

Bipolar Circuit Design for a 5000-Circuit VLSI Gate Array

Ever-increasing gate-array chip circuit count and size pose challenges to the circuit designer, who must optimize the cell layout and power dissipation of logic circuits. This paper describes the design of a bipolar gate array for a 5000-circuit microprocessor. The physical design data obtained after completion of automatic placement and wiring are presented. The distribution of projected circuit delay for the 4437 nets is then calculated and this information is used to determine the selective increase in circuit power to reduce wide spreads in turn-off and turn-on delays. It is shown that this technique improves power/performance and has implications for future VLSI designs. It was also found that voltage drop violations were manageable even with the long net lengths occurring in this large VLSI gate array. Several commonly used bipolar circuit types are compared on the basis of capacitance drive capability at low power.

Introduction

Gate array designs in IBM started with a 100-circuit gate array in the early 1970s and continued through gate arrays of 700 [1], 1300 (STL) [2], and 1500 circuits [3] in the late 1970s. This progression in chip circuit count is plotted in Fig. 1. (The earlier MST [4, 5] is also shown.) Other gate arrays are presented in [6]. This paper describes the circuit design of a VLSI chip which extends this progression to 5000 logic gates. It details technology selection, chip design, wiring data, and circuit performance for a 5000-circuit experimental microprocessor [7]. The microprocessor executes the full System/370 instruction set at 200 kips (kilo-instructions per second). This requires 5000 high-performance circuits on one chip, thereby encompassing the equivalent data flow of the System/370 Model 138 [8]. The microprocessor is designed on a gate array using design automation techniques which are an extension of those applied previously to less dense gate arrays [9].

The discussion of the technology selection phase of the project includes the key decisions on circuit type and cell design. Some unique aspects of the chip image are also described. The measured and calculated performance of the low-power T²L (Transistor-Transistor Logic) circuit is quantified. The wiring data (primarily net lengths) obtained after completion of automatic placement and wiring are presented as a statistical distribution for the 4437

nets. With the use of this length histogram and the circuit delay sensitivity to net length, a delay histogram is calculated. Also presented are the results of the voltage drop analysis, which is critical for large chips with small geometries using T²L circuits.

The greater density achieved in this chip illustrates the consequences of the need to reduce circuit power dissipation by operating at lower power supply voltages. The result is a greater sensitivity to load capacitance. This paper demonstrates a technique of selective increases in power based on the histograms of net lengths and the microprocessor interconnections which reduces the spread on the histogram of circuit delays and minimizes the skew between turn-off and turn-on delay. Comparisons of delay between various circuit types and recommendations for future circuits to be used for VLSI gate arrays are made on the basis of the delay histogram calculations.

The organization of the paper follows the development of the chip. After a brief description of the microprocessor function, the technology selection phase, which consumed approximately one man-year of effort over two months, is discussed. Following this are sections on the gate array chip image (which was designed for the use of automatic placement and wiring programs), the circuit performance, and power dissipation. The main

Copyright 1981 by International Business Machines Corporation. Copying is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the *Journal* reference and IBM copyright notice are included on the first page. The title and abstract may be used without further permission in computer-based and other information-service systems. Permission to republish other excerpts should be obtained from the Editor.

section of the paper deals with the physical design data obtained from the microprocessor design of the gate array. Quantitative effects of these data on chip performance are given, and the paper concludes with recommendations for future designs.

Microprocessor function

A brief discussion of the microprocessor function is provided as background information before the gate array design is described. The data flow and function of the microprocessor are described in detail in [8]. The microprocessor is capable of executing the full IBM System/370 instruction set at an average sustained rate of 200 kips. This is based on an average System/370 instruction requiring 50 machine cycles with each cycle requiring 100 ns.

This chip only contains the data flow portion since the control read only store (ROS), general-purpose registers, and, of course, memory are all located off this chip. The System/370 data flow chip (refer to Fig. 8, given later) contains an 8-bit Arithmetic Logical Unit (ALU), a 24-bit incrementer/decrementer, control registers and decoding circuits, and various working registers to store memory data, instructions, memory address, central processing unit status, interrupt condition, and ROS address. Among the signals entering and leaving the chip are a 54-bit ROS control word, a 16-bit ROS address, a 16-bit bidirectional memory data buss, and a 24-bit memory address. Also included as inputs are clock signals, status control, and various machine control functions such as stop, reset, memory, write, etc. Several extra busses such as the "Z" and "A" (in Fig. 8) were brought off chip to improve the test pattern coverage. Also included on the chip are circuits to perform the relocate function. The total input and output (I/O) signal count is 200. There are no PLAs or ROS contained on this chip. On-chip storage or registers are implemented with standard logic circuits configured as set-reset latches.

