

Communication in Computation

Robert W. Keyes

IBM T. J. Watson Research Center, Yorktown Heights, New York 10598

Received May 6, 1981

The need for communication is an essential factor shaping the physical nature of a computing system. A computing system functions by bringing a large number of separate elements to bear on a common problem. Coordinated operation of the elements requires a large amount of communication among them through many long wires. Some implications of communication for the physical characteristics of electrical computers are as follows. Variability in manufacturing and of the environment within a system causes the elements of a system to differ from one another in their responses to electrical signals. Signals must be large enough to be interpretable by any element of the system. Therefore, the need for rapid reliable communication implies that signals that are large in some sense are required. The combination of large signals and many wires implies high-power dissipation. Fitting the pattern of interconnection into minimum space is a very complex problem which is one of the most severe limitations on the utilization of integrated electronics in computation. Experience has led to the establishment of certain quantitative generalizations concerning the communication requirements in large systems. Similarity between these and neuronal communication is noted.

1. INTRODUCTION

Much attention has been devoted to the devices that can implement digital computational operations in physical form. The devices must combine digital, almost always binary, signal inputs in a nonlinear way to produce output signals that represent logical functions of the inputs, such as ANDs, ORs, and NORs. Solid state research and development efforts also focus on the inherent speed and power dissipation of devices. A variety of more or less fundamental physical restrictions on devices have been recognized.

The purpose of computers, however, is to bring many logic devices to bear on a common problem. This requires a large amount of communication among the devices. In fact, an examination of large computers reveals that

their physical characteristics are dominated by the requirements for communication among the devices.

Communication and information theory is generally aimed at optimizing the utilization of bandwidth and power. Communication in computation must be viewed in a somewhat different light. Computation requires that two or more streams of information interact. Thus, the times of their arrival at a device are important. The coding of data streams that allows efficient use of channel capacity in communication cannot be used in computation because methods of information processing through the interaction of two or more streams of coded data are not known. Further, frequent coding and decoding of data would consume power and space, both of which are scarce resources in a system. Any increase in the physical size of a system means that signals spend more time traveling from one place to another in comparison to the time spent in actually interacting in logic devices.

2. REQUIREMENTS OF COMMUNICATION

The need for communication imposes several physical constraints on an electronic computing system:

Electrical conductors must be provided to carry signals through the system. The physical dimensions of the conductors are limited by the ability of manufacturing processes to produce them, and by physical phenomena, namely, ohmic resistance and electromigration. The need for access for assembly and replacement of parts requires that some signal paths pass through connectors of macroscopic size in large systems. Providing communication plays a major role in determining the physical size of a system.

Signals must be large enough to overcome the variability in the response of receiving devices. Part of the variability derives from irreproducibility in manufacturing; for example, variations in doping level, surface conditions, and oxide thickness make control of the threshold voltage of an insulated-gate field-effect transistor to within a few tenths of a volt difficult. A more fundamental source of variability among devices is random fluctuations in the number of dopant atoms in small regions. Additional differences among devices, up to 0.1-V shifts in characteristics, arise from temperature differences in large systems. Thus signal amplitudes must be at least several tenths of a volt and may have to be a few volts. The signals are large compared to thermal noise and other sources of electrical noise inherent in devices. The variability among components plays a role akin to that of the usual electrical noise.

The electrical conductors may be regarded as transmission lines, but frequently effectively reduce to capacitors or inductors. Electrical signals on

the lines create electromagnetic fields in space, the energy of which is dissipated in the operation of the system. The combination of large signals and many wires implies high-power dissipation. Removal of the heat produced is a major consideration in the physical design of a system.

The pattern of interconnections in a computing system is not regular but very complex. In fact, useful generalizations can be obtained by assuming that it has a pseudorandom character and application of probability theory (Sutherland and Oestreicher, 1973; Heller et al., 1977; Donath, 1968; El Gamal, 1981). In detail the complexity of the communication pattern is so great that the design and physical implementation of a system is very difficult and time consuming.

