

**Michael J. Flynn** was born in New York, NY, on May 20, 1934. He received the B.S.E.E. degree from Manhattan College, New York, the M.S. degree from Syracuse University, Syracuse, NY, and the Ph.D. degree from Purdue University, West Lafayette, IN, in 1955, 1960 and 1961, respectively.

He joined IBM in 1955 and over a ten year tenure there, worked principally in the areas of circuit development and computer organization and design. He was a designer and manager of prototype versions of the IBM 7090 and 7094/II and later was the manager responsible for design and development of the System 360 Model 91 Central Processing Unit. He was a member of the faculty of Northwestern University and The Johns Hopkins University, Baltimore, MD, from 1970-74. In 1973-74 he was on leave from the University to serve as Vice President of Palyn Associates, Inc., San Jose, CA, a computer design firm where he continues to serve as Senior Consultant. Since January 1975, he has been a Professor of Electrical Engineering at Stanford University and since 1977, Director of Computer Systems Laboratory.

Dr. Flynn has served as Vice-President of the IEEE Computer Society and as an Associate Editor of the TRANSACTIONS ON COMPUTERS. He was

Founding Chairman of both the Association for Computing Machinery Special Interest Group on Computer Architecture and the IEEE Computer Society Technical Committee on Computer Architecture.



**Lee W. Hoebel** (S'74-M'77) was born in Kansas City, MO, on January 22, 1946. He received the B.A. degree in mathematics and economics from Rice University, Houston, TX, in June 1968, and the Ph.D. degree in electrical engineering from The Johns Hopkins University, Baltimore, MD, in 1979.

In June 1968, he joined the staff of The Johns Hopkins Applied Physics Laboratory, Silver Spring, MD. He took a leave of absence in the fall of 1971 to become a full-time graduate student of The Johns Hopkins University. He was a Research Assistant with the Department of Electrical Engineering, Stanford University, Stanford, CA, from 1975-1979. He has been a member of the technical staff at the IBM T. J. Watson Research Center, Yorktown Heights, NY, since 1979. His current research interests include program behavior, memory hierarchies, and computer architecture.

# The NYU Ultracomputer—Designing an MIMD Shared Memory Parallel Computer

ALLAN GOTTLIEB, RALPH GRISHMAN, CLYDE P. KRUSKAL, KEVIN P. MCAULIFFE,  
LARRY RUDOLPH, AND MARC SNIR

**Abstract**—We present the design for the NYU Ultracomputer, a shared-memory MIMD parallel machine composed of thousands of autonomous processing elements. This machine uses an enhanced message switching network with the geometry of an Omega-network to approximate the ideal behavior of Schwartz's paracomputer model of computation and to implement efficiently the important fetch-and-add synchronization primitive. We outline the hardware that would be required to build a 4096 processor system using 1990's technology. We also discuss system software issues, and present analytic studies of the network performance. Finally, we include a sample of our effort to implement and simulate parallel variants of important scientific programs.

**Index Terms**—Computer architecture, fetch-and-add, MIMD, multiprocessor, Omega-network, parallel computer, parallel processing, shared memory, systolic queues, VLSI.

Manuscript received January 13, 1982; revised June 28, 1982. This work was supported in part by the National Science Foundation under Grant NSF-MCS79-07804, and by the Applied Mathematical Sciences Program of the Department of Energy under Grant DE-AC02-76ER03077.

A. Gottlieb, R. Grishman, K. P. McAuliffe, and M. Snir are with the Courant Institute of Mathematical Sciences, New York University, New York, NY 10012.

C. P. Kruskal is with the Department of Computer Science, University of Illinois, Urbana, IL 61801.

L. Rudolph was with the Department of Computer Science, University of Toronto, Toronto, Ont., Canada. He is now with the Department of Computer Science, Carnegie-Mellon University, Pittsburgh, PA 15213.

## I. INTRODUCTION

WITHIN a few years advanced VLSI (very large scale integration) technology will produce a fast single-chip processor including high-speed floating-point arithmetic. This leads one to contemplate the level of computing power that would be attained if thousands of such processors cooperated effectively on the solution of large-scale computational problems.

The NYU "Ultracomputer" group has been studying how such ensembles can be constructed for effective use and has produced a tentative design that includes some novel hardware and software components. The design may be broadly classified as a general purpose MIMD machine accessing a central shared memory via a message switching network with the geometry of an Omega-network. (For related designs see [3], [37], [38], [42], and [44].)

The major thrust of this paper is to outline and justify, in some detail, the proposed hardware and present the analytic and simulation results upon which parts of the design are based. We also discuss system software issues and describe some of our ongoing efforts to produce parallel versions of important scientific programs (but the reader should see [13]

and [18] respectively for a more detailed treatment of these last two topics). Section II of the present report reviews the idealized computation model upon which our design is based; Section III presents the machine design; Section IV analyzes network performance; Section V highlights a parallel scientific program; and Section VI summarizes our results.

## II. MACHINE MODEL

In this section we first review the paracomputer model, upon which our machine design is based, and the fetch-and-add operation, which we use for interprocessor synchronization. After illustrating the power of this model, we examine some alternative models and justify our selection. Although the paracomputer model to be described is not physically realizable, we shall see in Section III that close approximations can be built.

### A. Paracomputers

An idealized parallel processor, dubbed a "paracomputer" by Schwartz [36] and classified as a WRAM by Borodin and Hopcroft [2], consists of autonomous processing elements (PE's) sharing a central memory. The model permits every PE to read or write a shared memory cell in one cycle. In particular, simultaneous reads and writes directed at the same memory cell are effected in a single cycle.

We augment the paracomputer model with the "fetch-and-add" operation (described below) and make precise the effect of simultaneous access to the shared memory. To accomplish the latter we define the *serialization principle*: The effect of simultaneous actions by the PE's is as if the actions occurred in some (unspecified) serial order. For example, consider the effect of one load and two stores simultaneously directed at the same memory cell. The cell will come to contain some one of the quantities written into it. The load will return either the original value or one of the stored values, possibly different from the value the cell comes to contain. Note that simultaneous memory updates are in fact accomplished in one cycle; the serialization principle speaks only of the effect of simultaneous actions and not of their implementation.

We stress that paracomputers must be regarded as idealized computational models since physical limitations, such as restricted fan-in, prevent their realization. In the next section we review the technique whereby a connection network may be used to construct a parallel processor closely approximating our enhanced paracomputer.

### B. The Fetch-and-Add Operation

We now introduce a simple yet very effective interprocessor synchronization operation, called fetch-and-add, which permits highly concurrent execution of operating system primitives and applications programs. The format of this operation is  $F\&A(V, e)$ , where  $V$  is an integer variable and  $e$  is an integer expression. This indivisible operation is defined to return the (old) value of  $V$  and to replace  $V$  by the sum  $V + e$ . Moreover, fetch-and-add must satisfy the serialization principle stated above: If  $V$  is a shared variable and many fetch-and-add operations simultaneously address  $V$ , the effect of these operations is

exactly what it would be if they occurred in some (unspecified) serial order, i.e.,  $V$  is modified by the appropriate total increment and each operation yields the intermediate value of  $V$  corresponding to its position in this order. The following example illustrates the semantics of fetch-and-add: Assuming  $V$  is a shared variable, if PE $_i$  executes

$$ANS_i \leftarrow F\&A(V, e_i),$$

and if PE $_j$  simultaneously executes

$$ANS_j \leftarrow F\&A(V, e_j),$$

and if  $V$  is not simultaneously updated by yet another processor, then either

$$\begin{array}{l} ANS_i \leftarrow V \\ ANS_j \leftarrow V + e_i \end{array} \quad \text{or} \quad \begin{array}{l} ANS_i \leftarrow V + e_j \\ ANS_j \leftarrow V \end{array}$$

and, in either case, the value of  $V$  becomes  $V + e_i + e_j$ .

For another example consider several PE's concurrently applying fetch-and-add, with an increment of 1, to a shared array index. Each PE obtains an index to a distinct array element (although one does not know beforehand which element will be assigned to which PE). Furthermore, the shared index receives the appropriate total increment.

Section III presents a hardware design that realizes fetch-and-add without significantly increasing the time required to access shared memory and that realizes simultaneous fetch-and-adds updating the same variable in a particularly efficient manner.

### C. The Power of Fetch-and-Add

Since in a parallel processor the relative cost of serial bottlenecks rises with the number of PE's, users of future ultra-large-scale machines will be anxious to avoid the use of critical (and hence necessarily serial) code sections, even if these sections are small enough to be entirely acceptable in current practice.

If the fetch-and-add operation is available, we can perform many important algorithms in a completely parallel manner, i.e., without using any critical sections. For example Gottlieb *et al.* [13]<sup>1</sup> present a completely parallel solution to the readers-writers problem<sup>2</sup> and a highly concurrent queue management technique that can be used to implement a totally decentralized operating system scheduler. (The queue management technique is reprinted in the Appendix of the present paper.) We are unaware of any other completely parallel solutions to these problems.<sup>3</sup> To illustrate the nonserial behavior of these algorithms, we note that given a single queue that is

<sup>1</sup> As explained in [12], the replace-add primitive defined in [13] and used in several of our earlier reports is essentially equivalent to the fetch-and-add primitive used in the present paper.