Technology selection

The first consideration in the design of this microprocessor chip was the selection of the type of semiconductor process. As stated in the previous section, the machine cycle objective was 100 ns. This objective and the logic path length of 23 stages required a worst-case internal circuit delay including loading (such as fan-out, dotting, and wire capacitance) of 4 ns or less. Because of this delay requirement, a bipolar process (as opposed to FET) was selected for the chip.

The next major selection was the chip design technique, namely custom design or gate array. In a gate array, the diffusions are fixed and only the metal layers can be changed from one part number to another. On the one




Figure 1 Progression of IBM gate array chips.

hand, custom design has the potential for greater density due to its freedom to vary the diffusions and thereby optimize the area for each function. Gate arrays offer the advantages of using the extensive and broad computer-aided design software base called the engineering design system [10]. Reference [10] enumerates the advantages of using gate arrays. For this chip, the most important advantage is the automation of the placement, wiring, and complete checking. This checking includes process ground rules such as metal spacings, logical to physical validation, and electrical rule violations. Since this data flow chip contains approximately 5000 circuits and 10 000 connections, automation and checking are essential because of the tremendous volume of data involved. The decision was made to design the microprocessor initially as a gate array and, if custom design techniques could reduce the chip size, to phase it in for a second pass as a cost reduction.

The next major decision was which logic circuit type could achieve the desired performance. SCS (Schottky Current Switch) and MTL (Merged Transistor Logic) were considered but rejected due to high cell area and low performance, respectively. Two circuit types could satisfy both the worst-case performance of 4 ns per stage and the required density of 5000 wired circuits on a 7-mm chip. These two circuits were T^2L and STL (Schottky Transistor Logic) [2]. STL is similar to I^2L (Integrated Injection Logic) but faster. The version of STL considered was the clamped one shown in Fig. 10 of Ref. [2].

Due to the strong interrelationship between wiring channels available in each direction of the individual cell or circuit, and the chip size necessary to wire 5000 circuits, many cell configurations had to be evaluated for each circuit type. This is especially true of STL, which uses "remote" Schottky barrier diodes (SBD) to save

Table 1 T²L versus STL (clamped version): Delay comparison based on computer simulation.

Loading conditions	Condition 1	Condition 2	Condition 3
Fan-out	1	3	1
Collector dotting	1	2	5
Total net wire capacitance (pF)	0.1	0.375	2.0
Delay simulation results			
(1) T ² L average delay (ns)	1.91	2.04	4.79
For STL the wire capacitance is located at the base or collector depending on the location of the remote SBDs			
Assumed portion of net wire capacitance (pF):			
At collector			1 0
At base			1 2
(2) STL average delay (ns (layout with 6 SBDs in the collector bed))	1.73	1.85	4.58 5.13
(3) STL average delay (ns (layout with 10 SBDs in the collector bed))	2.25	2.25	5.6 5.9

Table 2 Cell and chip size comparison for all four cases (6K circuits assumed per chip). The T²L cell is the cell which was adopted for the gate array.

Circuit type	Cell size (μm)	Cell wiring channels (Vert. × Hor.)	No. of cells per chip	Total wiring channels per chip	Chip size (mm)
T ² L	75 × 75	10 × 5	7640	1280	7.16
STL (6 SBDs)	117 × 97	10 × 4	9128	1212	7.16
STL (10 SBDs)	—	10 × 3.5	12,000	1304	7.5

wiring space. Also, STL circuit components are highly integrated with the cell wiring. A total of 35 cell layouts for STL and T²L were analyzed and these were reduced to the most promising four, two of each circuit type. Three of these cells were modeled in the circuit simulation program, ASTAP [11] (the two T²L cells have identical performance), and the performance was analyzed for various combinations of fan-out, capacitance at the collector and base for STL, and collector dotting. The results of this analysis are presented in Tables 1 and 2. The chip size required to wire 5000 circuits was compared for

the four cells, and yield was calculated on the basis of chip size and critical device areas. The overall results were as follows:

1. Performance for STL was within 10% of the T²L performance depending on the cell layout used for STL and the location of the remote SBD (refer to Table 1 for a summary of the data).
2. Chip size was comparable (within 10%), provided some of the new wiring concepts developed for STL [2] were also applied to T²L. Table 2 compares the cell and chip size and the wiring channels available for the three most promising cells corresponding to Table 1.
3. The estimated yield was 30% better for STL but the greater logic power of T²L would reduce the yield advantage of STL somewhat.
4. T²L uses a slightly simpler process, and the logical representation and physical design system are simpler.
5. T²L can utilize metal ground rule improvements more easily since intercircuit wiring is more isolated from the device (transistor) layout than in STL.