3. THE WIRE-DOMINATED CHIP

Communication among the 1500 gates or elementary logic circuits on a recently described logic chip required 4 m of wire (Blumberg and Brenner, 1979). If the gates are arranged on a square lattice the wire length is 2.8×10^4 gate-to-gate distances. Minimizing the space required by such large interconnection nets is an objective of design. A combination of experience and mathematical analysis has led to approximate quantification of certain of the communication requirements in electronic computers. One useful observation relates the average connection length on a complex logic chip to the number of logic gates on the chip (Donath, 1979, Heller et al., 1977). Accommodating the longer wires means that more space on the chip must be allowed for the wiring pattern. The number of wire channels required, K , can be estimated and is found to be related to N , the number of gates on the chip, in the way shown in Figure 1 (Heller et al., 1977). The figure means that one divides the chip into rectangular cells each containing one logic gate. Each cell, then, must contain K channels in which a wire might be fabricated to pass through the cell. Provision of too few channels decreases the probability that the desired interconnection pattern can be fabricated on the chip. For example, a chip containing 10^3 gates must allow space for about 20 wire channels. Ordinarily half of these would run in the X direction and half in the Y direction, and the two directions of wire would be in separate layers. The wire channels would probably occupy much more space in the cell than the semiconductor devices.

In fact, then, the physical nature of the wires leads directly to an approximation to the size of the chip. Let the width of a wire channel be w , the number of wire layers be M , and the area of a logic cell be A . Then the total channel length in a logic cell is

$$L = KA^{1/2} \quad (1)$$




Fig. 1. The number of wire channels needed per logic cell as a function of the number of logic gates on a chip (Heller et al., 1977).

However, if the device area is negligible MA is the area provided that accommodates the channels:

$$MA = Lw \quad (2)$$

Thus, from (1) and (2),

$$A = (Kw/M)^2 \quad (3)$$

The total area of the chip is NA .

The physical characteristics of the wiring and the voltage requirements of the semiconductor devices also determine the energy dissipation per operation in this approximation. Let a fraction f of the total available channels be occupied by wire with a capacitance C' per unit length. Then the amount of charge stored in the wire by the average logic gate during an operation is

$$Q = C'fL\Delta V = C'fK^2w\Delta V/M \quad (4)$$

ΔV is the signal voltage. If Q is drawn from a power supply of voltage V_B the energy dissipation, the power-delay product of the logic, is

$$U = C'fK^2wV_B\Delta V/M \quad (5)$$

The point is that U is determined by the physical characteristics of the communication network rather than by such prominent semiconductor

properties as mobility and limiting velocity. The principal demand on the devices is that they be small enough to require only a minor increase in the size of the wire-limited logic cell.

4. PACKAGING

Large systems are constructed from many chips. The art of interconnecting chips to form a large electronic computer is known as packaging. Communication is accomplished through a packaging hierarchy. Electronic devices and elementary logic circuits are fabricated on a chip, which contains wiring patterns. Some gates must communicate with gates on other chips so that a certain number of off-chip interconnection terminals, here called pins, must be provided. The chips are mounted on a multichip substrate called a module that must contain means for connecting to the chip pins, a wiring matrix that interconnects the chip pins, and a further set of pins of a different type that can carry away signals destined for other modules. The modules are in turn mounted on boards that must be able to receive the module connections and provide another matrix of interconnections among them. Means for removing heat by transferring it to some fluid that eventually carries it out of the system must also be provided. It must be understood that the details of the package parts and the names assigned to them differ widely among machines and manufacturers.

A few quantitative examples will illustrate the magnitude of the communication requirement. One hundred chips are mounted on a ceramic module 9×9 cm² in the IBM 3081 computer. Each chip has about 100 signal pins, and the module contains over 100 m of wire. The board on which the modules are mounted measures 60×70 cm² and contains 1 km of interconnections. Rather similar numbers characterize other computing systems (Beall, 1974; Wilson, 1978; Chiba, 1978; Werbizky et al., 1980).

The long wires in the package hierarchy mean that there will be long gate-to-gate transit times in some logic operations. These long paths limit the rate at which the computer can be cycled (Wu, 1978). An electronic signal can travel a few meters on the interconnections during one cycle of a modern computer. It is interesting to note that this is about the physical dimension of a central processing unit of a computer (Edwards, 1978).

