FCFS, and yet the throughput increase is only from 3.5 to 3.9 requests per second. In the analysis of the SSTF policy, the assumption that the n tracks to be accessed (for the requests present) will always be uniformly distributed between the inner and outer cylinder is questionable.

A new analysis of these algorithms would be considerable subject matter for another paper.

J. L. Smith
Research and Development Center
General Electric Company
Schenectady, N. Y.

B. LARGE-SCALE INTEGRATION

R68-21 System Architecture for Large Scale Integration—H. R. Beelitz, S. Y. Levy, R. S. Linhardt, and H. S. Miiller (*1967 Fall Joint Computer Conf., AFIPS Proc.*, vol. 31. Washington, D. C.: Thompson, 1967, pp. 185–200).

This paper proposes an approach to improving the regularity of computer logic for the more effective use of LSI which has been limited by the number of pins and wiring layers.

The authors' idea is to partition the computer logic into several processing modules each with a certain amount of local control and a set of buses interconnecting the modules. An instruction is executed by transfers of information between the modules, thereby simplifying centralized control and accordingly reducing the irregularity of the computer logic associated with control.

The approach is made in two steps: the register machine concept, followed by the LIMAC (large integrated monolithic array computer) concept.

In the register machine concept, each module, called an instruction execution unit (IEU) executes a particular type of instruction. Each IEU is provided with its own control, functional logic, and a register to interface the buses upon which instructions are transferred, operands fetched, and the results carried to be stored into memory.

To estimate the number of extra gates required for the additional regularity, the authors made trial designs for two existing machines (RCA 301 and RCA 4102) in the register machine form. Gate counts for the trial designs are the same order for the RCA 301 and three halves to twice as many for the RCA 4102 as those for the conventional designs.

The authors see improvement of the regularity in the register machine versions of both machines; however, the improvement in the RCA 4102 is less significant than that in the RCA 301 in which the gate count proportion for control is reduced from 65 to 25 percent. The differences in reduction are attributed to wider data paths and smaller control of the parallel-word RCA 4102 compared to those of the RCA 301 which is a character machine.

The second step, the LIMAC concept, is a microprogrammed version of the register machine concept, in which each module, called a function execution unit, handles a microinstruction or an elementary operation. Here the central control is reduced to a memory addressing control, and a machine instruction is executed by a string of information transfers between the function execution units.

Source and destination register names for a transfer are transmitted over the register identification bus and decoded by the recognition logic which enables the appropriate gates of the corresponding registers. The operation to be performed within the unit is specified by setting a suitable bit pattern into a register which controls the arithmetic matrix in the unit.

The advantage of the LIMAC concept is illustrated in terms of the pin-to-gate ratio in an experimental computer for navigational problems being built according to the concept. The ratios are 1:7.2 for 1000 gates and 1:2.6 for 100 gates in the LIMAC machine, and are described to be a factor of two or more greater than those of conventional architectures.

Details are given of the structure of the LIMAC, which is 16-bit fixed-point binary machine with 4096 words of 1-μs core memory and 256 words of 100-ns MOS memory for fast microprogram execution.

The final section of the paper concerns the circuit technology employed in the LIMAC machine including diffusion of chips as a collection of identical cells, inter- and intracell metallization of the cells for functional assignment, cell diagram, speed, and power dissipation.

The authors' approach to partitioning a machine into modules also suggests the possibility of increased throughput from concurrent operation of the modules in addition to the gains from the additional regularity of the machine logic.

Masao Kato
Dept. of Computer Science
University of Illinois
Urbana, Ill.

R68-22 Current Status of Large Scale Integration Technology—R. L. Petritz (*1967 Fall Joint Computer Conf., AFIPS Proc.*, vol. 31. Washington, D. C.: Thompson, 1967, pp. 65–85).

This paper combines a useful review of progress during the past year, a description of present programs and immediate goals, and some forecasts of ultimate goals. Although the author, as director of Texas Instruments' semiconductor research and development laboratories, deals primarily with programs and points of view at TI, he makes enough references to work at other companies to build up a reasonably comprehensive picture of the large-scale integration field. Four main areas are covered: bipolar chip technology, full-slice technology, MOS technology, and hybrid technology.

Biploar (i.e., n-p-n or p-n-p transistors as opposed to field effect devices) chip technology is, of course, common to essentially all present-day monolithic integrated circuit production. Typically, some 100 or so identical circuits, including metallized connections, are processed on a slice of silicon. Probing is done to identify and mark defective circuits, and the circuits are then scribed and separated into rectangular chips. The defective chips are discarded.

The author makes the point that bipolar chip technology has already progressed from the first to the second level of complexity. The first level includes the simpler elements of the TI series 54/74 TTL logic, such as gates and J–K flip-flops. Average figures are 4.2 gates per logic block on a chip of 2600 mil^2 area, or 620 mil^2 per gate. The second level, classified by the author as integrated electronic components (IEC) rather than integrated circuits (IC), is typified by a list of 16 recent additions to the 54/74 series, including a quadruple latch, an 8-bit shift register, and an active memory. Averages for the list are 25 gates per chip, area of chip 7000 mil^2, or 280 mil^2 per gate.

Extension of the single-chip technology to LSI levels of complexity (100 gates or more) will require larger chip areas and smaller circuits. The latter objective seems assured by predictable improvements in optics. Larger chip areas impose a serious limitation on yield; for randomly distributed defects, the expected yield would decrease exponentially with increasing area. However, careful experiments show a substantially more optimistic result, primarily because defects tend to occur in clumps. The author predicts that by the 1970's chip sizes of 62 500 mil^2 (one-fourth inch$\times$one-fourth inch) and gate areas of 250 mil^2 will be attained, for a resulting 250 gates per chip. Active bipolar memories of 500 bits per chip are also predicted.

For even higher levels of complexity, the full slice seems to be the best module. The method advocated by TI for overcoming yield problems is a combination of redundancy and discretionary wiring. Individual circuits (gates or active-memory cells) are laid out in a rectangular array, with extra elements in each row (or column). Probing