Request Combining in Multiprocessors with Arbitrary Interconnection Networks

Alvin R. Lebeck and Gurindar S. Sohi

Abstract— Several techniques have been proposed to allow parallel access to a shard memory location by combining requests. They have one or more of the following attributes: requirements for *a priori* knowledge of the request to combine, restrictions on the routing of messages in the network, or the use of sophisticated interconnection network nodes. We present a new method of combining requests that does not have the above requirements. We obtain this new method for request combining by developing a classification scheme for the existing methods of request combining. This classification scheme is facilitated by separating the request combining process into a two part operation: (1) determining the *combining set*, which is the set of requests that participate in a combined access; and (2) distributing the results of the combined access to the members of the combining set. The classification of combining strategies is based upon which system component, processor elements, or interconnection network performs each of these tasks. Our approach, which uses the interconnection network to establish the combining set and the processor elements to distribute the results, lies in an unexplored area of the design space. We also present simulation results to assess the benefits of the proposed approach.

Index Terms— Request combining, interconnection networks, hot spots, parallel prefix

I. INTRODUCTION

ARVIND and Iannucci [2] state that the design of a large-scale, share memory multiprocessor must address two basic issues.

1) It must tolerate long latencies for memory requests.
2) It must achieve unconstrained, yet synchronized, access to shared data.

Although several techniques—for example, caches and prefetching [24] and low-level context switching [25]—have been proposed to tolerate the latency of memory requests, heretofore the only known methods of allowing unconstrained yet synchronized access to shared data are implementations of *request combining*. The earliest published proposal for request combining was in the CHoPP system [28], where several read requests to a common memory location are combined in the interconnection network and are satisfied with only a single access of the memory location. When two read requests destined for the same memory location meet at a node of the network, the source of one of the requests is saved, and only one read request is forwarded to memory. When the response

This figure is adapted from [6].

When two *Fetch* & Φ requests combine, state is saved in a wait buffer until the response returns. This type of combining requires comparators, wait buffers, and an ALU in each network node.

Fig. 1. Request combining in the ultracomputer.

of the read request arrives at the node where combining took place, two responses are sent back toward the processors.

The idea of combining read requests, or *read combining*, in CHoPP was extended in the New York University (NYU) Ultracomputer to allow several types of requests to combine [6]. The Ultracomputer uses the *Fetch&*Φ primitive, where Φ is any associative and commutative operator. An enhanced interconnection network with the topology of an Omega network is proposed to perform combining on the *Fetch&*Φ primitive.

The Ultracomputer style of request combining is illustrated in Fig. 1. When a *Fetch&*Φ(X, e) request meets a *Fetch&*Φ(X, j) request at a network node, *combining* takes place. e is saved in a *wait buffer*, an ALU in the node computes $e\Phi j$, and the request *Fetch&*Φ$(X, e\Phi j)$ is forwarded to memory. When a response V is received from memory for the *Fetch&*Φ$(X, e\Phi j)$ request, *decombining* takes place. V is forwarded as a response to the *Fetch&*Φ(X, e) request, $e\Phi V$ is forwarded as a response to the *Fetch&*Φ(X, j) request, and e is removed from the wait buffer. Correct operation is guaranteed if the combining of requests satisfies the serialization principle: The final state of the system must be consistent with the servicing of all requests in some (unspecified) serial order [6].

There are three distinct features of the Ultracomputer style of request combining.

1) Requests are combined on the forward trip through the network.
2) State is saved in the network when requests are combined.
3) Requests are decombined on the return trip through the network.

Combining in a network node first requires a comparison to determine if two requests are combinable. The two requests must be *Fetch&*Φ requests to the same memory location. This requires the use of comparators in the network nodes. When two requests are combined, each request is removed

from the queue, the value of e is saved in the wait buffer, the operation $e\Phi j$ is carried out in the ALU, and the request *Fetch&Φ($X, e\Phi j$)* is placed in the queue to forward to memory. The wait buffer must be large enough to hold as many values as there are requests that can be combined at this node, else combining will not take place.