<sup>2</sup> Since writers are inherently serial, the solution cannot strictly speaking be considered completely parallel. However, the only critical section used is required by the problem specification. In particular, during periods when no writers are active, no serial code is executed.

<sup>3</sup> Reed and Kanodia [31] give a completely parallel solution to those readers-writers problems for which it suffices to detect, rather than prevent, the concurrent execution of multiple reader tasks with a single writer task. (After detecting such an occurrence, the reader tasks are "undone" and "re-done".)

neither empty nor full, the concurrent execution of thousands of inserts and thousands of deletes can all be accomplished in the time required for just one such operation. Other highly parallel fetch-and-add-based algorithms appear in [18], [24], [25], and [33].

#### D. Generalizing Fetch-and-Add

One can define a more general fetch-and- $\phi$  operation that fetches the value in  $V$  and replaces it with  $\phi(V, e)$ . Of course defining  $\phi(a, b) = a + b$  gives fetch-and-add. If  $\phi$  is both associative and commutative, the final value in  $V$  after the completion of concurrent fetch-and- $\phi$ 's is independent of the serialization order chosen.

We now show that two important coordination primitives, swap and test-and-set, may also be obtained as special cases of fetch-and- $\phi$ . (It must be noted, however, that the fetch-and-add operation has proved to be a sufficient coordination primitive for all the highly concurrent algorithms developed to date.) We use the brackets { and } to group statements that must be executed indivisibly and define test-and-set to be a value-returning procedure operating on a shared Boolean variable.

```
TestAndSet(V)
  {Temp ← V
   V ← TRUE}
  RETURN Temp.
```

The swap operation is defined as exchanging the values of a local variable  $L$  (which specifies a processor register or stack location) and a variable  $V$  stored in central memory

```
Swap(L, V)
  {Temp ← L
   L ← V
   V ← Temp}
```

It is easy to see that

TestAndSet( $V$ ) is equivalent to Fetch&OR( $V$ , TRUE).

Similarly, a swap operation can be effected by using the projection operator  $\pi_2$ , where  $\pi_2(a, b) = b$ ; i.e.,

Swap( $L, V$ ) is equivalent to  $L \leftarrow \text{Fetch}\&\pi_2(V, L)$ .

We conclude this discussion of fetch-and- $\phi$  by showing that this operation may be used as the sole primitive for accessing central memory. Specifically, we show how to obtain the familiar load and store operations as degenerate cases of fetch-and- $\phi$ . To load the local variable  $L$  from a variable  $V$  stored in central memory one simply executes

```
 $L \leftarrow \text{Fetch}\&\pi_1(V, *)$ 
```

where  $\pi_1(a, b) = a$  and the value of  $*$  is immaterial (and thus need not be transmitted). Similarly, to store the value of  $L$  into  $V$  one executes

```
 $* \leftarrow \text{Fetch}\&\pi_2(V, L)$ 
```

where the  $*$  indicates that the value returned is not used (and thus again need not be transmitted).

#### E. Alternate Machine Models

In this subsection we discuss several other heavily researched models of parallel processors and explain our choice of a large-scale MIMD shared memory machine.

One line of study pioneered by Kung (see, e.g., [27]), focuses on the great economic and speed advantages obtainable by designing parallel algorithms that conform well to the restrictions imposed by VLSI technology, in particular algorithms and architectures that lay out well in two dimensions. These "systolic" processor designs are already having a significant impact on signal processing, an impact that will doubtless increase dramatically over the next several years. However, for computations having complex control and data flow, the systolic architecture is less well suited. We do expect that VLSI systolic systems will be used for those subcomponents of our machine having regular control and data flow; the design of one such component, an enhanced systolic queue, is presented in Section III-D.

The current generation of supercomputers may be roughly classified as SIMD shared memory machines by considering their vector pipelines to be multiple processors each executing the same instruction [41]. Effective use of such machines is only attained by algorithms consisting primarily of vector operations. Although it is far from trivial to "vectorize" algorithms, such a program has been successfully undertaken at many supercomputer sites. Once again, however, some problems (especially those with many data dependent decisions) appear to resist effective vectorization. Rodrigue, Giroux, and Pratt [32] of Lawrence Livermore National Laboratory write:

Vector and array processors were designed with the idea of solving fluid-type problems efficiently. In general these machines do not lend themselves well to particle tracking calculations. For a scientific laboratory such as LLNL, the computer should be able to handle both forms of calculation, but it remains to be seen whether this goal will ever be achieved.

This goal is achieved by rejecting SIMD machines in favor of the MIMD paracomputer model, which our simulation studies have shown to be effective for both fluid-type [34] and particle tracking calculations [19].

Yet a third alternative model, specifically architectures derived from very general abstract "dataflow" models of parallel computation, have been pursued by other researchers (see the February 1982 Special Issue of *Computer* and the references contained therein). Recent work in this area has stressed the advantages of a purely applicative, side-effect-free programming language for the description of parallel computation. Although such dataflow machines have been discussed for several years, no completely satisfactory physical design has yet emerged. Without commenting on the relative merits of applicative programming, we note that Gottlieb and Schwartz [14] show how a dataflow language may be executed with maximal parallelism on our machine (see also [17]).

The final model we consider is a message passing alternative to shared memory. Except for very small systems, it is not

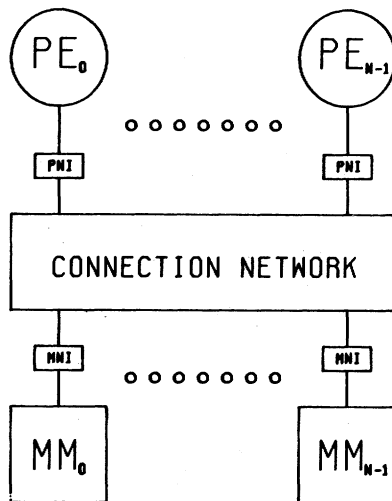


Fig. 1. Block diagram.


possible to have every PE directly connected to every other PE. Thus it may be necessary to route messages via intermediate PE's. In the original ultracomputer design of Schwartz [36] the programmer specified the routing explicitly. By tailoring algorithms to the particular interconnection geometry, one can obtain very high performance. However, we found such a machine to be significantly more difficult to program than one in which the entire memory is available to each PE (see [8], [10], [11], and [35]). If the geometry is hidden from the programmer by having the individual PE's perform the necessary routing, a more loosely coupled machine results. In recent years such machines have been much studied for distributed computing applications. Although message passing architectures are indeed quite attractive for distributed computing, we believe that for the applications we have emphasized, thousands of processors cooperating to solve a single large-scale scientific problem, the more tightly coupled model featuring high speed concurrent access to shared memory is more effective.

### III. MACHINE DESIGN

In this section we present the design of the NYU Ultracomputer, a machine that appears to the user as a paracomputer, and we justify our design decisions. As indicated above, no machine can provide the single-cycle access to shared memory postulated in the paracomputer model; our design approximates a paracomputer by using a message switching network with the geometry of the Omega-network of Lawrie [29]<sup>4</sup> to connect  $N = 2^D$  autonomous PE's to a central shared memory composed of  $N$  memory modules (MM's). Thus, the direct single cycle access to shared memory characteristic of paracomputers is replaced by an indirect access via a multi-cycle connection network. Each PE is attached to the network via a processor network interface (PNI) and each MM is attached via a memory network interface (MNI). Fig. 1 gives a block diagram of the machine.

After reviewing routing in the network, we show that an analogous network composed of enhanced switches provides

<sup>4</sup> Note that this network has the same topology as the rectangular SW banyan network of Goke and Lipovsky [7].

Fig. 2. Omega-network ( $N = 8$ ).

efficient support for concurrent fetch-and-add operations. We then examine our choice of network and local memory. To conclude this section we present a detailed design for the switches and describe the PE's, MM's, and network interfaces. As will be shown both the PE's and MM's are relatively standard components; the novelty of the design lies in the network and in particular in the constituent switches and interfaces.

#### A. Network Design

For machines with thousands of PE's the communication network is likely to be the dominant component with respect to both cost and performance. The design to be presented achieves the following objectives.

- i) Bandwidth linear in  $N$ , the number of PE's.
- ii) Latency, i.e., memory access time, logarithmic in  $N$ .
- iii) Only  $O(N \log N)$  identical components.
- iv) Routing decisions local to each switch; thus routing is not a serial bottleneck and is efficient for short messages.
- v) Concurrent access by multiple PE's to the same memory cell suffers no performance penalty; thus interprocessor coordination is not serialized. We are unaware of any significantly different design that also attains these goals.