One might wonder why all of the logic gates are not formed and interconnected on a single slice of silicon, instead of being packaged in this hierarchical fashion. One answer lies in the manufacturability of the communication network. Dense packing of circuits requires narrow lines and it is not possible to manufacture more than several meters of such lines without an unacceptable probability of a defect. Similar constraints are encountered at each level of packaging.

5. COMPLEXITY

An earlier section mentioned that the average line length (measured in logic gate-to-logic gate units) increases with the number of gates on the chip. A derivation of the result has been given by Donath (1979), who connects it with a relation known as "Rent's rule" (Landman and Russo, 1971; Chiba 1978). The rule relates the number of signal connections that must be made to a packaged portion of an information-processing system, for example, a chip or a multichip carrier, to the number of gates contained in the portion. It has the form

$$P = aN^b \quad (6)$$

P is the number of connections or pins and a and b are parameters. The value of b is about $2/3$, but is given somewhat differently by different authors. The rule must be viewed as an empirical one. Minimizing the number of pins is desirable; an off-chip connection is less reliable and much more expensive to fabricate than a connection on the chip and also requires high power to drive the long off-chip interconnection. There is a long signal delay in the off-chip path. Equation (6) represents the number of pins that chip designers find by experience is needed to permit utilization of most of the gates on a chip.

One way to look at the dependence of pins and wire lengths on level of integration is as follows. A chip designer must have a certain amount of flexibility in interconnecting gates in order to implement a logic diagram in hardware. In other words, the physical constraints must permit a degree of choice among a multiplicity of possible interconnection patterns. Increasing N , the number of gates on a chip, causes, according to equation (6), a decrease in the number of interconnections to allow a wire to reach more different possible destinations. It seems clear that such a trading of on-chip and off-chip wiring flexibility can only be valid over a limited range, however; chips must be able to transmit to and receive from other chips.

The question of how much flexibility in the form of a multiplicity of possible communication patterns must be provided in order to make a functioning large system can be regarded as one measure of the complexity of a system. The number of possible interconnection patterns on a chip can be estimated from a knowledge of the average wire length. Let a chip contain Q_k pairs of circuits that might be joined by an interconnection of length k (measured in units of the circuit lattice constant on a rectangular grid). For example, in a chip containing 5 circuits by 5 circuits $Q_1=40$, $Q_2=62$, etc. Let p_k be the number of wired connections of length k . Then

the total length of wire is

$$L_T = \sum k p_k \quad (7)$$

We seek the minimum length that will permit some number of ways, Ω , of wiring the chip. The value of Ω is

$$\Omega = \prod_k Q_k! / p_k! (Q_k - p_k)! \quad (8)$$

In addition, there is a fixed number of connections, N_c .

$$\sum p_k = N_c \quad (9)$$

N_c is the total number of connections and is proportional to the number of circuits on the chips and L_T is the maximum total wire length.

It will be recognized that the problem posed is equivalent to maximizing the entropy of the distribution of the probability that a wire site is occupied, while maintaining a fixed length of wire. The problem of congestion in wire channels, which the result shown in Figure 1 is intended to minimize, has been ignored. The minimization is easily carried out by setting $Q_k = 2Nk$, its limiting value when the number of circuits on the chip, N , is large and using Stirling's approximation and the method of Lagrangian multipliers. The results are conveniently described by introducing the average wire length L_{av}

$$L_{av} \equiv L_T / N_c \quad (10)$$

and a parameter z defined in terms of L_{av} by

$$z = (L_{av} - 1) / (L_{av} + 1) \quad (11)$$

The probability that a particular wire site of length k is occupied is

$$q_k = (N_c / 2N) (1 - z)^2 z^{k-1} \quad (12)$$

Also Ω turns out to have the value

$$\log \Omega = N_c \log [2N e z / N_c (1 - z)^2] - L_T \log z \quad (13)$$

A good approximation to this rather unwieldy result if the average length is at least 3 is the more transparent

$$\log \Omega = N_c(2 + \log L_{av}^2) + N_c \log(2Ne/N_c) \quad (14)$$

The first term of equation (14) comes from the choice among destinations that a wire might reach. It will be noted that the area (in number of logic gates) that can be reached by a connection no longer than L_{av} is $2L_{av}^2$. The second term arises from the entropy associated with the location of a wire. These results are approximate also in that no account is taken of the possibility that different wiring patterns may provide functionally equivalent interconnection schemes.