On the return trip, each returning request searches the wait buffer, and, if decombining must take place, appropriate actions are initiated. This implies that the return path must be identical to the forward path for decombining to take place, or the return path must have at least one node in common with the forward path: the node where the combining state is stored. Almasi and Gottlieb [1] give several examples of how such *hardware combining* can eliminate serial bottlenecks.

Several alternative proposals for request combining have appeared in the literature [8], [11], [19], [21], [29], [32]. The primary focus of these efforts is on reducing the cost of the combining network. This is accomplished either by altering the topology of the combining network or by requiring the system software to reduce the amount of contention for shared data.

This paper has two purposes. The first is to develop a taxonomy that can be used to categorize combining methods proposed to date. This allows us to enumerate the issues involved, and to make a comparison of known techniques for combining requests. The second is to propose a new approach to request combining, one that can be used with arbitrary interconnection topologies.

We develop a taxonomy in Section II, and classify existing methods using this taxonomy. We observe that one area of the design space, which we call interconnect-processor combining, has not been explored for arbitrary interconnection networks. We explore this in Section III, where we present a new scheme for request combining. The potential of the new combining scheme is evaluated in Section IV. Section V summarizes the paper and suggests directions for further research.

II. A TAXONOMY FOR REQUEST COMBINING

A. Parallel Prefix Computation and Request Combining

Kruskal, Rudolph, and Snir [14] observed that request combining is very similar to the problem of *parallel prefix computation* [15]. Given the elements $x_1, x_2, \cdots, x_n$, a prefix computation produces the following results:

$$\left\{ \begin{array}{ll} r_1 & = x_1 \\ r_n & = r_{n-1}\Phi x_n \end{array} \right\},$$

where Φ is any associative operator. Computing the results in parallel is termed a parallel prefix computation [15].

To examine the similarity between request combining and parallel prefix, consider an example in which four processors add a constant, C, to a shared variable X and receive the previous value of X. Assume that the processors simultaneously execute the atomic operation *Fetch&Add*(X, C). The values returned to the processors are X, $X+C$, $X+2C$, and $X+3C$. Regardless of the order in which requests are serviced, memory has the final value of $X+4C$. This is simply the set of results $\{r_1, r_2, r_3, r_4, r_5\}$ produced by a prefix computation on the set of elements $\{X, C, C, C, C\}$ with the addition operator (+).

Based on this observation, we see that arbitrary request combining is a two-part operation. The first part, or task of the combining method, is to determine the set of requests that are destined for the same memory location and need to be combined. We call this set of requests the *combining set*. The second task is to distribute the results of the combined access to the appropriate processors by performing a prefix computation on the combining set.

A prefix computation network, such as the one proposed for *scan* primitives by Blelloch [3], can be used to distribute the results of the combined access. In such a network, state is saved on the forward trip through the prefix computation network, and the results are distributed on the return trip, very similarly to the Ultracomputer approach toward combining. However, the use of a prefix computation network requires *a priori* knowledge of the combining set. Blelloch proposed the *scan* primitives for a single-instruction multiple-data (SIMD) paradigm where the elements on which to perform the prefix computation, in our case the combining set, are stored in an array. The array is distributed across the processors and participation in the prefix computation is based on the processors' *activation* status. Therefore, to use a prefix computation network to distribute the results of a combined access, the combining set must be established prior to insertion in the prefix computation network. Alternatively, in the Ultracomputer approach toward combining, the combining set is determined dynamically in the interconnect by comparing the addresses of requests on the forward trip through the network. The results are then distributed to the appropriate processors on the return trip through the network.

The two techniques use different system components to establish the combining set. The Ultracomputer uses the interconnect, whereas Blelloch uses *a priori* knowledge in the processor elements. We can therefore obtain a taxonomy for request combining by specifying which system component, processor elements, or interconnect performs each of the tasks involved in combining requests.