On the basis of these factors, T²L was selected as the circuit type, although the decision was a close one. A cell layout was designed and optimized based on discussions with the engineering design departments responsible for automated placement and wiring as described in [9] and [12]. The technology comparison and cell design phase took approximately two months due to the tradeoffs between cell design and circuit performance, chip image, and capability to wire 5000 circuits on a fixed-size chip.

Process and chip description

The semiconductor process is identical to the one developed for the bipolar random logic chips contained in the IBM 4300 computer series [1]. Its main features are a 2-μm-thick epitaxial layer, recessed oxide isolation, and three levels of metal. However, for this chip, the horizontal ground rules are based on smaller image sizes.

The gate array (masterslice) contains basic cells in a matrix of 96 rows by 92 columns. Figure 2 is an illustration of the basic chip image. Each cell can contain one circuit and occupies an area 75 μm by 75 μm. Note, in Fig. 2, that the off-chip drivers and receivers are interspersed with the internal circuits. They are laid out to allow wiring to pass through the row or column in the perpendicular direction. The drivers occupy twelve columns and the receivers four rows. There are a total of 7640 internal cells, which will be only approximately 65% occupied by circuits. The other 35% of the internal cells are unoccupied in order to provide additional wiring channel capacity.




Figure 4 Second-level-metal wiring for microprocessor.




Figure 5 Photograph of microprocessor chip after terminal metallurgy.

This circuit only requires a low power supply voltage of 1.7 V, and the off-chip receivers and drivers require a higher supply voltage of 3.4 V to generate larger voltage swings. The total nominal chip power dissipation is 2.3 W, of which 1.6 W is dissipated by the internal circuits.

Table 4 Wiring statistics for microprocessor.

Chip size:	7 mm
No. of internal circuits wired:	4705
No. of off-chip drivers:	96
No. of receivers:	122
No. of signal pads used:	200
Total first-metal wiring length:	2.40 m
Total second-metal wiring length:	3.40 m
Total no. of wiring vias:	33,516
No. of nets:	4437
No. of connections:	10,605

Part number wiring statistics

The experimental microprocessor was automatically placed and wired on the previously described gate array with the use of extensions to programs applied to other gate arrays [3, 9, 12]. Figure 4 shows the checkpoint of the data used to generate the mask for the second-level-metal layer. Figure 5 is a photograph of the chip showing the pads.

The wiring statistics from this 5000-circuit microprocessor are listed in Table 4, with the total interconnection length (including stubs) being equal to 5.8 m. This can be compared to 2.2 m of wiring for the 704-circuit gate array used in the IBM System/38 [1, 13]. The average net length is then equal to 5800/4437, or 1.3 mm. The average net length does not convey enough information to effectively design the circuits for optimum power/speed. The distribution of net lengths is important since lengths significantly greater than the average will cause performance degradation by causing a significant increase in turn-off (T_{off}) delay.

The automatic placement and wiring programs were used to compile the total net length including stubs for each of the 4437 nets, and Fig. 6 is a histogram of these data. Note the long tail of the distribution extending out beyond 1.0 chip edge in length (7 mm). Almost 6% of the nets have a length greater than 7 mm. The net length can be translated into total net capacitance, and then the distribution of circuit delay can be calculated based on the sensitivity of the circuit delay to load capacitance. These translations are performed using the formula for the transformation of a single random variable.

Given $y = F(x)$ and the probability density function (*pdf*) of x equal to $P_1(x)$, the equation for transformation is [14]

$$P_2(y) = P_1(x) \frac{dx}{dy} = \text{pdf of } y.$$




Figure 6 Histogram of total net wiring lengths for microprocessor.

For our particular case,

$C = \text{capacitance} = \text{const.} \times \text{length};$

$C = k \times l, k \text{ is positive.}$

Then,

$$P_2(C) = P_1\left(\frac{C}{k}\right) \times \frac{1}{k}.$$

Therefore, the *pdf* for C is identical in shape to the *pdf* of length (P_1) except the axis is transformed by the constant k . The magnitude is multiplied by $1/k$ but the plot in Fig. 6 uses a magnitude expressed in percentage and so the subsequent *pdfs* for capacitance and delay will do the same. This means that $\int p(x)dx \neq 1$.