1) *Routing in an Omega-Network:* The manner in which an Omega-network can be used to implement memory loads and stores is well known and is based on the existence of a (unique) path connecting each PE-MM pair. To describe the routing algorithm we use the notation in Fig. 2: both the PE's and the MM's are numbered using  $D$ -bit identifiers whose values range from 0 to  $N - 1$ ; the binary representation of each identifier  $x$  is denoted  $x_D \cdots x_1$ ; upper ports on switches are numbered 0 and lower ports 1; messages from PE's to MM's traverse the switches from left to right; and returning messages traverse the switches from right to left. A message is transmitted from PE( $p_D \cdots p_1$ ) to MM( $m_D \cdots m_1$ ) by using output port  $m_j$  when leaving the stage  $j$  switch. Similarly, to travel from MM( $m_D \cdots m_1$ ) to PE( $p_D \cdots p_1$ ) a message uses output port  $p_j$  at a stage  $j$  switch.

The routing algorithm just presented generalizes immediately to a  $D$ -stage network composed of  $k$ -input- $k$ -output switches (instead of the  $2 \times 2$  switches used above) connecting  $k^D$  PE's to  $k^D$  MM's: The ports of a switch are numbered 0 to  $k - 1$  and the identifiers are written in base  $k$ . Although the remainder of this section deals exclusively with  $2 \times 2$  switches, all the results generalize to larger switches, which are considered in Section IV.

2) *Omega-Network Enhancements*: To prevent the network from becoming a bottleneck for machines comprising large numbers of PE's, an important design goal has been to attain a bandwidth proportional to the number of PE's. This has been achieved by a combination of three factors (see Section IV for an analysis of network bandwidth).

i) The network is pipelined, i.e., the delay between messages equals the switch cycle time not the network transit time. (Since the latter grows logarithmically, nonpipelined networks can have bandwidth at most  $O(N/\log N)$ .)

ii) The network is message switched, i.e., the switch settings are not maintained while a reply is awaited. (The alternative, circuit switching, is incompatible with pipelining.)

iii) A queue is associated with each switch to enable concurrent processing of requests for the same port. (The alternative adopted by Burroughs [3] of killing one of the two conflicting requests also limits bandwidth to  $O(N/\log N)$ , see [26].)

Since we propose using a message switching network, it may appear that both the destination and return addresses must be transmitted with each message. We need, however, transmit only one  $D$  bit address, an amalgam of the origin and destination: When a message first enters the network, its origin is determined by the input port, so only the destination address is needed. Switches at the  $j$ th stage route messages based on memory address bit  $m_j$  and then replace this bit with the PE number bit  $p_j$ , which equals the number of the input port on which the message arrived. Thus, when the message reaches its destination, the return address is available.

When concurrent loads and stores are directed at the same memory location and meet at a switch, they can be combined without introducing any delay by using the following procedure (see [13] and [21]).

i) Load-Load: Forward one of the two (identical) loads and satisfy each by returning the value obtained from memory.

ii) Load-Store: Forward the store and return its value to satisfy the load.

iii) Store-Store: Forward either store and ignore the other.

Combining requests reduces communication traffic and thus decreases the lengths of the queues mentioned above, leading to lower network latency (i.e., reduced memory access time). Since combined requests can themselves be combined, the network satisfies the key property that any number of concurrent memory references to the same location can be satisfied in the time required for just one central memory access. It is this property, when extended to include fetch-and-add operations as indicated below, that permits the bottleneck-free implementation of many coordination protocols.

3) *Implementing Fetch-and-Add*: By including adders in




Fig. 3. Combining Fetch-and-Adds.

the MNI's, the fetch-and-add operation can be easily implemented: When  $F \& A(X, e)$  is transmitted through the network and reaches the MNI associated with the MM containing  $X$ , the value of  $X$  and the transmitted  $e$  are brought to the MNI adder, the sum is stored in  $X$ , and the old value of  $X$  is returned through the network to the requesting PE. Since fetch-and-add is our sole synchronization primitive (and is also a key ingredient in many algorithms), concurrent fetch-and-add operations will often be directed at the same location. Thus, as indicated above, it is crucial in a design supporting large numbers of processors not to serialize this activity.

Enhanced switches permit the network to combine fetch-and-adds with the same efficiency as it combines loads and stores: When two fetch-and-adds referencing the same shared variable, say  $F \& A(X, e)$  and  $F \& A(X, f)$ , meet at a switch, the switch forms the sum  $e + f$ , transmits the combined request  $F \& A(X, e + f)$ , and stores the value  $e$  in its local memory (see Fig. 3). When the value  $Y$  is returned to the switch in response to  $F \& A(X, e + f)$ , the switch transmits  $Y$  to satisfy the original request  $F \& A(X, e)$  and transmits  $Y + e$  to satisfy the original request  $F \& A(X, f)$ .

Assuming that the combined request was not further combined with yet another request, we would have  $Y = X$ ; thus the values returned by the switch are  $X$  and  $X + e$ , thereby effecting the serialization order " $F \& A(X, e)$  followed immediately by  $F \& A(X, f)$ ." The memory location  $X$  is also properly incremented, becoming  $X + e + f$ . If other fetch-and-add operations updating  $X$  are encountered, the combined requests are themselves combined, and the associativity of addition guarantees that the procedure gives a result consistent with the serialization principle.

Although the preceding description assumed that the requests to be combined arrive at a switch simultaneously, the actual design can also merge an incoming request with requests already queued for output to the next stage (see Section III-D).

To combine a fetch-and-add operation with another reference to the same memory location we proceed as follows.

i) Fetch&Add-Fetch&Add: As described above, a combined request is transmitted and the result is used to satisfy both fetch-and-adds.

ii) Fetch&Add-Load: Treat Load( $X$ ) as Fetch&Add( $X, 0$ ).

iii) Fetch&Add( $X, e$ )-Store( $X, f$ ): Transmit Store( $e + f$ ) and satisfy the fetch-and-add by returning  $f$ .

Finally, we note that a straightforward generalization of the above design yields a network implementing the fetch-and- $\phi$  primitive for any associative operator  $\phi$  [12].

4) *Other Considerations*: We now turn our attention to other issues concerning the proposed network design.

Since the introduction of queues in each switch leads to stochastic delays and the network is pipelined, it is possible for memory references from a given PE to distinct MM's to be satisfied in an order different from the order in which they were issued. This reordering can violate the serialization principle specified in our model (see [28]). A simple-minded solution to this problem is not to pipeline requests to read-write shared variables; however, this approach is overly conservative since most such request can be safely pipelined.

Since the analyses thus far obtained require the introduction of simplifying assumptions (see Section IV), and we are unable to perform faithful simulations of full 4096 PE networks, we cannot confidently predict the expected network latency. Our preliminary analyses and partial simulations have yielded encouraging results.

A potential serial bottleneck is the memory module itself. If every PE simultaneously requests a distinct word from the same MM, these  $N$  requests are serviced one at a time. However, introducing a hashing function when translating virtual addresses to physical addresses, assures that unfavorable situations occur with probability approaching zero as  $N$  increases. On average, the most requests directed at a single MM is (asymptotically) on  $\log N / \log \log N$  [44]. Furthermore, the probability that all MM's receive less than  $\log N$  requests (asymptotically) exceeds  $1 - 1/P(N)$  for any polynomial  $P$ .

The hardware complexity due to the decision to adopt a queued message switching network introduces significant processing at each stage. Although the internal cycle time of the switches may be important for today's technology, we expect that by the end of the decade any on-chip delay will be dominated by the chip-to-chip transmission delays. (Since the switch bandwidth will be pin limited, the added internal complexity will not increase the component count.)

### B. Local Memory

The negative impact of the large network latency can be partially mitigated by providing each PE with a local memory in which private variables reside and into which read-only shared data (in particular, program text) may be copied. Storing shared read-write data in the local memory of multiple PE's must, in general, be prohibited: the resulting memory incoherence would otherwise lead to violations of the serialization principle. We shall show in Section III-D that in certain special cases, this restriction may be relaxed.

One common design for parallel machines is to implement a separately addressable local memory at each PE, imposing upon compilers and loaders the onus of managing the two level store. The alternative approach, which we intend to implement, is the one conventionally used on uniprocessors: The local memory is implemented as a cache. Experience with uniprocessor systems shows that a large cache can capture up to 95 percent of the references to cacheable variables, effectively shifting the burden of managing a two level store from the software to the hardware (see [20]).

### C. The Switches

We now detail an individual network switch, which is es-

entially a  $2 \times 2$  bidirectional routing device transmitting a message from its input ports to the appropriate output port on the opposite side. The PE side sends and receives messages to and from the PE's via input ports, called  $\text{FromPE}_i$ , where  $i = 0, 1$ , and output ports, called  $\text{ToPE}_i$ . Similarly, the MM side communicates with the MM's via ports  $\text{FromMM}_i$  and  $\text{ToMM}_i$ . (Note that in our figures the To and From ports are coalesced into bidirectional ports.)

