

Continuous Models for Communication Density Constraints on Multiprocessor Performance

LANCE A. GLASSER, MEMBER, IEEE, AND CHARLES A. ZUKOWSKI, MEMBER, IEEE

Abstract—Fundamental limits on the communication capabilities of massively parallel multiprocessors are investigated. It is shown that in the limit of machines of infinite extent in which the number of processors per unit volume is constant and in which the communication bandwidth from each processor to its neighbors depends only on their separation distance, interprocessor communication must fall off faster than the fourth power of distance. For machines of finite size, communication energy density is used as a metric to compare various machine sizes and packaging densities. For instance, for machines with spherical symmetry and uniform communication requirements, the peak density depends on the number of processors to the 4/3 power and the number of processors per unit volume to the 2/3 power.

Index Terms—Communication limits, information density, massive parallelism, multiprocessors.

I. INTRODUCTION

MASSIVELY parallel fine-grained multiprocessors [1], [2] generally require massive amounts of interprocessor communication. Communication bandwidth is thought to strongly impact the performance of multiprocessors by imposing limits on the degree to which processors can exchange information and cooperate on a single problem. In this paper, we examine the physical limits imposed on interprocessor communication by the capacity of the communication medium. Other fundamental constraints that limit the performance of large digital multiprocessors can be found in [3]. We also present relationships between communication requirements and communication capacity under a number of different assumptions about the multiprocessor array. In the spirit of massive parallelism, we use continuous rather than discrete variables in the formulation of our interprocessor communication model. This choice makes the mathematics cleaner without sacrificing the physics.

To measure the capacity of a communication medium, we have chosen to use information density. We first derive this quantity, and motivate its use, by investigating fundamental limits on the communication capacity of any interconnection network for a multiprocessor array. We know, from informa-

tion theory, that the energy E_0 required to transmit one bit of information is on the order of kT , where k is Boltzmann's constant and T is absolute temperature [4], [5]. Thus, high data rates require large quantities of power, independent of how the information is transmitted or coded. Note that these data rates can arise from the requirements of a single channel or the superposition of a number of different channels that interfere. If we examine a transmission medium of fixed diameter, whether it be a length of coaxial cable, a microwave waveguide, or a fiber optic link, the information velocity is limited to c , the speed of light. If the "wire" diameter is D , then the energy density in the wire must be at least $BE_0/(cD^2)$, where B is the bit rate and c is the speed of light (an upper bound on the speed of transmission). All physical media, with the exception of absolute vacuum, have a maximum energy density that they can withstand without breaking down. Since electromagnetic energy is usually the information carrying method of choice, limits on information density must ultimately arise when electric forces caused by the information signal become significant when compared to the forces binding electrons to their atoms. These electric fields will cause the medium to behave nonlinearly, distorting the information. Large enough fields will cause destruction of the medium by tearing it apart. The largest electric fields in solids can be supported in materials such as mylar, polypropylene, and various ceramics and glasses. These have intrinsic breakdown strengths in the 50–200 kV/m range [6]. One can counteract this phenomenon by either making D larger or running several cables in parallel. We therefore argue that a fundamental quantity which is physically limited is information density within a cross section of space, i.e., bandwidth per unit area. The highest values of information density are found in optical fibers and VLSI chips, where values of $10^{19\pm 2}$ bits/(s·m²) can be observed. Usually communication densities are much lower.

While information densities may be far from their fundamental limits in current technologies, there will always be practical limits and costs associated with different densities. Therefore, it is important to study the general relationship between information density and interprocessor communication. Given an array of processors and a communication medium of bounded volume, the interprocessor communication requirements necessarily generate information density requirements on the medium. Limits on information density, therefore, impose limits on the degree to which processors can communicate and hence cooperate. It is these limits which we will discuss in this paper.

Manuscript received May 14, 1986; revised December 18, 1986. This work was supported in part by the National Science Foundation under Grants ECS-8118160 and CDR-84-21402, by the Defense Advanced Research Projects Agency under Contract N00014-80-C-0622, and by an IBM Graduate Fellowship.

L. A. Glasser is with the Department of Electrical Engineering and Computer Science and the Research Laboratory of Electronics, Massachusetts Institute of Technology, Cambridge, MA 02139.

C. A. Zukowski is with the Department of Electrical Engineering, Columbia University, New York, NY 10027.

IEEE Log Number 8717692.

