should be the most reliable component of all. A review of the data of the previous section indicates an opposite trend is occurring. This inverse trend can be explained if one studies the growth of the semiconductor industry. When the procurement of the semiconductor parts for the Apollo Guidance Computer began, the diode and transistor industry had long since reached maturity in high volume, low cost production. At the same time the integrated circuit in the form of the TO-47 single nor gate was being procured from production lines fabricating low-to-medium volume quantities. And for even more contrast, the dual nor gate was supplied by a company building up a capability and directing its production goal toward Apollo requirements and high reliability in general.

There is a decided difference in a company which meets the reliability requirements of a general commercial market, and one which produces a high reliability device exclusively for the military and space markets. The transistors and diodes represent a device which was produced by a general production line but the procurement requirements had little effect on the basic fabrication techniques. Traceability lot control, even in some cases special inspections were instituted but these added requirements did not essentially affect the device fabrication. When general production greatly exceeds the volume of the high reliability military and space market, a separate group, usually from Quality Control, institutes the special non-standardized high reliability requirements. This group complies to the traceability requirements, and performs the screening and test but does not have an active part in the actual device fabrication. This is not as effective as an organization whose sole purpose is to produce the very best device humanly possible and where requirements are directly executed by those fabricating the device. It is interesting to note that, with the same requirement imposed, the transistors which exhibited higher failure rates cost more than the dual integrated circuits.

In defense of the manufacturer these obstacles can only be overcome if the systems requiring high reliability are designed with standardization of parts such that large procurements are possible. The manufacturer can then set up independent high reliability lines or organize his production such that special care in producing the parts is profitable.

PRESENT RELIABILITY PROCEDURES

The previous sections have shown that high reliability is attainable if th vendor's approach is to build the reliability in. At this point,we should critically review what the current popular techniques for achieving reliability really accomplish and if the price paid is justified.

It is not the intent of this section to claim that any of the described technique are not useful, but that a judicious evaluation must prove that the procedures ar effective in attaining the defined goals.

Screening

A screen and burn-in procedure is a generally accepted method for decreasing part failure rates. The effectiveness of the screening depends upon the basic underlying assumption that the failure modes must exhibit a decreasing failure rate with time under a given stress condition. Therefore the stress must just trigger these infant mortalities without shortening the life of the remaining population. To do this, failure modes and mechanisms must be known.

If screening is the only method of achieving high reliability, a few problem become immediately obvious:

1. All failure modes do not have a decreasing failure rate with time, as for example, failures due to corrosion of metallization, contact resistance problems and loose conducting particles[1,2].

2. Few, if any, failure modes can be completely triggered by any one condition and stress level. If one increases the stress level to attempt to increase the probability of removing all failures, then one runs the risk of weakening the devices and introducing new failure modes. An example of this occur in the process of screening gold ball bond failures. Indications to date are that permanent damage due to fatigue has been introduced at a given acceleration level exceeding 20,000 G's[4].

3. The attempt to screen out failures is complicated by the fact that differen failure modes have different stress dependencies. Unpowered baking is generally used as an integrated circuit screen for unstable surfaces and gross metallization corrosion problems. However, if 200°C is exceeded the "purple plague" is introduced and contact resistance at oxide cuts is accelerated. In this way the life of the IC is shortened without screening all failures. On the other hand, power operating acts as a screen for

unstable surfaces, thermal resistance and contact problems, but the test may weaken the integrated circuit due to noncatastrophic oxide breakdown. Further, under a power operating stress, the integrated circuit's useful life could be shortened due to acceleration of aluminum corrosion or "disappearing aluminum"[2,3] at oxide steps and scratches.

The fact must be faced that failure modes which are not screenable exist. Customers will learn that attempts to "test-in" high reliability (i.e., $0.001\%/10^3$ hours) through screening alone is ineffective and extremely expensive.

Step Stressing

Another approach for reliability prediction, referred to as step stressing, is based on the assumption that the failure mode under stress test is a continuous log-normal function of time and temperature. Although failure modes might exist which exhibit such behavior, there are a host more which don't. Since the total failure rate is the sum of all the failures generated by each failure mode, step stressing which does not trigger all failure modes becomes far too inaccurate a prediction tool for high reliability parts. Note that all the integrated circuit field failures were not generated by thermal stressing. The failure due to disappearing aluminum at the oxide step was excited by local heating caused by excessive current density where local temperatures were generated up to the melting point of aluminum. The particle failures failed intermittently. Infrequent conducting particle failures are triggered only by mechanical, not thermal, stressing. The underbonding and lead wire failures are also triggered by mechanical stressing.

Since step stressing is a sample test, another basic assumption underlying the effectivity of the procedure is that the sample must be truly representative of the entire population. For high reliability prediction the sample size must approach 100 % in order to detect all the nonpredominant but contributing failure modes. Because step stressing cannot trigger each and every failure mode in proper proportion it is inadequate for failure rate prediction in the high reliability area. Finally. in any step stress procedure, extreme care must be taken not to trigger failure modes which will never occur at use conditions[5].