As indicated above, we associate a queue with each output port. The head entry is transmitted when the switch at the adjacent stage is ready to receive it (the message might be delayed if the queue this message is due to enter is already full).

To describe the process whereby requests are combined in a switch, we view a request as consisting of several components: function indicator (i.e., load, store, or fetch-and-add), address, and data. The address itself consists of the amalgamation of part of the PE number and part of the MM number, and the internal address within the specified MM. For ease of exposition, we consider only combining homogeneous requests (i.e., requests with like function fields); it is not hard to extend the design to permit combining heterogeneous requests. For each request,  $R\text{-new}$ , that enters a ToMM queue<sup>5</sup>, we search the requests already in this queue using as key the function, MM number, and internal address from  $R\text{-new}$ .<sup>6</sup> If no request matches  $R\text{-new}$ , then no combining is possible and  $R\text{-new}$  simply remains the tail entry of the output queue. Otherwise, let  $R\text{-old}$  denote the message in the ToMM queue that matches  $R\text{-new}$ . Then, to effect the serialization  $R\text{-old}$  followed immediately by  $R\text{-new}$ , the switch performs the following actions: The addresses of  $R\text{-new}$  and  $R\text{-old}$  are placed into a Wait Buffer (to await the return of  $R\text{-old}$  from memory) and  $R\text{-new}$  is deleted from the ToMM queue. If the request is a store then the datum of  $R\text{-old}$  (in the ToMM queue) is replaced by the datum of  $R\text{-new}$ . If the request is a fetch-and-add then the datum of  $R\text{-old}$  is replaced by the sum of the two data. In addition, for fetch-and-adds, the datum of  $R\text{-old}$  is sent to the wait buffer. Thus, each entry sent to the wait buffer consists of the address of  $R\text{-old}$  (the entry key); the address of  $R\text{-new}$ ; and, in the case of a combined fetch-and-add, a datum. (Note that stores and fetch-and-adds can both be implemented by using an ALU that receives the data of  $R\text{-old}$  and  $R\text{-new}$  and returns either the sum of the two numbers or just  $R\text{-new}$ .)

Before presenting the actions that occur when a request returns to a switch from a MM, we make two remarks. First, we will use two Wait Buffers (one associated with each ToMM queue) since access to a single wait buffer would be rate limiting. Second, the key of each entry is the Wait Buffer uniquely identifies the message for which it is waiting since the PNI is to prohibit a PE from having more than one outstanding reference to the same memory location.

After arriving at a FromMM port, a returning request,  $R\text{-ret}$ , is both routed to the appropriate ToPE queue and used to search associatively the relevant wait buffer. If a match

<sup>5</sup> Although we use the term queue, entries within the middle of the queue may also be accessed.

<sup>6</sup> The design of the ToMM queue, permitting this search and subsequent actions to be performed with minimal delay, is detailed in Section III-C1.

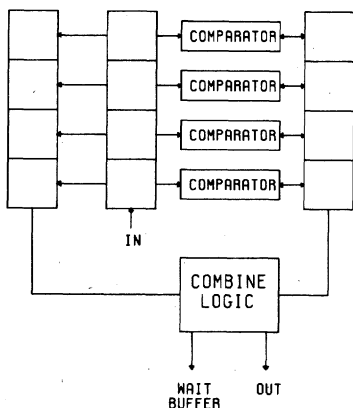


Fig. 4. Systolic ToMM queue.

occurs, the entry found,  $R$ -wait, is removed from the buffer and its function indicator, PE and MM numbers, and address are routed to the appropriate ToPE queue. If the request was a load, the data field is taken from  $R$ -ret; if a fetch-and-add, the  $R$ -wait data field is added to the  $R$ -ret data field.

To summarize the necessary hardware, we note that in addition to adders, registers, and routing logic, each switch requires two instances of each of the following memory units. For each unit we have indicated the operations it must support.

- i) ToMM-queue: Entries are inserted and deleted in a queue-like fashion, associative searches may be performed, and matched entries may be updated.
- ii) ToPE-queue: Entries may be inserted and deleted in a queue-like fashion.
- iii) Wait-buffer: Entries may be inserted and associative searches may be performed with matched entries removed.

Note that it is possible for more than two requests to be combined at a switch. However, the structure of the switch is simplified if it supports only combinations of pairs since a request returning from memory could then match at most one request in the wait buffer, eliminating the need for contention logic. Another advantage of not supporting multiple combinations within one switch is that it permits the pipelined implementation of the ToMM queue described below.

The switch can be partitioned into two essentially independent components, each implementing a unidirectional switch. The communication between the two components is restricted to the information pertaining to combined messages, that is, the information sent from the ToMM queues to the wait buffers. Since requests are combined relatively infrequently, the link between the two components can have a small bandwidth. We are currently investigating other possible partitions for a switch while noting that its increased functionality hinders a bit-slice implementation.

1) *The ToMM Queue*: As illustrated in Fig. 4 our ToMM queue is an enhancement of the VLSI systolic queue of Guibas and Liang [15]. We first describe the queue-like behavior of this structure and then explain how the necessary searching is accomplished.

Items added to the queue enter the middle column, check the adjacent slot in the right column, and move into this slot if it is empty. If the slot is full, the item moves up one position

in the middle column and the process is repeated. (Should the item reach the top of the middle column and still be unable to shift right, the queue is declared full.) Meanwhile, items in the right column shift down, exiting the queue at the bottom.

Before giving the enhancements needed for searching, we make four observations: the entries proceed in a FIFO order; as long as the queue is not empty and the switch in the next stage can receive an item, one item exits the queue at each cycle; as long as the queue is not full a new item can be entered at each cycle<sup>7</sup>; items are not delayed if the queue is empty and the next switch can receive them.

The queue is enhanced by adding comparison logic between adjacent slots in the right two columns, permitting a new entry moving up the middle column to be matched successively against all the previous entries as they move down the right column.<sup>8</sup> If a match is found, the matched entry moves (from the middle column) to the left column, called the "match column." Entries in the match column shift down at the same rate as entries on the right column of the queue. A pair of requests to be combined will therefore exit their respective columns at the same time and will thus enter the combining unit simultaneously.

Note that it is possible to reduce the width of the ToMM queue by having each request split into several successive entries. If requests are transmitted between switches as a series of successive packets, a ToMM queue with a width matching the size of these packets would avoid the assembly and disassembly of messages, resulting in a complete pipelining of the message processing. The smaller size of comparators and adders may also result in faster logic. A detailed description of the VLSI switch logic appears in [40].

#### D. The Network Interfaces

The PNI (processor-network interface) performs four functions: virtual to physical address translation, assembly/disassembly of memory requests, enforcement of the network pipeline policy, and cache management. The MNI (memory-network interface) is much simpler, performing only request assembly/disassembly and the additions operation necessary to support fetch-and-add. Since the MNI operations as well as the first two PNI functions are straightforward, we discuss only pipelining policy and cache management.

Before detailing these two functions, we note two restrictions on pipelining memory request (i.e., issuing a request before the previous one is acknowledged). As indicated above, pipelining requests indiscriminately can violate the serialization principle (Section III-A4), and furthermore, pipelining requests to the same memory location is not supported by our current switch design (Section III-C).

Since accessing central memory involves traversing a mul-

<sup>7</sup> The number of cycles intervening between successive insertions must, however, be even (zero included).

<sup>8</sup> Actually, an item is matched against half of the entries moving down the right column. Since we packetize messages (see below) we can arrange for each item to be matched against its corresponding item in each request moving down the right column (this is particularly easy if each message consists of an even number of packets). If, however, an entire request is contained in one packet, then one needs either twice as many comparators or two cycles for each motion.

tistage network, effective cache management is very important. To reduce network traffic a write-back update policy was chosen: writes to the cache are not written through to central memory; instead, when a cache miss occurs and eviction is necessary, updated words within the evicted block are written to central memory. Note that cache generated traffic can always be pipelined.

In addition to the usual operations described above, which are invisible to the PE, our cache provides two functions, release and flush, that must be specifically requested and can be performed on a segment level or for the entire cache. We now show that judicious use of release and flush further reduces network traffic.

The release command marks a cache entry as available without performing a central memory update. This enables a task to free cache space allocated to virtual addresses that will no longer be referenced. For example, private variables declared within a begin-end block can be released at block exit. Thus, the release operation reduces network traffic by lowering the quantity of data written back to central memory during a task switch. Moreover, if (prior to a task switch) another virtual address maps to a released cache address, no central memory update is necessary.

Release also facilitates caching shared read-write data during periods of read-only access: if a set of tasks sharing read-write data can guarantee that during a period of time no updates will occur, then the data is eligible for caching for the duration of this period. Subsequently, the data must be released and marked uncacheable to insure that no task uses stale data.