II. STATEMENT OF PROBLEM

Consider an array of infinitesimal processing elements that are exchanging information while carrying out a computation. In general, the array has a density of processors ρ that is a function of position. The rate at which information is flowing from a processor at point \vec{q} to one at point \vec{s} , denoted $I(\vec{q}, \vec{s})$, can also vary as a function of position. For simplicity, we assume that all functions discussed are continuous and bounded. This section presents the machinery needed to analyze the interprocessor communication requirements of such an array.

Depending on the physical implementation of the system, the array might have any number of dimensions up to three. To avoid treating each case separately, the equations presented here use a general number of dimensions K , and are valid for any positive integer K . Processors in two-dimensional space might represent an array on a planar integrated circuit. Large systems, consisting of many VLSI subsystems, are constructed to take advantage of all three physical dimensions. Thus, $K = 3$ must be taken as fundamental, although one might, as an engineering decision, restrict oneself to $K \approx 2$. Machines with $K > 3$, while interesting to contemplate, are not physically realizable. A 16-dimensional hypercube interconnect network, for instance, can only be built in 3-space.

First consider the information flowing out from a single processing element. If all communication between every two processors is assumed to be through the shortest path, the flux density of bandwidth from a single processor is a radial vector field. The use of the term flux density appeals to the analogy of electromagnetic field theory. That is, each processor emits, along a radial line, that information required by other processors in that direction. Near the transmitter, the information density due to the transmitter must contain the information intended for all receivers in a given direction. Further away from the transmitter, the information flux decreases as various receivers remove data. $\vec{F}(\vec{q}, \vec{s})$ represents the information flux density at point \vec{q} originating from a processor at point \vec{s} . $f(\vec{q}, \vec{s})$ denotes the scalar magnitude of the flux and has the units bits/(s·proc·m^{K-1}). We have the simple relation

$$\vec{F}(\vec{q}, \vec{s}) = \frac{(\vec{q} - \vec{s})}{|\vec{q} - \vec{s}|} f(\vec{q}, \vec{s}). \quad (1)$$

The function $f(\vec{q}, \vec{s})$, for a given processor at point \vec{s} , must be a decreasing function of the distance $|\vec{q} - \vec{s}|$ because other processors can only receive information from the processor at \vec{s} . In any dimension larger than one, the surface area available for communication increases with distance. Therefore, when nonzero, f must fall with distance at least as fast as $|\vec{q} - \vec{s}|^{1-K}$, the rate seen in the absence of receiving processors. Fig. 1 illustrates the relevant geometries. In the absence of receivers, the total information traveling outward from \vec{s} is independent of distance but the information density is not.

To represent a physical computer with discrete wires in terms of our continuous model, one must take a spatial average. For instance, if n discrete wires, each of area A and a bit rate B , intersect a surface of area W with a uniform

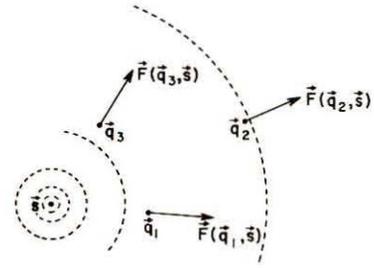


Fig. 1. The information flux leaving a point \vec{s} is a decreasing function of distance.

distribution, this can be modeled as a uniform information density of nB/W at the surface.

The flux from a single processor can be viewed as a vector field, but the contributions from different processors do not add as vectors. If two processors are exchanging information at equal rates, the total communication taking place is not zero. As a result, our definition of the flux magnitude $f(\vec{q}, \vec{s})$ assumes that it is always positive and incoming information is represented in the positive outgoing flux from other processors.

To find the total communication density at a particular point, we add the magnitudes of the information fluxes there originating from all processors. The total information flux through any point \vec{q} is given by

$$\Phi(\vec{q}) = \int_{\vec{s} \in \mathbf{R}^K} f(\vec{q}, \vec{s}) \rho(\vec{s}) d\vec{s}. \quad (2)$$

Φ represents the density of bits per second passing through an infinitesimal area of space and has the units of bits/(s·m^{K-1}). Said another way, Φ represents information rate density. It is this density which, we argued in the Introduction, is a physically significant and limited quantity. Equation (2) can be simplified by assuming spherical symmetry about $\vec{q} = 0$ and considering only the bandwidth density at the origin. As a result of symmetry, the only necessary coordinate is r , the distance from the origin. Equation (2) then becomes

$$\Phi(0) = C_K \int_0^\infty (r^{K-1}) f(0, r) \rho(r) dr \quad (3)$$

where

$$C_K = \begin{cases} 2 & \text{for } K=1 \\ 2\pi & \text{for } K=2 \\ 4\pi & \text{for } K=3 \\ \frac{(K-3)C_{K-1}C_{K-2}}{(K-2)C_{K-3}} & \text{in general.} \end{cases} \quad (3.1)$$