B. Classification of Existing Request Combining Strategies

Based upon how the combining set is determined and how the results of the prefix computation are computed and distributed, the design space for request combining can be divided into four regions (Fig. 2). Interconnect-interconnect combining (IIC) covers schemes in which the interconnection network determines the combining set and distributes the results. In processor-interconnect combining (PIC), the processors establish the combining set, and the interconnection network distributes the results. Schemes in which the processors perform both tasks are classified as processor-processor combining (PPC). Finally, interconnect-processor combining (IPC) indicates that the interconnection network determines the combining set, and the processors distribute the results. In the following subsections, we discuss existing methods of request combining according to the region of the design space to which they belong.

Interconnect-Interconnect Combining (IIC): The CHoPP [28] and the NYU Ultracomputer [6] methods of request

Result Computation and Distribution	Determines Combining Set	
	Processors	Interconnect
Processors	Processor-Processor Combining (PPC)	Processor-Interconnect Combining (PIC)
Interconnect	Interconnect-Processor Combining (IPC)	Interconnect-Interconnect Combining (IIC)

Request combining is a two-part operation: establishing a combining set and distributing the results by performing a prefix computation. The techniques for combining requests are classified based on which system component, processors or interconnect, performs each of the tasks.

Fig. 2. Request combining design space.

TABLE I
REQUEST COMBINING COMPARISON

Request Combining Issues				
Combining Classification	Determining Combining Set		Result Computation and Distribution	
	a priori Knowledge	Restricted Routes	Interconnect Sophistication	Latency
IIC	No	Yes	High	O(L)
PIC	Yes	Yes	High	O(L)
PPC	Yes	No	NA	O(HL)
IPC	No	No	Low	O(LlogS)

NA = Not Applicable
H = Height of Software Combining Tree
S = Size of Combining Set
L = Latency of Interconnection Network

The existing techniques of request combining do not cover the entire design space and require either *a priori* knowledge or restrict the routes of messages. Techniques that use the interconnection network to distribute results require a high degree of sophistication in the network nodes. Although there are currently no implementations, IPC does not have the above limitations, but has the potential drawback of increased latency to complete the combining operation.

■ Read & Write Connection
+ Read Connection

Fig. 3. Priority chain implementation.

combining are instances of IIC. The interconnection network determines the combining set and distributes the results. The IBM RP3 [21] researchers proposed the basic ideas of the Ultracomputer method of combining for their implementation. However, the RP3 has two interconnection networks, one network that combines requests and one that services noncombining requests. A distinction is made between noncombinable and potentially combinable requests (typically synchronization requests), and the interconnect dynamically determines the combining set of the potentially combinable requests.

Two alternative techniques for IIC are presented by Tzeng [29] and Hsu and Yew [11]. Tzeng separates the interconnect into a routing section and a combining section. It is assumed that requests that may combine are distinguished from noncombining requests by examination of the opcode. Such requests are directed to the combining section of the network. The combining section of the network then determines the combining set. Hsu and Yew propose a single-stage shuffle-exchange combining network in addition to a noncombining network. These two proposals are similar to the IBM RP3 method of combining, but the difference in topology of the combining network reduces the hardware cost of the networks. However, the basic technique of request combining in all three schemes, the IBM RP3, Hsu and Yew, and Tzeng, is the same as the technique used in the Ultracomputer.
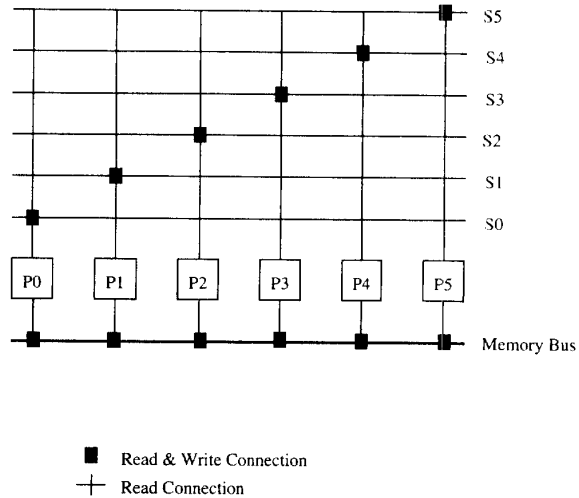
An interesting example of IIC read combining can be found in schemes with hierarchical cache/memory structures, such as cache-only memory architecture (COMA) machines [7], [27], or nonuniform memory access (NUMA) machines with hierarchical caches [18], [23], [30]. Here read combining can be implemented by using a technique similar to the CHoPP method of read combining. A read miss of a cache block at one level of the hierarchy causes a request to be propagated to the next-higher level in the hierarchy. Subsequent read misses of the same cache block at the same level in the hierarchy cause state to be saved. This state allows the data to be forwarded to the appropriate requesting processors when the response arrives from the higher level of the hierarchy.