The flush facility, which enables the PE to force a write-back of cached values, is needed for task switching since a blocked task may be rescheduled on a different PE. To illustrate another use of flush and release, consider a variable  $V$  that is declared in task  $T$  and is shared with  $T$ 's subtasks. Prior to spawning these subtasks,  $T$  may treat  $V$  as private (and thus eligible to be cached and pipelined) providing that  $V$  is flushed, released, and marked shared immediately before the subtasks are spawned. The flush updates main memory, the release insures that the parent task will not use stale data, and marking  $V$  shared enables  $T$ 's subtasks to reference  $V$ . Once the subtasks have completed,  $T$  may again consider  $V$  as private and eligible for caching. Coherence is maintained since  $V$  is cached only during periods of exclusive use by one task.

To increase performance one may define a combined flush-and-release command that can be implemented using only a single scan through the cache directory.

### E. The Processors and Memory Modules

The MM's are standard components consisting of off the shelf memory chips. The PE's, however, need to be a (slightly) custom design since we require the fetch-and-add operation. Moreover, to utilize fully the high bandwidth connection network, a PE must continue execution of the instruction stream immediately after issuing a request to fetch a value from central memory. The target register would be marked "locked" until the requested value is returned from memory; an attempt to use a locked register would suspend execution.

Note that this policy is currently supported on large scale computers and is becoming available on one chip processors [30]. Software designed for such processors would attempt to prefetch data sufficiently early to permit uninterrupted execution and a cache designed for such processors would not lock up despite outstanding cache misses (see [23] for one such design).

If the latency remains an impediment to performance, we would hardware-multiprogram the PE's as in the CHOPP design [42] and the Denelcor HEP machine [4]. Note that  $k$ -fold multiprogramming is equivalent to using  $k$  times as many PE's—each having relative performance  $1/k$ . Since, to attain a given efficiency, such a configuration requires problems of larger size, we view multiprogramming as a last resort.

Although we have not given sufficient attention to I/O, we note that I/O processors can be substituted for arbitrary PE's in the system. More generally, since the design does not require homogeneous PE's, a variety of special purpose processors (e.g., FFT chips, matrix multipliers, voice generators, etc.) can be attached to the network.

### F. Machine Packaging

We conservatively estimate that a machine built in 1990 would require four chips for each PE-PNI pair, nine chips for each MM-MNI pair (assuming a 1 Mbyte MM built out of 1 Mbit chips), and two chips for each 4-input-4-output switch (which replaces four of the  $2 \times 2$  switches described above). Thus, a 4096 processor machine would require roughly 65 000 chips, not counting the I/O interfaces. Note that the chip count is still dominated, as in present day machines, by the memory chips, and that only 19 percent of the chips are used for the network. Nevertheless, most of the machine volume will be occupied by the network, and its assembly will be the dominant system cost due to the nonlocal wiring required.

It is possible to partition an  $N$  input,  $N$  output Omega network built from  $2 \times 2$  switches into  $\sqrt{N}$  "input modules" and  $\sqrt{N}$  "output modules." An input module consists of  $\sqrt{N}$  network inputs and the  $\sqrt{N} (\lg N)/4$  switches that can be accessed from these inputs in the first  $(\lg N)/2$  stages of the network.<sup>9</sup> An output module consists of  $\sqrt{N}$  network outputs and the  $\sqrt{N} (\lg N)/4$  switches that can be accessed from these outputs in the last half of the network. Moreover, as noted by Wise [45], it is possible to arrange the switches of each module so that, between any two successive stages, all lines have the same length (Fig. 5). Finally, if the input boards are stacked vertically on one rack, the output boards are stacked vertically on another rack, and the two racks are stacked one atop another, such that the boards on one rack are orthogonal to the boards on the other rack, then all off board lines will run nearly vertically between the two sets of boards as illustrated in Fig. 6. The same strategy can be used for networks built of  $k \times k$  switches. (Figs. 5 and 6 are reprinted with permission from [45].)

We propose using this layout for a 4K processor machine

<sup>9</sup> We use  $\lg$  for the base 2 logarithm.






Fig. 5. Layout of network on boards.




Fig. 6. Packaging of network boards.

constructed from the chips described at the beginning of this section. This machine would include two types of boards: "PE boards" that contain the PE's, the PNI's, and the first half of the network stages and "MM boards" that contain the MM's, the MNI's, and the last half of the network stages. Using the chip counts given above, a 4K PE machine built from two chip  $4 \times 4$  switches would need 64 PE boards and 64 MM boards, with each PE board containing 352 chips and each MM board containing 672 chips. Since the PE chips will be near the free edge of the PE board and the MM chips will be near the free edge of the MM board, I/O interfaces can be connected along these edges. Bianchini and Bianchini [1] have considered machines built from single chip  $2 \times 2$  switches. Their analysis shows that an air cooled 4096 PE ultracomputer can easily be packaged into a  $5 \times 5 \times 10$  ft enclosure.

#### IV. COMMUNICATION NETWORK PERFORMANCE

Since the overall ultracomputer performance is critically dependent on the communication network and this network is likely to be the most expensive component of the completed machine, it is essential to evaluate the network performance carefully so as to choose a favorable configuration.

##### A. Performance Analysis

Although each switch in the network requires a significant amount of logic, it appears feasible to implement a  $2 \times 2$  switch on one chip using today's technology. Further, we believe it will be feasible in 1990 technology to implement  $4 \times 4$  or even  $8 \times 8$  switches on one chip. It seems, however, that the main restriction on the switch performance will be the rate at which information can be fed into and carried from the chip, rather than the rate at which that information can be processed within the chip. The basic hardware constraint will therefore be the

number of bits that can be carried on or off the chip in one unit of time (one cycle).

Suppose that 400 bits can be transferred on or off the chip in one cycle (which we estimate, for 1990 technology, to be on the order of 25 ns). If each message transmitted through the network consists of approximately 100 bits (64 bits of data), then a  $2 \times 2$  switch needs two cycles for the transfer of the 800 bits involved in the relaying of two messages in each direction. It is, however, possible to pipeline the transmission of each message, so that the delay at each switch is only one cycle if the queues are empty.

The chip bandwidth constraint does not determine a unique design for the network. It is possible to replace  $2 \times 2$  switches by  $k \times k$  switches, time multiplexing each line by a factor of  $k/2$ . It is also possible to use several copies of the same network, thereby reducing the effective load on each one of them and enhancing network reliability. We present performance analyses of various networks in order to indicate the tradeoffs involved.

A particular configuration is characterized by the values of the following three parameters.

- i)  $k$ —the size of the switch. Recall that a  $k \times k$  switch requires  $4k$  lines.
- ii)  $m$ —the time multiplexing factor, i.e., the number of switch cycles required to input a message. (To simplify the analysis we assume that all the messages have the same length.)
- iii)  $d$ —the number of copies of the network that are used.

The chip bandwidth limits the  $k/m$  ratio, which we may thus take to be independent of the network configuration. Note that for any  $k$  a network with  $n$  inputs and  $n$  outputs can be built from  $(n \lg n)/(k \lg k)$   $k \times k$  switches and a proportional

number of wires. Since our network contains a large number of identical switches, the network's cost is essentially proportional to the number of switches and independent of their complexity. Thus the cost of a configuration is  $d(n \lg n)/(k \lg k)$  (we are neglecting the small cost of interfacing the  $d$  copies of the network).

In order to obtain a tractable mathematical model of the network we have made the following simplifying assumptions.

- 1) Requests are not combined.
- 2) Requests have the same length.
- 3) Queues are of infinite size.
- 4) Requests are generated at each PE by independent, identically distributed, time-invariant random processes.
- 5) MM's are equally likely to be referenced.

Let  $p$  be the average number of messages entered into the network by each PE per network cycle. If the queues at each switch are large enough ("infinite queues") then the average switch delay is approximately

$$1 + \frac{m^2 p (1 - 1/k)}{2(1 - mp)}$$

cycles (see [26]; similar results can be found in [16] and [6]). The average network traversal time (in one direction)  $T$  equals the number of stages times the switch delay plus the setting time for the pipe, i.e.,

$$T = \frac{\lg n}{\lg k} \left( 1 + \frac{m^2 p (1 - 1/k)}{2(1 - mp)} \right) + m - 1.$$

Let us note the following facts.

1) The network has a capacity of  $1/m$  messages per cycle per PE. That is each PE cannot enter messages at a rate higher than one per  $m$  cycles, and conversely the network can accommodate any traffic below this threshold. Thus, the global bandwidth of the network is indeed proportional to the number of PE's connected to it.

2) The initial 1 in the expression for the switch delay corresponds to the time required for a message to be transmitted through a switch without being queued (the switch service time). The second term corresponds to the average queueing delay. This term decreases to zero when the traffic intensity  $p$  decreases to zero and increases to infinity when traffic intensity  $p$  increases to the  $1/m$  threshold. The surprising feature of this formula is the  $m^2$  factor, which is explained by noting that the queueing delay for a switch with a multiplexing factor of  $m$  is roughly the same as the queueing delay for a switch with a multiplexing factor of one, a cycle  $m$  times longer, and  $m$  times as much traffic per cycle.