We now present the relationship between the information flux density $\vec{F}(\vec{q}, \vec{s})$ and the communication bandwidth $I(\vec{q}, \vec{s})$. Here, the actual information flow between processors is assumed independent. The special case of broadcast, which is a violation of this assumption, will be studied at the end of Section III. Recall that $I(\vec{q}, \vec{s})$, with units of bits/(s·proc²),

In the special case of broadcast, in which each processor broadcasts its message to all other processors, the minimum information density is obtained when each processor's message visits each receiving processor only once. One might imagine the information from a processor radiating out in all directions, with each processor repeating the message to those beyond it. In this case $f(\vec{q}, \vec{s})$ is proportional to $\rho^{(K-1)/K}$, the density of processors at \vec{q} within a space of $K - 1$ dimensions that is perpendicular to the vector $\vec{q} - \vec{s}$. In two dimensions, this corresponds to the density of processors along the circle centered at \vec{s} and passing through \vec{q} . By substituting this flux density into (3), we get

$$\Phi(0) \propto C_K \rho(0)^{(K-1)/K} \int_0^\infty r^{K-1} \rho(r) dr. \quad (15)$$

Along with (13), this implies

$$\Phi(0) \propto N. \quad (16)$$

Broadcast is indeed a more efficient, although more limited, form of communication.

IV. COMMUNICATION SCALING

Assuming that a processor array is finite, (6) also indicates how communication density scales with bandwidth requirements and array size. Here we use Φ as a measure of the engineering difficulty of implementing the communication network. Engineering experience widely suggests that the difficulty of building a computer increases as the speed of the wires are increased and as those wires are packed more closely together. The backplane of a CRAY supercomputer [8] is an excellent example of a network in which Φ is very large and expensive. The maximum value of Φ a technology can achieve is, we suggest, one reasonable figure of merit for an interconnect technology.

Consider the special case where ρ is constant over a finite array of radius R . Assume as before that I is only a function of distance d . Let $I(d)$ have the constant value $I_0 a^{-M}$ for $0 < d \leq a$ and be $I_0 d^{-M}$ for $d > a$, where a is some constant such that $0 < a < R$. From (6), we arrive at

$$\Phi(0) = C_K \rho_0^2 \int_0^R \int_r^{r+R} (r^{*K-1}) I(r^*) dr^* dr. \quad (17)$$

Evaluating (17), we find that

$$\Phi(0) = C_K I_0 \rho_0^2 (\beta + \delta R^{K-M+1}), \quad (18)$$

where

$$\beta = \frac{M a^{K-M+1}}{(K+1)(M-K-1)}, \quad (18.1)$$

$$\delta = \frac{2(2^{K-M}-1)}{(M-K)(M-K-1)}, \quad (18.2)$$

when M is not equal to K or $K + 1$. Consistent with the results of the previous section, if communication falls off with an order larger than $K + 1$, $\Phi(0)$ approaches asymptotically to the finite value of $C_K I_0 \rho_0^2 \beta$ as R is increased towards infinity.

If M is smaller than $K + 1$, (18) determines the rate at which $\Phi(0)$ approaches infinity with increasing radius R .

In the special case where $M = 0$ (I a constant function of distance), we find that $\beta = 0$. If the number of processors N is held constant, we have

$$\Phi(0) = \left(\frac{2K(2^K-1)}{C_K(K+1)} \right) N^2 I_0 R^{1-K} \quad (19)$$

giving rise to the three following relationships:

$$\Phi(0) \propto N^{(K+1)/K} \rho_0^{(K-1)/K}, \quad (20)$$

$$\Phi(0) \propto N^2 R^{1-K}, \quad (21)$$

$$\Phi(0) \propto R^{K+1} \rho_0^2. \quad (22)$$

From (20), we can see how communication density decreases with the spreading of a processor array. With uniform communication among a fixed number of processors, the communication density falls off as the volume to the power $(1 - K)/K$. For a two-dimensional array on a planar integrated circuit, every factor of four increase in the area produces a factor of two reduction in the communication density. If, on the other hand, the processor density is held constant and the number of processors is increased by a factor of four, the communication density is increased by a factor of eight.

Equations (19)–(22) allow us to understand better the constraints on making a machine, in which I is constant, larger. If we keep the technology constant, and by this we mean keeping $\Phi(0)$ constant, and restrict the discussion to $K = 3$, then we discover that the radius of the machine must increase linearly with N and the density must fall off as N^2 . In a given technology, a machine that has twice the number of processors occupies eight times the volume, assuming the technology is equally strained in both cases.

