P. E. Fox
W. J. Nestork

Design of Logic Circuit Technology for IBM System/370 Models 145 and 155

Abstract: Monolithic System Technology (MST) is a microelectronic circuit family consisting of high-density monolithic circuit chips on ceramic substrates. Using current-switch emitter follower (CSEF) logic, the technology is designed to provide a balanced reflection of cost and performance requirements. One version, MST-2, has been developed for use primarily in the middle range of System/370 – the Models 145 and 155. The typical package usually consists of a single multi-circuit chip on a 16-pin module substrate, and the design embodies a number of recent interconnection and packaging developments that make it well suited to large-scale automated production.

Introduction

The IBM System/370 Models 145 and 155 are in the broad middle range of the data processing market spectrum, where performance and cost are of approximately equal importance. This fact had an important influence on all the design decisions made in the development of the new family of logic technology used in their construction. This family is called Monolithic System Technology (MST) and, as its name implies, is an overall approach to higher circuit density and performance through the use of monolithic circuit chips and a coordinated packaging system, both of which are well adapted to large-scale production. Shown at the right in Fig. 1 is a standard MST logic chip mounted on its one-half-inch-square ceramic module. Also shown, for contrast, is an SLT logic module[1].

MST is a further evolutionary step from its predecessor technologies, SLT and ASLT[2-4]. In addition to the monolithic chip, advances in chip joining and in higher level packaging technologies have been included. These changes have been incorporated in such a way as to take maximum cost advantage of existing automated production facilities.

The general approach has been to use a single monolithic chip on a modified ceramic SLT 16-pin module. The circuit used for general logic is of the current switch emitter follower type with both collector dotting (implicit AND) and emitter follower dotting (implicit OR). A speed-vs-power design point has been chosen to permit cooling with forced room-temperature air.

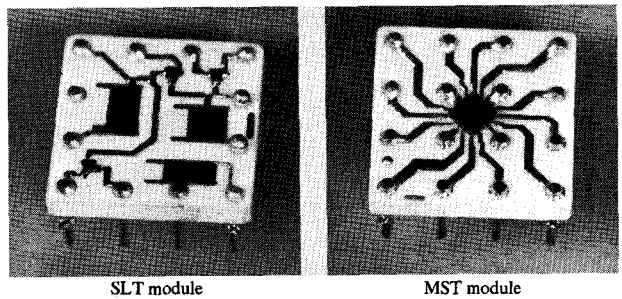

Figure 1 SLT (left) and MST modules.

This paper provides an overview of the technology and attempts to give the reasoning behind the more important design decisions. Processing techniques are reviewed, but their details will be found in the references.

Logic circuit

The principal logic circuits considered for this technology were 1) DTL, which had been used in SLT, 2) TTL, and 3) current switch emitter follower (CSEF). After a weighing of the advantages and disadvantages of each, CSEF (see Fig. 2) was chosen for three primary reasons:

1) *Performance*—The CSEF offers excellent speed and power characteristics. The fact that it can be designed for circuit propagation delays of less than

5 nsec with low power dissipation is important for a system having a high circuit density in an air-cooled environment.

2) *Logic power*—This circuit offers system designers a highly flexible logic function. By the use of the implied AND and implied OR functions, it is possible to have three levels of logic in one circuit. Actual implementation of machine logic has shown that 15 to 30 percent fewer circuits are required in CSEF than in DTL or TTL.

3) *Experience and compatibility*—Current switch circuits have been used by IBM for more than ten years[5], and the CSEF for more than five years[3]. Experience with the ASLT thick-film version of the CSEF in the System/360 Model 91 and associated memories was excellent. In addition, the advantages of this circuit led to its choice for newer high-performance systems and their memories. Thus, signal-level compatibility became an important factor. Also, the decision to use a related technology[6] in high-volume machines such as System/3 led to an excellent situation with respect to a manufacturing base.*