We now use these formulas to compare the performance of different configurations. Let us assume that using  $k \times k$  switches the time multiplexing factor  $m$  equals  $k$ . Using  $d$  copies of the network reduces the effective load on each copy by a factor of  $d$ . Thus the average transit time for a network consisting of  $d$  Omega-networks composed of  $k \times k$  switches is

$$T = \frac{\lg n}{\lg k} \left( 1 + \frac{k(k-1)p}{2(d-kp)} \right) + k - 1$$

cycles, where  $p$  is, as before, the average number of messages sent to the network by each PE per cycle. As expected, delays decrease when  $d$  increases. The dependency on  $k$  is more subtle. Increasing  $k$  decreases the number of stages in the network, but increases the pipelining factor, and therefore increases the queueing delays and the pipe setting delay.

We have plotted in Fig. 7 the graphs of  $T$  as a function of the traffic intensity,  $p$ , for different values of  $k$  and  $d$ . We see that for reasonable traffic intensities a duplicated network composed of  $4 \times 4$  switches yields the best performance (but see the next paragraph for a preliminary evaluation for networks consisting of two-chip switches). A network with  $8 \times 8$  switches and  $d = 6$  also yields an acceptable performance, at approximately the same cost as the previous network. Since the bandwidth of the first network is  $d/k = 0.5$  and the bandwidth of the second is 0.75, we see that for a given traffic level the second network is less heavily loaded and thus should provide better performance for traffic with high variance. Of course a final determination of an optimal configuration requires more accurate assessments of the technological constraints and the traffic distribution. The pipelining delays incurred for large multiplexing factors, the complexity of large switches, and the heretofore ignored cost and performance penalty incurred with interfacing many network copies, will probably make the use of switches larger than  $8 \times 8$  impractical for a 4K PE parallel machine.

The previous discussion assumed a one chip implementation of each switch. By using the two chip implementation described at the end of Section III-C, one can nearly double the bandwidth of each switch while doubling the chip count. As delays are highly sensitive to the multiplexing factor  $m$ , this implementation would yield a better performance than that obtained by taking two copies of a network built of one chip switches. (It would also have the extra advantage of decreasing the gate count on each chip.) Thus, the ultimate choice may well be one network built of  $4 \times 4$  switches, each switch consisting of two chips.

We now return to the five assumptions listed above. The first two assumptions, that all messages are of equal (maximal) length and traverse the entire network, are clearly conservative. In practice, messages that do not carry data (load requests and store acknowledgements) would be shorter, and merged messages do not each traverse the entire network.

Simulations have shown that queues of modest size ( $\leq 8$ ) give essentially the same performance as infinite queues.

Although the requests generated by PE's cooperating on a single problem are not independent, the presence of a large number of PE's and a number of different problems will tend to smooth the data. On the other hand, even in a large system the pattern of requests by a single PE will be time dependent and further analytic and simulation studies are needed to determine the effect of this deviation from our assumed model.

Finally, by applying a hashing function when translating from virtual to physical addresses, the system can ensure that each MM is equally likely to be referenced.




Fig. 7. Transit times for different configurations.

### B. Network Simulations

Our discussion of the possible configurations for the communication network still lacks two essential ingredients: an assessment of the traffic intensity we expect to encounter in practical applications, and an assessment of the impact of the network delay on the overall performance.

We routinely run parallel scientific programs under a paracomputer simulator (see [9]) to measure the speedup obtained by parallelism and to judge the difficulty involved in creating parallel programs (see Section V). A recent modifi-

cation allows us to simulate an approximation to the proposed network design rather than an ideal paracomputer: Since an accurate simulation would be very expensive, we used instead a multistage queueing system model with stochastic service time at each stage (see [39]), parameterized to correspond to a network with six stages of  $4 \times 4$  switches, connecting 4096 PE's to 4096 MM's. A message was modeled as one packet if it did not contain data and as three packets otherwise. Each queue was limited to fifteen packets and both the PE instruction time and the MM access time were assumed to equal twice

TABLE I  
NETWORK TRAFFIC AND PERFORMANCE

	avg. CM access time	idle cycles	idle cycles per CM load	memory ref. per instr.	shared ref. per instr.
1	8.94	37 percent	5.3	0.21	0.08
2	8.83	39 percent	4.5	0.19	0.08
3	8.81	22 percent	4.9	0.25	0.05
4	8.85	19 percent	3.5	0.24	0.06

the network cycle time. Thus the minimum central memory access time, which consists of the MM access time plus twice the minimum network transit time, equals eight times the PE instruction time.

We have monitored the amount of network traffic generated by several scientific programs under the pessimistic assumption that no shared data is cached and the optimistic assumption that all references to program text and private data are satisfied by the cache. The programs studied were as follows.

- i) A parallel version of part of a NASA weather program (solving a two dimensional PDE), with 16 PE's.
- ii) The same program, with 48 PE's.
- iii) The TRED2 program described in Section V, with 16 PE's.
- iv) A multigrid Poisson PDE solver, with 16 PE's.

Table I summarizes simulations of these four programs. The time unit is the PE instruction time. In these simulations the number of requests to central memory (CM) are comfortably below the maximal number that the network can support and indeed the average access time is close to the minimum. (Since each PE was a CDC 6600-type CPU, most instructions involved register-to-register transfers.) Specifically, only one instruction every five cycles for the first two programs (and one every four for the last two) generated a data memory reference.<sup>10</sup> Moreover only one data memory reference out of 2.6 in the first two programs, and one reference out of five for the last two programs were for shared data. We note that the last two programs were designed to minimize the number of accesses to shared data. As a result the number of idle cycles was significantly higher for the first two programs. Since the code generated by the CDC compiler often prefetched operands from memory, the average number of idle cycles per load from central memory was significantly lower than the central memory access time.

We conclude that were these studies repeated on actual hardware the traffic intensity would be low ( $p < .04$ ), and prefetching would mitigate the problem of large memory latency. The first conclusion, however, must be strongly qualified. The simulator we used is much less sensitive to fluctuations in the network traffic than an actual network would be. Moreover, we have ignored both cache generated traffic and the effect of operating system programs.

## V. SIMULATIONS AND SCIENTIFIC PROGRAMMING

As indicated above we use an instruction level paracomputer simulator to study parallel variants of scientific programs.

<sup>10</sup> Since for the first two programs, the PE's were idle (waiting for a memory reference to be satisfied) approximately 40 of the time, five cycles corresponds to approximately three instructions.

Applications already studied include radiation transport, incompressible fluid flow within an elastic boundary, atmospheric modeling, and Monte Carlo simulation of fluid structure. Current efforts include both extending the simulator to model the connection network more faithfully and running programs under a parallel operating system scheduler.

The goals of our paracomputer simulation studies are, first, to develop methodologies for writing and debugging parallel programs and second, to predict the efficiency that future large scale parallel systems can attain. As an example of the approach taken, and of the results thus far obtained, we report on experiments with a parallelized variant of the program TRED2 (taken from Argonne's EISPACK library), which uses Householder's method to reduce a real symmetric matrix to tridiagonal form (see [22] for details).

An analysis of the parallel variant of this program shows that the time required to reduce an  $N \times N$  matrix using  $P$  processors is well approximated by

$$T(P, N) = aN + dN^3/P + W(P, N)$$

where the first term represents "overhead" instructions that must be executed by all PE's (e.g., loop initializations), the second term represents work that is divided among the PE's, and  $W(P, N)$ , the waiting time, is of order  $\max(N, \sqrt{P})$ . We determined the constants experimentally by simulating TRED2 for several  $(P, N)$  pairs and measuring both the total time  $T$  and the waiting time  $W$ . (Subsequent runs with other  $(P, N)$  pairs have always yielded results within one percent of the predicted value.) Table II summarizes our experimental results and supplies predictions for problems and machines too large to simulate (these values appear with an asterisk). In examining this table, recall that the efficiency of a parallel computation is defined as

$$E(P, N) = T(1, N)/(P * T(P, N)).$$

Although we consider these measured efficiencies encouraging, we note that system performance can probably be improved even more by sharing PE's among multiple tasks. (Currently the simulated PE's perform no useful work while waiting.) If we make the optimistic assumption that all the waiting time can be recovered, the efficiencies rise to the values given in Table III.