Processor-Interconnect Combining (PIC): An alternative approach to reducing the cost of the combining network is to determine the combining set prior to inserting the elements into the network. This is the approach taken by schemes that fall under the classification of PIC in the design space.

As with schemes that perform IIC, schemes performing PIC use the interconnection network to distribute the results of combined requests. However, the processor elements, not the interconnection network, determine the combining set. Blelloch's prefix computation network [3], discussed in Section II-A, and the control network of the Thinking Machines CM-5 [17] fall into this category.

Address	
Value	
Head	Tail

Fig. 4. *Fetch* & Φ message format.

```
if (m1.address == m2.address && m1.type == Fetch& Φ
    && m2.type == Fetch& Φ)
    link_msg.dest := m1.tail
    link_msg.head := m2.head
    link_msg.type := LINK
    m1.tail := m2.tail
    m1.value := m1.value Φ m2.value   /* if ALUs in nodes */
    SEND_FORWARD(m1);
    SEND_REVERSE(link_msg);
    DELETE(m2)
fi
```

When two combinable requests arrive at a node of the interconnect, the messages are merged together and the respective combining sets are linked together.

Fig. 5. Merging two messages at an interconnect node.

Another form of PIC in an SIMD paradigm is proposed by Lipovski and Vaughan [19]. This implementation uses a modified carry-lookahead circuit to implement a prefix computation network that distributes the results. The combining set is determined by which processing elements are currently active. The prefix computation network may be extended for operation in a multiple-instruction multiple-data (MIMD) paradigm, though the authors do not explicitly state how this might be done.

An alternative technique for combining requests in a MIMD paradigm, proposed by Harrison [8], uses a synchronous prefix computation network. All requests at the same stage in the network combine. Therefore, the entire combining set must be inserted into the network in the same time slot. This is accomplished by broadcasting information about the combinable locations to the processors. Based on this information, each processor determines the correct time slot for its request.

In addition to being restricted to use *a priori* knowledge of the combining set, the methods discussed in this section make use of a (parallel) prefix computation network to distribute results. Such networks are conceptually very similar to the Ultracomputer combining network: State must be saved on the forward trip through the network, and requests are decombined on the return trip. Recall that this requires sophisticated interconnect nodes and restricts the routes of the return messages. These potentially undesirable features (associated with the result distribution) can be eliminated if we use the processor elements, to perform the (parallel) prefix operation and distribute the results.

Processor-Processor Combining (PPC): In schemes classified as IIC or PIC, the interconnect, as viewed from a share memory location, forms a tree. The memory module is the root of the tree and the processors are the leaves of the tree. The nodes of a combining tree are realized by the implicit storage in the interconnect nodes, i.e., wait buffers. Another

TABLE II
SIMULATION PARAMETERS

Simulation Parameters and Values	
Parameter	Value
message length	1
uniform request rate (h)	0.2 - 1.00
hot request rate (r)	0.0 - 0.32
outstanding uniform requests	unlimited
outstanding hot requests	1
switch buffers	8
processor buffers	4
memory buffers	4
memory cycle time	1
total # of requests	1,000,000
# of processors	256
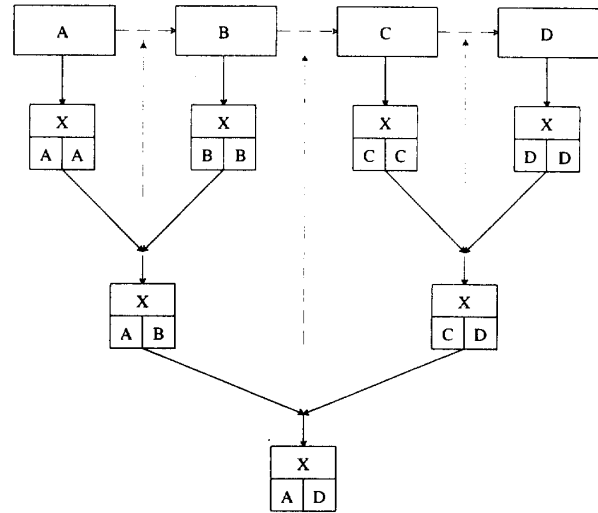
Combining requests involves merging the two messages into one message, which is forwarded to memory. Also, the two lists are linked together by sending a message back to the processors.