Fixing of Processes

Another approach used frequently in attempts to achieve high reliability is the freezing of a process by the customer after a level of development has been reached. After critical processes have been negotiated, contracts are drawn to the mutual agreement of the customer and vendor that the agreed-upon critical processes will not change. In a previous section, inadvertent changes in a dynamic "state of the art" process have been discussed. When a process is not completely understood nor thoroughly documented such that inadvertent changes occur, it becomes evident that attempts to legislate reliability through fixed processing documents are by themselves ineffective. The general techniques presently employed for the sake of increasing reliability are a rash of requirements and controls placed on the vendor. These reliability requirements create a great increase in device cost but do not necessarily increase the reliability commensurate with the cost increase. These requirements have actually impeded corrective action which was known to be necessary.

Sample Testing

Group B and sample testing in general, can only be a gross indicator of reliability. This indicator is adequate for many commercial applications, but has never yielded significant information for the semiconductor parts discussed in this report. To become an effective indicator the sample sizes must increase and reject levels decrease to a point where the cost becomes more expensive than one-hundred-percent testing.

SUMMARY

All the techniques discussed in the previous section can contribute to improvements but in themselves do not possess the sensitivity to predict high reliability. If the cogent features of the above techniques are combined, utilizing knowledge of failure modes and contributing causes, then there is a possibility of increasing the measurement sensitivity. The flight processing specifications[1] were developed incorporating features of the described techniques and the presented data substantiates that they possess the required sensitivity.

Two more prerequisites remain before a high reliability product can be delivered. One of these is the selection of responsible, capable vendors who will:

1. Be motivated to produce and strive to maintain a quality item. This is not as easy as it sounds. The basic problem with a large manufacturer is motivation. He is usually unwilling to alter his assembly process to conform to some non-standardized high reliability procedure. When the production manager is producing millions of devices per month, it is difficult to attract his attention to a 10,000-piece order and much less to a 200-piece order. Although the dollar value of the small, high reliability orders can be quite significant and can influence the marketing department, the pressure of volume delivery is the stronger influence on production.

2. Strive to maintain a consistent process.

3. Be responsive to the study and elimination of problems as they are discovered.

4. Make process changes only after extensive study to determine all the ramifications of the change.

5. Maintain an in-house sponsored reliability program which detects many of the inadvertent changes and long-time-dependent failure modes.

The other prerequisite is the influence created by a judicious, discerning customer who will:

1. Attempt to standardize parts in order to make it worthwhile for the vendor to instrument high reliability production lines.

2. Attempt to buy continuously so that the vendor can maintain a continuous line.

3. Maintain a fair price and not allow competition to jeopardize quality.

4. Develop an efficient monitoring system with rapid feedback to the vendor. Because of changes which occur even with the best of vendors (see Fig. 2), continuous vendor monitoring is required to maintain a consistently well controlled product.

5. Must not request unreasonable or excessive requirements that cannot be instrumented by a good vendor. The requirements must take into consideration the vendor's constraints in producing the parts.

6. Allow a reasonable period of time such that the required procedures can be adequately instrumented by the vendor.

Only by combining a sensitive reliability detection technique, maintaining dependable vendors, being a perceptive and judicious customer, and maintaining a good working relation between customer and vendor, can one hope to build the high reliability into the product.

REFERENCES

1. L. D. Hanley J. Partridge, and E. C. Hall, "The Application of Failure Analysis in Procuring and Screening of Integrated Circuits", Physics of Failure in Electronics, Vol. 4 (June 1966) p. 95.

2. J. Partridge, L. D. Hanley, E. C. Hall, "Progress Report on Attainable Reliability of Integrated Circuits for Systems Application," Microelectronics and Large Systems, 1965, p. 21.

3. W. M. Berger, R. S. Keen, G. L. Schnable, "Reliability Phenomena in Aluminum Metallizations" Physics of Failure in Electronics, Vol. 4 (June 1966), p. 1.

4. Private Communications with Fairchild Semiconductor.

5. J. Partridge, "On the Extrapolation of Accelerated Stress Conditions to Normal Stress Conditions of Germanium Transistors", Physics of Failure in Electronics, Vol. 2, (1964) p. 208.

E-1944

DISTRIBUTION LIST

Internal

M. Adams (MIT/GAEC)
J. Alekshun
R. Battin
P. Bowditch/F. Siraco
G. Cherry
E. Copps
R. Crisp
J. Dahlen
J. B. Feldman
P. Felleman
S. Felix
J. Flanders
J. Gilmore
Eldon Hall
D. Hanley
T. Hemker (MIT/NAA)
D. Hoag
F. Houston
M. Johnston
A. Laats
L. Larson
T. M. Lawton (MIT/MSC)
D. Lickly
L. Martinage
G. Mayo
R. McKern
James Miller
John Miller
J. Nevins
J. Nugent
J. Sabo
J. Partridge (10)
R. Ragan
G. Schmidt
W. Schmidt
N. Sears
J. Shillingford
G. Silver (MIT/KSC)
W. Stameris
W. Tanner
R. Weatherbee
R. Woodbury
W. Wrigley
Apollo Library (2)
MIT/IL Library (6)