## VI. CONCLUSION

Until now the goal of building high performance machines has been achieved at the price of increasingly complex hardware structures, and ever more exotic technology. It is our belief that the NYU Ultracomputer approach offers a simpler alternative, which is better suited to advanced VLSI technology: high performance is obtained by assembling large quantities of identical computing components in a particularly effective manner. The 4096 PE Ultracomputer that we envision has roughly the same component count as found in today's large machines. The number of different component types, however, is much smaller, each component being a sophisticated one chip VLSI system. Such machines would be three orders of magnitude faster and would have a main storage

TABLE II  
MEASURED AND PROJECTED EFFICIENCIES

N \ PE	Reduction of Matrices to Tridiagonal Form				
	16	64	256	1024	4096
16	62 percent	26 percent	7 percent	1 percent*	0 percent*
32	87 percent	60 percent	25 percent	6 percent*	1 percent*
64	96 percent	86 percent	59 percent	27 percent*	7 percent*
128	99 percent*	96 percent*	86 percent*	59 percent*	24 percent*
256	100 percent*	99 percent*	96 percent*	86 percent*	58 percent*
512	100 percent*	100 percent*	99 percent*	96 percent*	85 percent*
1024	100 percent*	100 percent*	100 percent*	99 percent*	96 percent*

TABLE III  
PROJECTED EFFICIENCIES

N \ PE	Reduction of Matrices to Tridiagonal Form (without waiting time)				
	16	64	256	1024	4096
16	71 percent	37 percent	12 percent	3 percent	0 percent
32	90 percent	69 percent	35 percent	12 percent	3 percent
64	97 percent	90 percent	68 percent	35 percent	12 percent
128	99 percent	97 percent	90 percent	68 percent	35 percent
256	100 percent	99 percent	97 percent	90 percent	68 percent
512	100 percent	100 percent	99 percent	97 percent	90 percent
1024	100 percent	100 percent	100 percent	99 percent	97 percent

three orders of magnitude larger than present day machines.

Our simulations indicate that the NYU Ultracomputer would be an extremely powerful computing engine for large numerical applications. The low coordination overhead and the large memory enable us to use efficiently the high degree of parallelism available. The usefulness of these machines is not restricted to numerical computations: they are true general-purpose computers. Finally, our limited programming experience, on an instruction level simulator, indicates that the manual translation of serial codes into parallel Ultracomputer codes is a relatively straightforward task.

To demonstrate further the feasibility of the hardware and software design we intend to construct a 64 PE prototype, which will use commercial microprocessors and memories together with custom-built VLSI components for the network.

## APPENDIX

### Management of Highly Parallel Queues

Since queues are a central data structure for many algorithms, a concurrent queue access method can be an important tool for constructing parallel programs. In analyzing one of their parallel shortest path algorithms, Deo *et al.* [5] dramatize the need for this tool.

However, regardless of the number of processors used, we expect that algorithm PPDM has a constant upper bound on its speedup, because every processor demands private use of the Q.

Refuting this pessimistic conclusion, we show in this appendix that, although at first glance the important problem of queue management may appear to require use of at least a

few inherently serial operations, a queue can be shared among processors without using any code that could create serial bottlenecks. The procedures to be shown maintain the basic first-in first-out property of a queue, whose proper formulation in the assumed environment of large numbers of simultaneous insertions and deletions is as follows: if insertion of a data item  $p$  is completed before insertion of another data item  $q$  is started, then it must not be possible for a deletion yielding  $q$  to complete before a deletion yielding  $p$  has started.

In the algorithm below we represent a queue of length  $Size$  by a public circular array  $Q[0:Size-1]$  with public variables  $I$  and  $D$  pointing to the locations of the next items to be inserted and deleted (these correspond to the rear and front of the queue respectively). Initially  $I = D = 0$  (corresponding to an empty queue).

We maintain two additional counters,  $\#QI$  and  $\#Qu$ , which hold lower and upper bounds respectively for the number of items in the queue, and which never differ by more than the number of active insertions and deletions. Initially  $\#QI = \#Qu = 0$ , indicating no activity and an empty queue. The parameters  $QueueOverflow$  and  $QueueUnderflow$  appearing in the program shown below are flags denoting the exceptional conditions that occur when a processor attempts to insert into a full queue or delete from an empty queue. (Since a queue is considered full when  $\#Qu \geq Size$  and since deletions do not decrement  $\#Qu$  until after they have removed their data, a full queue may actually have cells that could be used by another insertion.) The actions appropriate for the  $QueueOverflow$  and  $QueueUnderflow$  conditions are application dependent: One possibility is simply to retry an offending insert or delete; another possibility is to proceed to some other task.

Critical section-free Insert and Delete programs are given below. The insert operation proceeds as follows: First a test-

increment-retest (TIR) sequence is used to guarantee the existence of space for the insertion, and to increment the upper bound  $\#Qu$ . If the TIR fails, a QueueOverflow occurs. If it succeeds, the expression  $\text{Mod}(\text{Fetch\&Add}(I, 1), \text{Size})$  gives the appropriate location for the insertion, and the insert procedure waits its turn to overwrite this cell (see [13]). Finally, the lower bound  $\#Ql$  is incremented. The delete operation is performed in a symmetrical fashion; the deletion of data can be viewed as the insertion of vacant space.

Procedure Insert (Data,  $Q$ , QueueOverflow)

```

If TIR ( $\#Qu, 1, \text{Size}$ ) Then {
  MyI  $\leftarrow$  Mod(Fetch&Add( $I, 1$ ), Size)
  Wait turn at MyI
  Q[MyI]  $\leftarrow$  Data
  Fetch&Add( $\#Ql, 1$ )
  QueueOverflow  $\leftarrow$  False }

```

```

Else QueueOverflow  $\leftarrow$  True

```

End Procedure

Procedure Delete (Data,  $Q$ , QueueUnderflow)

```

If TDR ( $\#Ql, 1$ ) Then {
  MyD  $\leftarrow$  Mod(Fetch&Add( $D, 1$ ), Size)
  Wait turn at MyD
  Data  $\leftarrow$  Q[MyD]
  Fetch&Add( $\#Qu, -1$ )
  QueueUnderflow  $\leftarrow$  False }

```

```

Else QueueUnderflow  $\leftarrow$  True

```

End Procedure

Boolean Procedure TIR ( $S, \text{Delta}, \text{Bound}$ )

```

If  $S + \text{Delta} \leq \text{Bound}$  Then
  If Fetch&Add( $S, \text{Delta}$ )  $\leq$  Bound Then
    TIR  $\leftarrow$  true
  Else { Fetch&Add( $S, -\text{Delta}$ )
        TIR  $\leftarrow$  false }

```

```

Else TIR  $\leftarrow$  False

```

End Procedure

Boolean Procedure TDR ( $S, \text{Delta}$ )

```

If  $S - \text{Delta} \geq 0$  Then
  If Fetch&Add( $S, -\text{Delta}$ )  $\geq 0$  Then
    TDR  $\leftarrow$  True
  Else { Fetch&Add( $S, \text{Delta}$ )
        TDR  $\leftarrow$  false }

```

```

Else TDR  $\leftarrow$  False

```

End Procedure


Although the initial test in both TIR and TDR may appear to be redundant, a closer inspection shows that their removal permits unacceptable race conditions. This point is also expanded in Gottlieb *et al.* [13] where other fetch-and-add based software primitives are presented as well.

It is important to note that when a queue is neither full nor empty our program allows many insertions and many deletions to proceed completely in parallel with no serial code executed. This should be contrasted with current parallel queue algorithms, which use small critical sections to update the insert and delete pointers.