Similarly, the delay *pdf* is transformed from the capacitance *pdf* using a constant which is the delay sensitivity. The sensitivity of the T_{off} delay to load capacitance is 1.5 ns per pF for $R_L = 8 \text{ k}\Omega$ and 0.75 ns per pF for $R_L = 4 \text{ k}\Omega$ for one active load switching on a net. These sensitivities were obtained from computer simulations using ASTAP and can also be calculated using an exponential charging equation, since the output transistor is in the off state. More active loads will reduce the T_{off} delay since active loads supply charging current during part of the up-going transition.

Histogram of circuit delays based on net lengths

With wire capacitances of 0.63 pF per mm and 0.2 pF per mm for first- and second-metal wiring, respectively, a ca-




Figure 7 Histogram of wiring capacitances for microprocessor.

pacitance histogram can be calculated from the net length histograms. It is plotted in Fig. 7. Of course, the step height of this histogram is the same and it also has a long tail with approximately 6% of the nets or 265 of them having a total wiring capacitance greater than 3.0 pF. This high capacitance will slow down those 265 nets and degrade the performance of the microprocessor.

After the length and capacitance histograms have been obtained, the question arises as to which sections of the microprocessor function contain the high-capacitance nets and may therefore suffer increased delay. Figure 8 contains a very simplified diagram of the chip indicating the location of the nets above 5 pF. Reference [8] contains a much more detailed functional diagram and description of the operation of the microprocessor. Table 5 groups the nets with C greater than 5 pF by type and from-to designation.

Several general characteristics can be observed from the table and diagram. First, the high-capacitance nets occur as expected between functions, that is, between pages of logic. The microprocessor logic occupies approximately 50 pages. Only one net above 5 pF occurred within a function or page which averages approximately




Figure 8 Microprocessor block diagram showing location of high-capacitance nets. Circled numbers indicate number of nets above 5 pF.




Figure 9 Turn-off delay histogram.

Table 5 Location and type of high-capacitance nets.

No. of nets	Function type			Location	
	Data (buss)	Control	Timing	From	To
16	✓ (Z)	—	—	ALU	Working registers
5	—	✓	—	Control decoders	Working registers
2	✓ (A)	—	—	Working registers	ALU
2	✓ (MD)	—	—	MD buss	Working registers
1	✓ (W)	—	—	Working registers	Working registers
2	—	✓	—	Control registers	Decoders
2	—	—	✓	Timing circuits	Working registers
1	—	—	✓	Within working registers	
31	21	7	3	Totals	

100 circuits. The most heavily loaded nets occurred in the ALU output busses (Z_1 and Z_2), and all eight bits in each buss exceeded 5 pF, with the 16 nets ranging from 6.1 to 12.0 pF.

Second, as the ALU output busses indicate, the larger the function being driven and the higher the fan-out (FO) consisting of many loads distributed across pages, the longer the nets will be. It is therefore possible to identify the majority of high-capacitance nets before placement and wiring by just examining the data flow for nets between pages with high FO . These high-capacitance nets can be called nets with high dispersal. The driving circuits for these nets can be powered up, or special circuits can be used with low sensitivity to FO and capacitance. In the microprocessor of this paper, long nets occurred in the output buss of the ALU, which drives a total of 19 working registers.

These circuits with high capacitance load and critical delay are modified to have a collector pull-up resistor of 4 k Ω (refer to Fig. 3). The 13% of the circuits with a capacitance load greater than 1 pF use the 4-k Ω pull-up resistor. Figure 9 is a distribution of turn-off delay for all the circuits on the chip. The delay variation includes only the

variation in turn-off delay due to the load capacitance each internal circuit must drive on chip. Six percent of the circuits have a turn-off delay equal to two times average even with the 4-k Ω pull-up resistor. If the 8-k Ω pull-up had been used exclusively, the turn-off delay for 6% of the nets would be greater than three times the average. The turn-on delay distribution is narrower since its sensitivity to load capacitance is less.

The increase in average power for each circuit due to the 4-k Ω pull-up resistor is equal to 0.17 mW for 13% of the circuits. The total chip power increases by only $0.13 \times 4700 \times 0.17 = 0.1$ W or 4.5%. The selective use of the 4-k Ω pull-up resistor, therefore, gives a good power/performance gain.