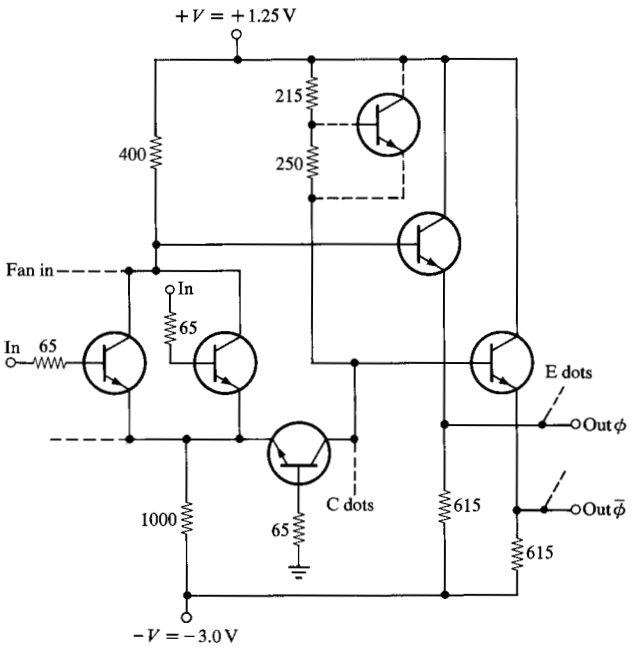

Figure 2 Basic CSEF circuit. (Resistance values in ohms.)

Circuit configuration and characteristics

The circuit shown in schematic form in Fig. 2 is the basic logic circuit of this technology. Its maximum logic capability is listed in Table 1. The implied OR function at the output, called emitter dotting, is obtained by connecting the outputs of other similar circuits as shown in Fig. 2. One common 615-ohm emitter load resistor is used. The implied AND function is obtained by dotting in-phase (ϕ) current switch collectors. Again, a single common load resistor is used. In the collector case, however, it is necessary to add a transistor clamp, shown with broken lines in Fig. 2, to prevent saturation of the grounded-base transistors when two or more collectors conduct simultaneously. It should be noted that these collectors are in separate current switches, and hence each can contribute one unit of current when it conducts. In contrast, the paralleled input transistors collectively pass only one unit of current, no matter how many conduct.

Although not shown in Fig. 2, the out-of-phase (input transistor) collector uses a clamp similar to that of the in-phase device when a two-input EXCLUSIVE OR function is formed from two single-input current switches. This EXCLUSIVE OR has out-of-phase collectors, in-phase collectors, and emitter followers dotted together in pairs.

Instead of the emitter stabilization capacitor of ASLT it was judged to be more efficient, in a monolithic environment, to use base-resistance stabilization. The 65-ohm underpass resistors insure that the input impedance is limited to values such that oscillation is statistically unlikely with the transmission-line configurations permitted by the wiring rules discussed below.

Table 1 Maximum logic capability, MST-2.

Fan-in	4
Fan-out	10
Emitter dots	4
Collector dots	4

The resistor values and power supplies shown in Fig. 2 were determined by the techniques of statistical design previously discussed[7]. Advantage was taken of the tracking of resistors and of transistor V_{BE}'s on the same chip[8]. Also included in the design was the requirement that any output be able to drive either an unterminated line or a 90-ohm line terminated by a 90-ohm resistor to ground. This fact permits any logic circuit to be used as a line driver, greatly simplifying the task of the system designer.

We should mention at this point that the 90-ohm impedance level represented a fundamental design decision. A lower value would have required higher power in the emitter follower. This is because the signal swing in CSEF circuits is not much affected by power level, but is determined primarily by V_{BE} and component tolerances. Because of the finite power gain of the emitter

*The version of the technology used in System/3 is called MST-1, while that used in Models 145 and 155 is called MST-2. Both are made on the same manufacturing line with the same process.

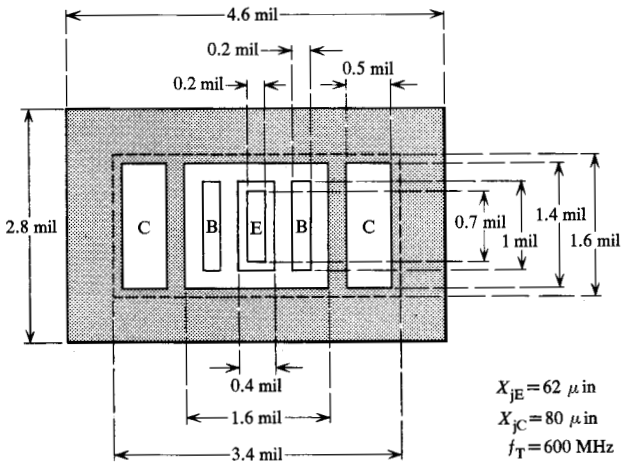

Figure 3 Dimensions of basic transistor.

follower, the current switch power also would have gone up. Together, these increases could have jeopardized the room-temperature air cooling requirement. On the other hand, line impedances higher than 90 ohms are hard to obtain mechanically. Also, a further reduction in power would have led to excessively large collector-load RC time constants. This would have caused increased propagation delays. On the basis of these considerations, 90 ohms was chosen as the best overall compromise.