As an example, the parameters of a chip that contained an entire computer central processing unit are given in Table I (Dansky, 1980). Equation (14) yields $\log \Omega = 6.5 \times 10^4$ or $\Omega = 370$ per interconnection. Even when the distribution of wire lengths does not have the optimal form, the extremum property insures that a reasonable approximation to Ω is obtained.

A complete large computer contains 10^5 or 2×10^5 logic gates and many different kinds of chips. The communication among them is also very complex, being accomplished by up to 30 layers of signal wiring on the multichip packages. For example, a module may hold 100 chips, each with 100 pins. The pins can be connected in pairs in about $\exp 10^5$ ways. (Wiring length restrictions will reduce this slightly.) The point of these numbers is to illustrate the magnitude of the problem of physical design and layout, the difficulty of which is one of the limits to the application of integrated electronics in computation.

6. BIOLOGICAL ANALOGS

Questions as to how fundamental some of the characteristics of the communication system are naturally arise. Clearly, the variability among

TABLE I. Parameters of a Computer CPU Chip^a

Quantity	Value
Number of logic gates (N)	4705
Circuit-to-circuit distance	0.075 mm
Number of interconnections (N_c)	11,000
Total wire length	5.8 m
Average interconnection length	0.53 mm
Dimensionless interconnection length (L_{av})	7.0

^aDansky, 1980.

devices that must be overcome by large voltages could be reduced by stricter selection, a higher standard of acceptability for devices. An economic factor is involved. It may be possible to reduce the average wire length on complex chips or to circumvent Rent's rule by improved design procedures of new views of logic. Thus, it is interesting to note some analogies between features of communication in computers and in biological systems.

Vertebrate nervous systems provide another example that confirms the dominance of complex information-processing systems by the physical provision of communication. Cerebral neurons are characterized by long dendrites and axons that allow communication with many other neurons, often many thousands of others. Trends such as Rent's rule and the need to devote a larger fraction of available space to communication channels as the size of a system increases have parallels. Figure 2 plots numbers related to certain connections in the human nervous system and compares them with Rent's rule as found in computer logic. Of course, the suggested equivalence of the neural cells to a logic gate as used in a computer is open to question,

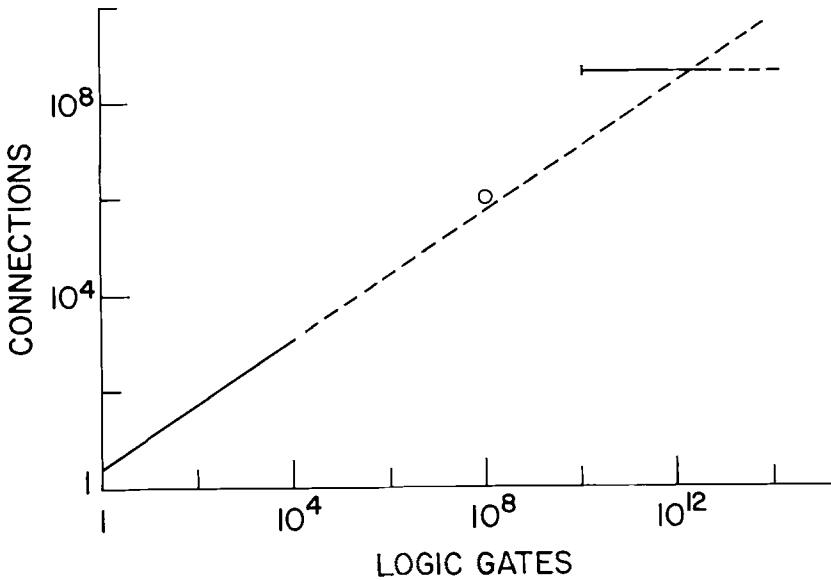


Fig. 2. Rent's rule, relating the number of interconnections needed to a partition of a system containing N logic elements to N . The solid line is the result of experience with electronic computers. The points approximate the optic nerve (10^8 photoreceptors in the human eye, 10^6 fibers in the optic nerve) and the human brain (10^{10} neurons per hemisphere, equivalent to some larger number of logic gates) and 2×10^8 fibers in the *corpus callosum*. (Keyes, R. W., *Proceedings IEEE*, **69**, pp. 267-278, 1981. Copyright 1981, IEEE.)