## REFERENCES

- [1] R. Bianchini and R. Bianchini, Jr., "Wireability of the NYU Ultracomputer," Courant Institute, NYU, NY, Ultracomputer Note 43, 1982.
- [2] A. Borodin and J. E. Hopcroft, "Routing, merging, and sorting on parallel models of computation," in *Proc. 14th Annu. Ass. Comput. Mach. Symp. Theory Computing*, 1982, pp. 338-344.
- [3] Burroughs Corp., "Numerical aerodynamic simulation facility feasibility study," NAS2-9897, Mar. 1979.
- [4] Denelcor Corp., "Heterogeneous element processor principles of operation," Denver, CO, Pub. 9 000 001, 1981.
- [5] N. Deo, C. Y. Pang, and R. E. Lord, "Two parallel algorithms for shortest path problems," in *Int. Conf. Parallel Processing*, 1980, pp. 244-253.
- [6] D. Dias and J. R. Jump, "Analysis and simulation of buffered delta networks," *IEEE Trans. Comput.*, vol. C-30, pp. 273-282, 1981.
- [7] L. R. Goke and G. J. Lipovsky, "Banyan networks for partitioning multiprocessor systems," in *Proc. 1st Annu. Symp. Comput. Arch.*, 1973, pp. 21-28.
- [8] A. Gottlieb, "PLUS—A PL/I based Ultracomputer simulator I," Courant Institute, NYU, NY, Ultracomputer Note 10, 1980.
- [9] ———, "WASHCLOTH—The logical successor to soapsuds," Courant Institute, NYU, NY, Ultracomputer Note 12, 1980.
- [10] ———, "PLUS—A PL/I based Ultracomputer simulator II," Courant Institute, NYU, NY, Ultracomputer Note 14, 1980.
- [11] A. Gottlieb and C. P. Kruskal, "MULT—A multitasking Ultracomputer language with timing, I & II," Courant Institute, NYU, NY, Ultracomputer Note 15, 1980.
- [12] ———, "Coordinating parallel processors: A partial unification," *Comput. Arch. News*, pp. 16-24, Oct. 1981.
- [13] A. Gottlieb, B. Lubachevsky, and L. Rudolph, "Basic techniques for the efficient coordination of very large numbers of cooperating sequential processors," *ACM TOPLAS*, Jan. 1983.
- [14] A. Gottlieb and J. T. Schwartz, "Networks and algorithms for very large scale parallel computations," *Computer*, vol. 15, pp. 27-36, Jan. 1982.
- [15] L. J. Guibas and F. M. Liang, "Systolic stacks, queues, and counters," in *Proc. Conf. Advanced Research VLSI*, Jan. 1982.
- [16] R. G. Jacobsen and D. P. Misunas, "Analysis of structures for packet communications," in *Proc. Int. Conf. Parallel Processing*, 1977.
- [17] S. D. Johnson, "Connection networks for output-driven list multiprocessing," Dept. Comput. Sci., Indiana Univ., Indiana, PA, Tech. Rep. 114, 1981.
- [18] M. Kalos, "Scientific calculations on the Ultracomputer," Courant Institute, NYU, NY, Ultracomputer Note 30, 1981.
- [19] M. Kalos, G. Leshem, and B. D. Lubachevsky, "Molecular simulations of equilibrium properties," Courant Institute, NYU, NY, Ultracomputer Note 27, 1981.
- [20] K. R. Kaplan and R. V. Winder, "Cache-based computer systems," *Computer*, vol. 6, pp. 30-36, 1973.
- [21] D. Klappholz, private communication, 1981.
- [22] D. Korn, "Timing analysis for scientific codes run under WASHCLOTH simulation," Courant Institute, NYU, NY, Ultracomputer Note 24, 1981.
- [23] D. Kroft, "Lockup-free instruction fetch/prefetch cache organization," in *Proc. 8th Annu. Symp. Comput. Arch.*, 1981, pp. 81-88.
- [24] C. P. Kruskal, "Upper and lower bounds on the performance of parallel algorithms," Ph.D. dissertation, Courant Institute, NYU, NY, 1981.
- [25] ———, "Replace-add based paracomputer algorithms," in *Proc. Int. Conf. Parallel Processing*, 1982.
- [26] C. P. Kruskal and M. Snir, "Some results on multistage interconnection networks for multiprocessors," Courant Institute, NYU, NY, Ultracomputer Note 41, 1982; also in *Proc. Princeton Conf. Inform. Sci. Syst.*, 1982.
- [27] H. T. Kung, "The structure of parallel algorithms," in *Advances in Computers*, vol. 19, M. C. Yovits Ed. New York: Academic, 1980, 65-112.
- [28] L. Lamport, "How to make a multiprocessor computer that correctly executes multiprocess programs," *IEEE Trans. Comput.*, vol. C-28, pp. 690-691, 1979.
- [29] D. Lawrie, "Access and alignment of data in an array processor," *IEEE Trans. Comput.*, vol. C-24, pp. 1145-1155, 1975.
- [30] G. Radin, "The 801 minicomputer," in *Symp. Architectural Support Programming Languages and Operating Systems*, 1982, pp. 39-47.

- [31] D. P. Reed and R. K. Kanodia, "Synchronization with eventcounts and sequencers," *Commun. Ass. Comput. Mach.*, vol. 22, pp. 115-123, 1979.
- [32] G. Rodrigue, E. D. Giroux, and M. Pratt, "Perspectives on large-scale scientific computing," *Computer*, vol. 13, pp. 65-80, Oct. 1980.
- [33] L. S. Rudolph, "Software structures for ultraparallel computing," Ph.D. dissertation, Courant Institute, NYU, NY, 1981.
- [34] N. Rushfield, "Atmospheric computations on highly parallel MIMD computers," Courant Institute, NYU, NY, Ultracomputer Note 22, 1981.
- [35] J. T. Schwartz, "Preliminary thoughts on Ultracomputer programming style," Courant Institute, NYU, NY, Ultracomputer Note 3, 1979.
- [36] J. T. Schwartz, "Ultracomputers," *ACM TOPLAS*, pp. 484-521, 1980.
- [37] H. J. Siegel and R. J. McMillen, "Using the augmented data manipulator in PASM," *Computer*, vol. 14, pp. 25-34, 1981.
- [38] B. J. Smith, "A pipelined shared resource MIMD computer," in *Proc. Int. Conf. Parallel Processing*, 1978, pp. 6-8.
- [39] M. Snir, "NETSIM—Network Simulator for the Ultracomputer," Courant Institute, NYU, NY, Ultracomputer Note 28, 1981.
- [40] M. Snir and J. Solworth, "The ultraswitch—A VLSI network node for parallel processing," Courant Institute, NYU, NY, Ultracomputer Note 39, 1982.
- [41] H. S. Stone, "Parallel computers," in *Introduction to Computer Architecture*, H. S. Stone, Ed. Chicago, IL: SRA, 1980, pp. 318-347.
- [42] H. Sullivan, T. Bashkow, and D. Klappholz, "A large scale homogeneous, fully distributed parallel machine," in *Proc. 4th Annu. Symp. Comput. Arch.*, 1977, pp. 105-124.
- [43] R. J. Swan, S. H. Fuller, and D. P. Siewiorek, "Cm\*—A modular, multi-microprocessor," in *Proc. AFIPS Conf.*, vol. 46, 1977, pp. 637-644.
- [44] D. S. Wise, "Compact layout of Banyan/FFT networks," in *Proc. CMU Conf. VLSI Systems and Computations*, Kung, Sproull, and Steele, Eds. Rockville, MD: Computer Science Press, 1981, pp. 186-195.
- [45] G. H. Gonet, "Expected length of the longest probe sequence in hash code searching," *J. Ass. Comput. Mach.*, pp. 289-304, Apr. 1981.



**Ralph Grishman** was born in New York, NY, on January 6, 1948. He received the A.B. and Ph.D. degrees in physics from Columbia University, New York, in 1968 and 1973, respectively.

Since 1973 he has been on the faculty at New York University, NY, where he is currently an Associate Professor in the Department of Computer Science. His main research area is computational linguistics, but he also has a long-standing interest in computer system design and construction. Before becoming involved with the Ultracomputer group, he designed a medium-scale scientific computer which is now in production use at New York University.

Dr. Grishman is a member of the Association for Computational Linguistics and the Association for Computing Machinery.



**Clyde P. Kruskal** was born on May 25, 1954. He received the A.B. degree in mathematics and computer science from Brandeis University, Waltham, MA, in 1976, and the M.S. and Ph.D. degrees in computer sciences from New York University in 1978 and 1981, respectively.

Currently, he is a Visiting Research Assistant Professor in the Department of Computer Science, University of Illinois, Urbana. His research interests include the analysis of sequential and parallel algorithms, and the design of parallel computers.




**Kevin P. McAuliffe** was born in Newark, NJ on March 16, 1956. He received the B.A. and M.S. degrees in computer science from New York University, NY, in 1978 and 1980 respectively.

He is currently pursuing the Ph.D. degree in computer science at New York University, where he also lectures in the Department of Computer Science. His current interests are computer architecture and operating systems.



**Larry Rudolph** received the B.S. degree from Queens College, City University of New York, NY, in 1976 and the M.S. and Ph.D. degrees from the Courant Institute, New York University, New York, in 1978 and 1981, respectively.

He spent the 1981-82 academic year as a Post-doctoral Fellow in the Department of Computer Science, University of Toronto, Ont., Canada. Currently, he is a Research Associate at Carnegie-Mellon University, Pittsburgh, PA. His research interests include parallel/distributed computing, robust parallel approximation schemes, and computational geometry.



**Marc Snir** was born in Paris, France, on October 10, 1948. He received the B.Sc. degree and the Ph.D. degree in mathematics from the Hebrew University of Jerusalem, Israel, in 1972 and 1979, respectively.

In 1979 he was Research Fellow in the Department of Computer Science, University of Edinburgh, Scotland. Since 1980 he has been an Assistant Professor in the Department of Computer Science, New York University, where he is actively involved in the Ultracomputer project. His

research interests include parallel processing, interconnection networks, complexity of parallel algorithms, inductive logic, and VLSI design.



**Allan Gottlieb** was born in New York, NY, on August 2, 1945. He received the B.S. degree from Massachusetts Institute of Technology, Cambridge, MA, in 1967, and the M.A. and Ph.D. degrees in mathematics from Brandeis University, Waltham, MA, in 1968 and 1973, respectively.

During 1971-1972 he was an Acting Instructor at the University of California, Santa Cruz, and during 1972-1973 he was an Instructor at the State College of Massachusetts, North Adams. From 1973-1979 he was with York College, City

University of New York, first as an Assistant Professor and then as an Associate Professor. Since 1979 he has been working on the NYU Ultracomputer project at the Courant Institute, NYU, where he is now a Research Associate Professor. His academic interests include parallel processors and parallel algorithms, algorithmic analysis, operating systems, differential topology, and computer chess.

Dr. Gottlieb is a member of the Association for Computing Machinery, the American Mathematical Society, the New York Academy of Sciences, and the Society of Sigma Xi, and is an affiliate member of the IEEE Computer Society.