When account is taken of the effect of emitter dots, the average power of a typical configuration of the circuit of Fig. 2 is about 30 mW. For the unloaded in-phase output, propagation delay (raw-circuit delay) varies from 3.5 to 5.9 nsec, depending on the number of collector dots. Out-of-phase delays are in the lower part of this range and, of course, depend on fan-in. In-the-environment delays are of more interest and significance, however. A study carried out on more than 800 logic stages in the Model 155 showed an average total delay of 7.5 nsec per stage. This was made up of 4.5 nsec for the raw circuit, and 3.0 nsec for the line-loading and wire (including cables).

Semiconductor structure

The silicon structure decided upon begins from a fairly conventional base, to which are added a number of new features. Arsenic n^+ subcollectors are diffused into 20 ohm-cm p-type wafers. After deposition of the 0.09 ohm-cm n-type epitaxial layer, p^+ isolation and p base diffusions are made with boron. Emitters are phosphorus.

All resistors except base underpasses are base diffusions. To obtain the desired stabilization resistance in the underpasses, they are made with subcollectors, over which an emitter diffusion is added at each end for contacts. The basic npn transistor, shown in Fig. 3, is used in all current switches. For the emitter follower and some collector clamps, either two of these transistors in parallel or a larger transistor with two of the same emitters is used. Emitter current in the current switch is 2.5 mA, while total emitter follower up-level current is 10 mA when the output is terminated. The two emitters are used to reduce the current density in the metal emitter stripes. Additional design features to limit current density are enlarged contacts at the positive-potential ends of resistors, and redundant pads for both $+V$ and $-V$ power supply to the chip.

The concern about current density in the metal was, of course, due to the electromigration phenomenon in thin metal films[9]. In the final version of the MST process, the metal used is copper-aluminum[10]. The total effect of these precautions has resulted in the reduction to a negligible level of the metal-migration contribution to failure rate.

The MST chip uses an improved glass encapsulating scheme that is well adapted to monolithic circuit processing. Over the final layer of thermally grown oxide, there is placed a layer of rf-sputtered SiO_2[11]. Contact hole opening, Cu-Al evaporation, and subtractive etching follow. After metal sintering, a final overlay of rf-sputtered SiO_2 is applied to protect the chip lands. Only one layer of metal is used. Holes are then opened in the overlay to accept the evaporated Cr-Cu-Au plus Sn-Pb pads which form part of the chip-to-module contact system (see below). The result is a passivated and sealed chip that needs no further moisture protection.

Packaging

The first level of MST packaging is the ceramic module shown in Fig. 1. The chip-to-module contact system used in MST is better adapted to monolithic circuit chips with their large number of contacts than is the metal ball system used in SLT. Known as controlled-collapse reflow chip joining[12], the new process depends upon a controlled volume of tin-lead solder in a controlled area to make electrical and mechanical contact with the module lands. The contribution to this volume from the chip pad is controlled in the pad evaporation process. The means of control for the module land contribution can be seen in Fig. 4, which shows the ceramic substrate just before the chip is applied. Note the glass dams, which are screened over the chip-contacting fingers near their ends. Because the dams are not wettable, they limit the amount of solder from the module available to each contact to the amount picked up by the finger tip during the module Sn-Pb dip operation. During the solder reflow step, after the chip has been positioned, the surface tension of the molten solder on all the pads pulls the chip into final alignment with the module fingers. It also determines the chip-to-module stand-off spacing with excellent accuracy.