so the figure should be regarded as illustrating a trend rather than having quantitative significance. In particular, however, the continuation of the trend to such large numbers in the natural system indicates that there may be functional advantages in avoiding the division of complex systems into distinct "functional units" which can communicate with one another through a very few interconnections.

An observation reminiscent of Figure 1 has been made for cells of the mammalian cerebral cortex (Jerison, 1973). A comparison of the brains of several species shows that the larger the brain the fewer neurons it contains per unit volume, more spacing being occupied by dendrites and axons rather than cell bodies. The dendrites are longer and can reach more other cells (Bok, 1959).

Again, because of the great functional difference between computer logic elements and neurons one cannot take the numbers involved too seriously, but their qualitative similarity suggests that basic trends in communication in information-processing systems are involved.

7. CONCLUSION

Computers containing a large number of elements require an extensive network of communication among them. The communication is supplied by a large amount of wire on integrated circuit chips and by a packaging hierarchy to interconnect the chips. Large signals are needed to overcome unavoidable differences among the logic elements. The space occupied by wires and interconnections and the power dissipation and delays caused by large signals in the communication system dominate the physical nature and the performance of computing systems. Some of the trends found in electronic computers are paralleled by trends in mammalian brains.

REFERENCES

- Beall, R. J. (1974). "Packaging for a supercomputer," paper 18/3, *1974 INTERCON Technical Papers*. IEEE, New York.
- Blumberg, R. J., and Brenner, S. (1979). "A 1500 gate random logic LSI masterslice," *IEEE Journal of Solid State Circuits*, **14**, 818-822.
- Bok, S. T. (1959). Elsevier, Amsterdam. pp. 185-232. *Histonomy of the Cerebral Cortex*.
- Chiba, T. (1978). "Impact of the LSI on high-speed computer packaging," *IEEE Transactions on Computers*, **27**, 319-325.
- Dansky, A. H. (1980). "Bipolar circuit design for VLSI gate arrays," pp. 674-677. *Proceedings ICC 80*. IEEE, New York.
- Donath, W. E. (1968). "Statistical properties of the placement of a graph," *SIAM Journal of Applied Mathematics*, **16**, 376-387.

- Donath, W. E. (1979). "Placement and average interconnection lengths of computer logic," *IEEE Transactions of Circuits and Systems*, **26**, 272–277.
- Edwards, N. P. (1978). Unpublished.
- El Gamal, A. A. (1981). "Two-dimensional stochastic model for interconnection in integrated circuits," *IEEE Transactions on Circuits and Systems*, **28**, 127–135.
- Heller, W. R., Mikhail, W. F., and Donath, W. E. (1977). "Prediction of wiring space requirements for LSI," pp. 32–42. *Proceedings 14th Design Automation Conference*, June 20–22, 1977, New Orleans.
- Jerison, H. J. (1973). *Evolution of the Brain and Intelligence*, pp. 68 and 69. Academic Press, New York.
- Landman, B. S., and Russo, R. L. (1971). "On a pin versus block relationship for partitions of logic graphs," *IEEE Transactions on Computers*, **20**, 1469–1479.
- Sutherland, I., and Oestreicher, D. (1973). "How big should a printed circuit board be?" *IEEE Transactions on Computers*, **22**, 537–542.
- Werbizky, G. G., Winkler, P. E., and Haining, F. W. (1980). "Packaging technology for the IBM 4300 processors," *1980 Spring COMPCON*, pp. 304–311. IEEE, New York.
- Wilson, E. A. (1977). "True liquid cooling of computers," *1977 National Computer Conference Proceedings*, Montvale, New Jersey, pp. 341–348. AFIPS Press.
- Wu, L. C. (1978). "VLSI and mainframe computers," *Spring '78 Compcon Digest*, pp. 26–29. IEEE, New York.