Voltage drop checking and results

The previous sections have quantified the effects of long nets on delay. Another important effect of long nets with high fan-out is the dc voltage drop between the internal circuit and the loads being driven. This occurs at a down-level state for T²L or STL. In circuits such as SCS [3] or CSEF (Current Switch Emitter Follower), the ON current in the load is small, since this current is divided by beta at the circuit's base input point. As geometries are reduced and the widths of the metal interconnections are reduced to achieve higher wiring density, the voltage drops occurring in interconnections between T²L-type circuits will increase and these nets must be checked. The power levels are decreasing as more circuits are placed on chip and this compensates for the increasing net resistance.

For this chip, the widths of first- and second-metal interconnection lines are 3.6 μ m and 4.0 μ m, respectively. The load current is equal to 0.15 mA per fan-out (FO). The noise tolerance analysis allowed for an 80-mV drop in the interconnection lines between internal circuit and fan-out. Therefore, for a net the allowable resistance is

$$R = 80 / (FO \times 0.15) = 530 / FO.$$

For a driver at one corner of a chip and four loads at the opposite corner, assuming a Manhattan type of routing, the total resistance would be equal to 126 Ω . This is less than the allowable maximum of $530/4 = 132 \Omega$ and so this net would be permitted. Figure 10 is a histogram of the FO for the 4437 nets contained on the chip. As the above example illustrated, only nets with $FO > 4$ can be in violation. A total of 21 nets exceeded the limit of 80 mV.

Seven of the nets were rewired manually. The other 14 nets were not revised, either because they were in marginal violation of the rule or because due to logic orthogonality the current was below the worst-case assumption.




Figure 10 Fan-out distribution for microprocessor. Total number of nets = 4437.

Two of the seven nets required the driver to be moved closer to the loads. The availability of 35% unoccupied cells made it easy to find empty locations in the vicinity of the loads. The rewiring, replacement, and rechecking took approximately six man-weeks due to the complexity of the nets and the total chip data content.

In summary, due to the low power dissipation of the T²L circuit used and the availability of unoccupied cells on this chip, the voltage drop violations were manageable even with the 3.6- μ m-wide lines and the long nets occurring on this large gate array.

Discussion of results and recommendations

Continued advances in chip circuit count with restriction in chip power dissipation will make necessary further reductions in internal circuit power dissipation below 0.25 mW. This will lead to greater sensitivity to load capacitance for conventional logic circuits including T²L, CSEF, I²L, etc., all of which are examples of circuits operated at low power supply voltage levels (1.0 to 3.0 V) with active (low-impedance) drive in one direction only.

Active drive in one direction only means that the delay in the opposite direction will be "RC" time constant limited for high capacitance loading. For VLSI, R must be larger to keep dc power low. Typical values of resistance are between 4 and 10 k Ω for collector drive or emitter




Figure 11 Comparison of circuit delay sensitivity *versus* load capacitance. Power dissipation = 0.4 mW for all circuit types.

follower drive, yielding time constants in the 20- to 50-ns range for the slow transition. The slow transitions are positive going for collector drivers and negative going for emitter follower drivers. These two types of drivers are called one-direction-active. Push-pull drivers usually consist of npn transistors connected on top of one another, emitter to collector, with the top device performing as an emitter follower and the bottom device as a grounded emitter with collector drive. Reference [15] describes a push-pull T²L-type circuit which requires only 1.7 V.

Figure 11 is a graph of delay for the slowest transition case for commonly used bipolar circuits. The delay is plotted as a function of load capacitance for the different circuit types all designed at 0.4 mW average power dissipation. Note the wide disparity in delay at 5 pF load capacitance. T²L has the advantage over SCS and CSEF because these two circuits dissipate power in the off state. Also, T²L in the slow transition (up-going output) does not have far to go to reach the input threshold voltage of the load circuit. The best performance can be achieved by push-pull circuits which can be estimated to be in the range shown crosshatched in Fig. 11.

To lower the sensitivity of these circuits to load capacitance, the power should be increased on a selective basis using the results of the placement and wiring programs. The resistor changes must not cause any additional global wiring blockages so the intercircuit wiring will be unaffected.

Another method for narrowing the delay distribution is to selectively change the circuit type to one with active drive in both directions. One should do this before place-

ment and wiring, based on a study of the logic description of the interconnections between major logic functions.