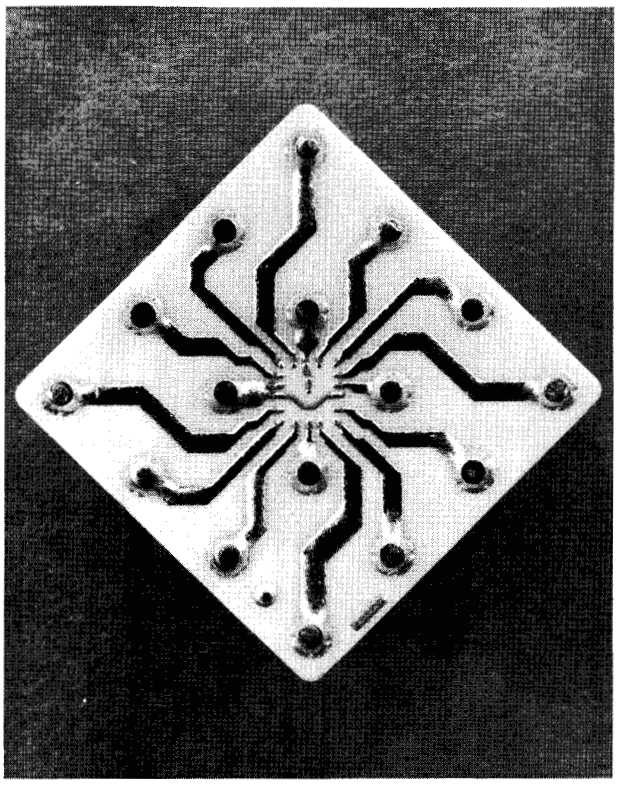

Figure 4 Substrate ready for chip joining. (Note glass dams near ends of land "fingers.")

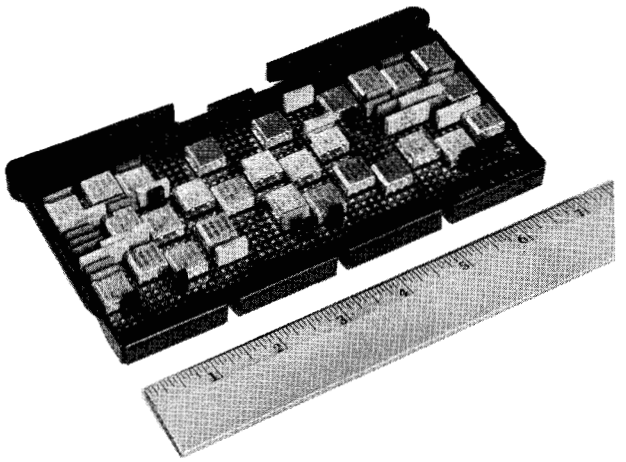

Figure 5 MST-2 4-wide card. Plastic levers along edge opposite sockets are for extraction. Small components among modules are resistors and by-pass capacitors.

Figure 6 MST-2 board with three 4-wide cards. Metal superstructure is for card retention and extraction.

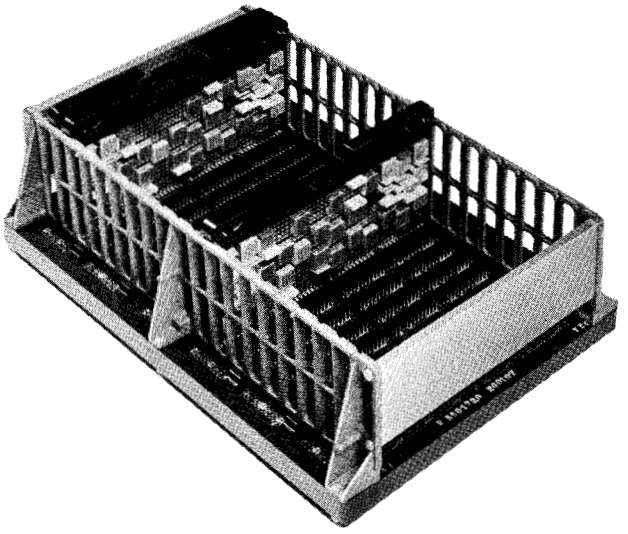

Further study of Fig. 4 shows the redundant power supply contacts mentioned above. There are three for $+V$ and three for $-V$. The total number of pads is thus 20. The 4×4 array of module pins on a 0.125-inch grid provides three power supply pins and 13 input/output pins. As in SLT, the completed substrate assemblies receive a protective aluminum can and silicone encapsulants to increase the corrosion resistance of the module lands and chip-to-module contacts.