V. CONCLUSION

We have presented a continuous model for communication density in large multiprocessor arrays. Physical systems, regardless of their architecture, must contend with the limits we described. While our analysis is only valid for straight line communication, the average value of Φ can only increase when line-of-sight communication is not used. This is because Φ represents an information rate density. If information takes the straight line route, then it occupies the minimum volume over a minimum time of flight and hence contributes minimally to a volumetric and time average.

We have observed that Φ is a good measure of the difficulty and cost of building a machine's interconnect technology. The use of Φ is further justified by the fact that it does have a physical maximum based on material properties and information theory. This paper is essentially a discussion of the implications which follow from identifying communication density as a fundamental quantity. Equations (1)–(6) are a presentation of the general formulation of the theory and the remainder of the paper investigates the theory's application to finite and infinite machines. For infinite machines, we discover the intuitively appealing result that there are limits on our ability to focus them on a single problem. For finite

machines, the degree to which communication density must increase as the number of processors grows has been quantified in Section IV.

The purpose of this research is to allow the computer architect to think about machine tradeoffs when only general requirements on locality are specified—and before a network topology is chosen. If, for instance, locality is not exploited (I constant), then one can immediately make statements about how the size of the machine grows with processor number. Both the strength and weakness of our approach is that it deals with the functionality, rather than the architecture, of a network and hence the statements we can make are very general.

REFERENCES

- [1] W. D. Hillis, *The Connection Machine*. Cambridge, MA: MIT Press, 1985.
- [2] S. E. Fahlman, G. E. Hinton, and T. J. Sejnowski, "Massively parallel architectures for AI: Netl, Thistle, and Boltzmann machines," in *Proc. Nat. Conf. Artif. Intell.*, Washington DC, 1983, pp. 109-113.
- [3] R. W. Keyes, "Fundamental limits in digital information processing," *Proc. IEEE*, vol. 69, pp. 267-278, 1981.
- [4] C. E. Shannon and W. Weaver, *The Mathematical Theory of Communication*. Urbana, IL: Univ. Illinois Press, 1949.
- [5] R. P. Feynman, R. B. Leighton, and M. Sands, *The Feynman Lectures on Physics*. Reading, MA: Addison-Wesley, 1975.
- [6] J. J. O'Dwyer, *The Theory of Electrical Conduction and Breakdown in Solid Dielectrics*. Oxford, England: Clarendon, 1973, also C. Cooke, private communication.
- [7] W. Crowther, J. Goodhue, E. Starr, R. Thomas, W. Milliken, and T. Blackadar, "Performance measurements on a 128-node butterfly parallel processor," in *Proc. 1985 Int. Conf. Parallel Processing*, IEEE Comput. Soc. Press, Aug. 20-23, 1985.
- [8] R. M. Russell, "The CRAY-1 computer system," *Commun. ACM*, vol. 21, pp. 63-72, 1978.



Lance A. Glasser (S'73-M'79) received the B.S. degree in electrical engineering from the University of Massachusetts, Amherst, in 1974 and the S.M. and Ph.D. degrees from the Massachusetts Institute of Technology, Cambridge, in 1976 and 1979, respectively.

He is currently an Associate Professor in the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology. Since joining the faculty, his research interests have shifted from microwaves and picosecond optics, to high-performance VLSI circuits and systems. In addition to numerous articles and patents, he has written, with Dan Dobberpuhl, *The Design and Analysis of VLSI Circuits*, (Reading, MA: Addison-Wesley).

Dr. Glasser is the 1986 recipient of the ASEE Frederick Emmons Terman Award.



Charles A. Zukowski (S'83-M'85) was born in Buffalo, NY, on August 17, 1959. He received B.S., M.S., and Ph.D. degrees in electrical engineering from the Massachusetts Institute of Technology, Cambridge, in 1982, 1982, and 1985, respectively.

Between 1979 and 1982 he worked as a student at the IBM T. J. Watson Research Center, and from 1982-1985 he studied under an IBM fellowship. Since 1985 he has been an Assistant Professor in the Department of Electrical Engineering, Columbia University, New York, NY. He is the author of the book *The Bounding Approach to VLSI Circuit Simulation* (Hingham, MA: Kluwer, 1986). His research interests include VLSI circuit analysis, VLSI circuit design, CAD, and computer architecture.

Dr. Zukowski is a member of Tau Beta Pi, Eta Kappa Nu, and Sigma Xi.