Fig. 6. Establishing the combining set as a linked list.

```
for all k in parallel do
    forward[k] := initial[k]
    while (forward[k] != null) do
        val[forward[k]] := val[k] + val[forward[k]]
        forward[k] := forward[forward[k]]
    endwhile
endfor
```

This algorithm, adapted from [10], can be used to distribute the results to the members of the combining set.

Fig. 7. All partial sums of a linked list.

alternative is to use explicit storage in memory to construct the combining tree. This method of request combining, called *software combining* in the literature [5], [31], [32], is classified as PPC, because the processors bear full responsibility for the combining of requests. The processors establish the combining set and distribute the results, and there are no demands on the network at all.

In software combining, one shared location is divided into L locations that constitute the storage for the nodes of the combining tree. Requests are combined as the processors traverse the combining tree. The result is that $\frac{S}{L}$ processors access each of the L locations, rather than S processors
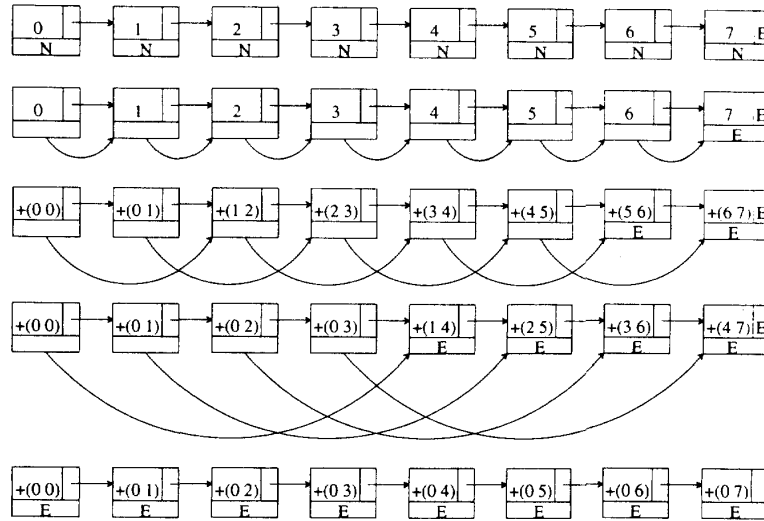
This figure is adapted from Fig. 5 in [10]. It shows parallel partial prefix sum computation for a list with eight members. The values E and N indicate End and NULL, respectively.

Fig. 8. Partial prefix sum computation.

accessing the single location. However, the L locations (nodes of the combining tree) must be distributed across the memory modules in order to alleviate excessive contention for a single memory module.

Yew, Tzeng, and Lawrie show how software combining can be used for barrier operations [32]. Goodman, Vernon, and Woest [5], and Johnson [12], extend the work of Yew, Tzeng, and Lawrie to carry out arbitrary $Fetch\&\Phi$ operations with a software combining tree. Tang and Yew also provide several algorithms for traversing a combining tree where the type of memory access determines which algorithm is chosen (e.g., barrier synchronization, semaphore, or read combining) [31].

A consequence of implementing the combining tree with explicit memory locations is flexibility in the type of memory access. In addition to variable types of memory accesses, software combining permits the use of networks with arbitrary topologies and relatively unsophisticated nodes.