The second level of packaging is an epoxy-glass card, Fig. 5, with two internal voltage planes (one shared between ground and $-V$) and two surface signal planes. Plated-through holes are on a 0.125-inch grid to match the module pins. Cards are available in one, two, or four units of width of 1-3/4 inches. Height of each size is a uniform 2-3/4 inch plus contact tabs. There are 24 tabs, 12 on each side of the card in each unit width, spaced 0.125 inch O.C. Two of these are used for $+V$, one for ground, and one for $-V$, leaving 20 for logic I/O.

The 4-wide card has its power planes brought out to match the standard module power pins at 36 module sites. The actual number of modules on a particular card depends, of course, on the logic being implemented and the ratio of spare positions desired. The narrower cards have proportionately fewer module sites. The two copper signal planes are subtractively etched into transmission-line interconnections with a characteristic impedance of approximately 90 ohms. The higher impedance required for a dual line-impedance scheme[3] at the basic 90-ohm level would have resulted in excessive dielectric thickness or overly narrow conductors.

The third level package, the back panel or board, Fig. 6, has multiple voltage planes and either two or four signal planes. The choice for a given application is determined principally by the circuit density desired.

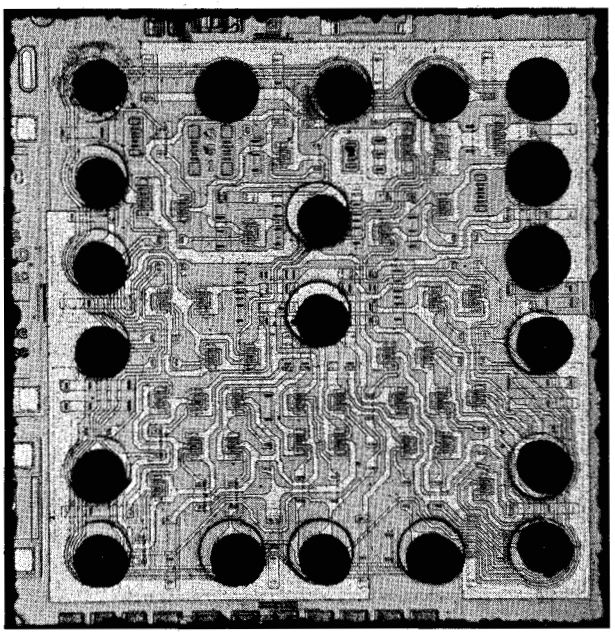

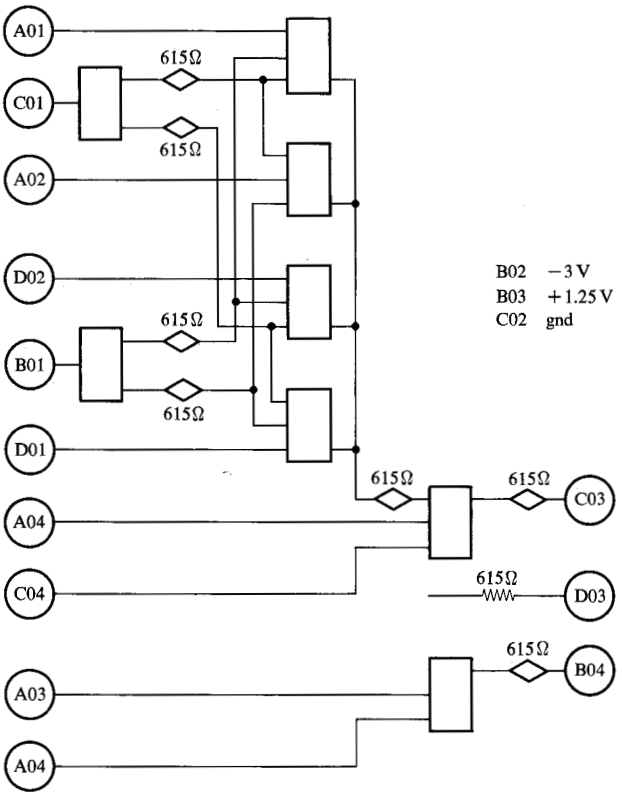

Figure 7 Photograph and logic block diagram of functional chip. In the block diagram each rectangle is a current switch with inputs at left, outputs at right, in-phase output at bottom, out-of-phase at top. The diamonds are emitter followers. Resistance values are for load resistors, with all 615-ohm loads connected to −3V.