Summary

Comparisons between circuits have been made on the basis of requirements of VLSI microprocessor design. The interrelationship among the microprocessor interconnections, net length, and performance has been quantified. This paper has described the following new features and design issues for a VLSI microprocessor: 1) A microprocessor implemented on a bipolar gate array; 2) Circuit selection and design for this VLSI implementation; 3) A low-power T²L circuit with a 0.7-pJ speed/power product; 4) Variable placement of unoccupied cells for optimized wiring channel location; and 5) Identification and solution of wide delay distributions due to high-capacitance nets. It has been shown that selectively increasing the circuit power based on wiring results reduces the size of the tail on the turn-off delay histogram and minimizes the skew between turn-off delay and turn-on delay. In doing this chip power increases only slightly because only a small percentage of the nets are involved.

Acknowledgments

The design, processing, and testing of this microprocessor chip was a cooperative effort between many areas of IBM in East Fishkill, NY and Boeblingen, West Germany. The author wishes to acknowledge the contributions of some of the people involved, namely C. Davis and his department, for the design and test patterns for the microprocessor; Departments 772 and 141, for the circuit and chip design of the gate array and coordination of hardware fabrication; F. Kirk's department, for the generation and checking of the physical data; and M. Ecker, for the module design. Also, N. Nan's departments wrote the software for the wiring programs, and in particular, J. Lee generated the wiring length data. The first chips were processed by S. Beck and tested by B. Schmidt in Boeblingen.

References and note

1. J. Pomeranz, R. Nijhuis, and C. Vicary, "Customized Metal Layers Vary Standard Gate-Array Chip," *Electronics* **52**, No. 6, 105-108 (1979).
2. J. Z. Chen, W. B. Chin, T.-S. Jen, and J. Hutt, "A High-Density Bipolar Logic Masterslice for Small Systems," *IBM J. Res. Develop.* **25**, 142-151 (1981, this issue).
3. R. J. Blumberg and S. Brenner, "A 1500 Gate, Random Logic, Large Scale Integrated (LSI) Masterslice," *IEEE J. Solid-State Circuits* **SC-14**, 818-822 (1979).
4. P. E. Fox and W. J. Nestork, "Design of Logic Circuit Technology for IBM System/370 Models 145 and 155," *IBM J. Res. Develop.* **15**, 384-390 (1971).
5. R. R. Wilcox, "A High Speed Circuit Application," *IEEE International Convention Digest*, March 1971, New York, NY, pp. 130-137.
6. J. G. Posa, "A Special Report on Gate Arrays," *Electronics* **53**, No. 21, 145-158 (1980).

7. Portions of this paper appeared in the article by A. H. Dansky entitled "Bipolar Circuit Design for VLSI Gate Arrays," which appeared in the *Proceedings of the IEEE International Conference on Circuits and Computers ICC '80*, Port Chester, NY, October 1980, pp. 674-677.
8. C. Davis, G. Maley, R. Simmons, H. Stoller, R. Warren, and T. Wohr, "S/370 Bipolar Gate Array Micro-Processor Chip," *Proceedings of the IEEE International Conference on Circuits and Computers ICC '80*, Port Chester, NY, October 1980, pp. 669-673.
9. N. Nan and M. Feuer, "A Method for the Automatic Wiring of LSI Chips," *Proceedings of the 1978 IEEE International Symposium on Circuits and Systems*, New York, May 1978, pp. 11-16.
10. T.-S. Jen and N. Nan, "Gate Array Experiences in IBM," *Proceedings of the 1980 Electro Professional Program*, Boston, MA, May 1980, Session 22-3.
11. *Advanced Statistical Analysis Program (ASTAP) Program Reference Manual*, Pub. No. SH20-1118-0, IBM Corporation, Data Processing Division, White Plains, NY 10604.
12. M. Feuer, K. H. Khokhani, and D. A. Mehta, "The Layout and Wiring of a VLSI Microprocessor," *Proceedings of the IEEE International Conference on Circuits and Computers ICC '80*, Port Chester, NY, October 1980, pp. 678-679.
13. E. Bloch, "VLSI for the 1980's," *Circuits Manufacturing* **18**, 16-26 (1979).
14. W. B. Davenport and W. L. Root, *An Introduction to the Theory of Random Signals and Noise*, McGraw-Hill Book Co., Inc., New York, 1958, pp. 34-35.
15. A. H. Dansky, "Push-pull T²L Internal Circuit," *IBM Tech. Disclosure Bull.* **23**, 1431-1432 (1980).

Received September 9, 1980; revised November 21, 1980

The author is located at the IBM General Technology Division laboratory, East Fishkill Facility, Hopewell Junction, New York 12533.