Module set

For this technology, it was decided that a limited set of standard module part numbers was the best overall approach. Two implications of this decision were evident. First, a well-chosen set of unit-logic part numbers would be required for use in random-logic areas of systems. Second, to make up for the relatively low density of the unit logic, we would need some relatively high density functional part numbers that might be called medium-scale integration (MSI). The resulting set of standard MST-2 logic modules contains 25 part numbers. The density range is 2 to 16 circuits per module. Some of these part numbers were wired without emitter load resistors so that their outputs can be dotted into other circuits. In the Models 145 and 155, the weighted average density is 5.5 to 6.5 circuits per module. Figure 7 contains a photograph of the chip for a module of above-average density. Also shown is a block diagram of its logic.

In the latest production version, most module part numbers have a single chip and use the simple space transformer substrate of Figs. 1 and 4. All of the corresponding chip part numbers are made with a common diffusion and contact-hole pattern. Part number differences are therefore largely determined by the metal subtractive etch pattern. A few of the higher density part numbers use two chips from the same master chip.

In addition to the standard logic module set, there are about 15 modules available as a logic support set. These part numbers include such items as single-shot multivibrators, indicator drivers, and circuits to convert to the standard system I/O interface. They employ a mixture of monolithic and hybrid technology. Also available with compatible form factor is a collection of by-pass capacitors and resistors. The most used of these is the four-pin module carrying three 90-ohm resistors used for line termination. There is also a similar module with three emitter load resistors. The latter is used only when an emitter load is not available on any of the modules being dotted together.

Wiring rules and delay prediction

Restrictions on the configuration of loads on transmission line circuit interconnections are required for two reasons:

1) to improve the predictability of delay, and
2) to prevent the formation of spurious signal pulses on a line by excessively-large reflections.

In this technology, the wiring-rule problem is complicated by the presence of both terminated and unterminated nets.* In both types of nets, the primary limitations

*The term "net" is used to refer to the assemblage of lines, stubs and loads driven by a circuit output.

on the attainment of the two objectives listed above are 1) the loss of wireability with a consequent loss of circuit density, and 2) the increase in delay due to added wire. For the unterminated case, the most effective wiring restriction is the total length of the net. In most cases, 12 inches is allowed. For the most disadvantageous spacing of loads, this must be reduced to 8 inches.

When it is not possible to arrange a net to meet the unterminated requirements, it is necessary to terminate the net. The rules for terminated nets permit a cluster, defined as a concentration of loads. Although there is no restriction on the number of loads in a cluster, it is necessary to restrict its length. Any remaining loads that will not fit into a cluster must be distributed along the line beyond the cluster. In the "distributed section," there is a minimum spacing between loads determined by the allowable reduction in line impedance caused by the capacitance of the loads. Thus, at the expense of wiring space for "forced wire," any configuration of loads can be interconnected in the distributed section of a net.

The cycle times of the Models 145 and 155 represent a very close match with the performance capability of the MST-2 circuits. It was therefore necessary to provide to the logic designer a practical means for computing delay and tolerance on delay for any chain of logic he might want to assemble. With such a tool, he would not have to waste circuit capability to cover excess uncertainty in what his actual delays would be.

The calculation of logic chain delay was handled by the addition of a module (raw-circuit) component, and a line-loading component. The former is tabulated by module part number and by logic path within the module. The latter is obtained from a heuristically-derived formula using line length, number and location of loads, and type of net.

The problem of tolerance on delay was attacked by first considering the case of a relatively long logic chain. It was reasoned that, in such a chain, there would normally be enough randomness that a statistical approach could be used. It was then assumed that the actual delays of a large number of these chains, when normalized to their predicted delays, would result in a normal statistical distribution. This was later proven experimentally to be true.

The next step was to generate a formula for computing the standard deviation of delay for any logic chain. The general approach was to make a root-sum-square addition of the various components of tolerance with heuristically-determined weighting factors. The resulting expression was too complicated for manual use, but was well adapted for use in a computer program that had been written for computation of delay. With that addition, the program produced the nominal, maximum, and minimum delays (to the instructed number of standard deviations)

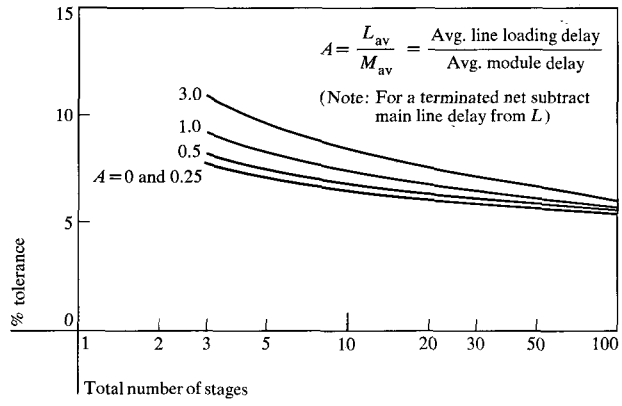

Figure 8 One-sigma delay tolerance vs number of stages in the logic chain.

when it was fed a fairly complete description of the physical implementation of the chain. When one considers the fact that a typical Model 155 contains over 10,000 nets, the utility of such a program can be better appreciated.

To give the logic designer an approximate method for manual evaluation of a chain, as well as to provide an appreciation of the relative importance of the components of tolerance, the curves of Fig. 8 were produced. The curves start with the three-stage chain because it was determined that the statistical approach cannot be applied to one and two-stage chains. For these cases, special tolerances are assigned that range from ±45% to as much as a factor of two for three standard deviations. It was also found necessary to set up a list of criteria which a long chain must meet if it is to qualify for the statistical treatment. For those few chains that do not qualify, the larger two-stage tolerance must be applied.

Conclusions

Monolithic System Technology is a coordinated over-all approach to the problem of producing monolithic integrated electronic components and applying them to systems for the broad mid-range of the data processing product line. Engineering trade-offs have resulted in a desired balance of cost against performance. Experience to date has shown that MST is extremely reliable and can be produced with good yields.

Acknowledgments

It is impossible, in an undertaking of the magnitude of MST, to list all of the many contributors. Mention should certainly be made, however, of J. L. Walsh and E. J. Rymaszewski, who were responsible, at different times, for the over-all direction of the program.

References

1. E. M. Davis, W. E. Harding, R. S. Schwartz and J. J. Corning, "Solid Logic Technology: Versatile, High-Performance Microelectronics," *IBM J. Res. Develop.* **8,** 102 (1964).

2. J. L. Langdon and E. J. VanDerveer, "Design of a High-Speed Transistor for the ASLT Current Switch," *IBM J. Res. Develop.* **11,** 69 (1967).
3. R. F. Sechler, A. R. Strube and J. R. Turnbull, "ASLT Circuit Design," *IBM J. Res. Develop.* **11,** 74 (1967).
4. R. H. F. Lloyd, "ASLT: An Extension of Hybrid Miniaturization Techniques," *IBM J. Res. Develop.* **11,** 86 (1967).
5. H. S. Yourke, "Millimicrosecond Current Switching Logic Circuits," *IRE Trans. Circuit Theory* **CT-4,** 236 (1957).
6. O. Bilous and E. J. Rymaszewski, "Medium Density Monolithic Logic Technology," presented at the *1969 International Electronic Circuit Packaging Symposium* (Wescon), Aug. 18–22, San Francisco.
7. D. H. Chung and J. A. Palmieri, "Design of ACP Resistor-Coupled Switching Circuits," *IBM J. Res. Develop.* **7,** 190 (1963).
8. O. Bilous, I. Feinberg and J. L. Langdon, "Design of Monolithic Circuit Chips," *IBM J. Res. Develop.* **10,** 370 (1966).
9. W. E. Mutter, "Electromigration in Metal Films," presented at the *1967 Electrochemical Society Spring Meeting*, May 7–12, Dallas. Abstract No. 62. p. 96.
10. F. M. d'Heurle, "The Effect of Copper Additions on Electromigration in Aluminum Thin Films," *Metallurgical Trans.* **2,** 683 (1971).
11. W. A. Pliskin, P. D. Davidse, H. Lehman and L. I. Maissel, "Properties of Insulating Thin Films Deposited by RF Sputtering," *IBM J. Res. Develop.* **11,** 461 (1967).
12. L. F. Miller, "Controlled Collapse Reflow Chip Joining," *IBM J. Res. Develop.* **13,** 239 (1969).

Received April 26, 1971

The authors are located at the IBM Components Division Laboratory, East Fishkill, New York 12533.