A requirement of implementing a software combining tree for access to a shared location is that the shared location must be known prior to program execution. Furthermore, if the latency of the combined access is to be minimized, the combining tree must be balanced; this requires a priori knowledge of the number of requests that may combine [31]. Moreover, because the combining tree is created based on the maximum number of requests that may combine, the latency to complete the combining operation is influenced by this maximum number. If only one request is accessing the shared location, it must traverse the entire combining tree. For example, in a balanced (software combining) tree of height H, the single request must perform H memory accesses, each of which must traverse the interconnection network.

If the processors are responsible for establishing the combining set, they require a priori knowledge. If the burden of determining the combining set is placed back on the interconnect, the need for a priori knowledge is eliminated.

```
while chum ≠ null do
    send(chum,myinc_val)
    receive(rev,Φ_val)
    value := Φ_val + value
    myinc_val := Φ_val + myinc_val
    send(rev,chum)
    receive(chum,newchum)
    chum := newchum
endwhile
```

The SIMD version for computing the all partial sums of a linked list can be written in an MIMD fashion that uses only point-to-point synchronization [16].

Fig. 9. MIMD distribution of results.

Interconnect-Processor Combining (IPC): Heretofore there are no proposed methods (with the exception of some special cases that we discuss in Section III) of request combining that use an arbitrary interconnection network to determine the combining set and the processor elements to distribute the results. In order to decide if such a scheme is worthy of investigation, the next section compares the issues of implementing combining under each of the classifications.

C. Issues in Combining Requests

There are several aspects to request combining that we have touched upon in the previous discussion that we reiterate to motivate IPC. They are:

* a priori knowledge of the combining set,
* restrictions placed on the routes of messages,
* sophistication of the interconnection network, and
* latency of the combining operation.

The need for a priori knowledge of the combining set requires the programmer to specify this information. Restrictions placed on the routing of messages limits the choices for the interconnection network topology. Sophistication of the interconnection network impacts both the cost and performance of the system. A high degree of sophistication may increase the design time and may either increase the latency

Fig. 10. (a) All requests.

for noncombining requests or require the addition of a second network. The latency of the combining operation is the time from when a processor generates a combinable request to the time when the result of the request is received.

In the next two sections, we look at how each of these issues is affected by which system component establishes the combining set and which component distributes the results. Table I summarizes the following discussion.

Determining the Combining Set: The processor elements require *a priori* knowledge of the combinable locations in order to establish the combining set. For example, the nodes of a software combining tree [31] are defined during algorithm design. In contrast, the interconnect determines the combining set dynamically by comparing destination addresses of messages. The consequence of introducing comparators is a small increase in the sophistication of the nodes of the interconnect, which may result in a slight increase in the latency to complete the combining operation.

Result Computation and Distribution: Placing the responsibility of result distribution on the interconnect has two disadvantages. The first is an increase in the sophistication of the interconnect as a result of wait buffers, and decombining logic in each interconnect node. The second, and perhaps more important, disadvantage is that the route that a return message may travel is restricted, because it must visit the nodes where state was saved on the forward trip through the network.

The primary advantage of using the processor elements to distribute the results of the combining operation is that no restrictions are placed on the routes that messages may travel (no requirement to visit a particular node). Another advantage is that the sophistication of the interconnect nodes is not increased (no need for wait buffers). However, the sophistication of the processor elements (more accurately, the processor-network interface) may increase somewhat to handle the protocol needed to distribute results.

The latency of the combining operation, measured in the number of steps needed to carry out the operation, is a disadvantage of using the processors to distribute results. In Section III, we show that the latency of distributing the results is logarithmic with respect to the number of requests in the combining set, assuming the system does not have broadcast capability.

Based on the above discussion, we feel further investigation of IPC, the as yet unexplored area of the design space, is worthwhile. Such schemes would use the interconnect to determine the combining set and use the processor elements to distribute the results. The above discussion also points out the following potential advantages of IPC combining.

1) No *a priori* knowledge of the combining set is required, because the interconnect dynamically determines the combining set.

(b)

Fig. 10. (b) Hot requests. Average latency vs. maximum bandwidth: No combining.

2) No restrictions on the routes of messages are required, because the processor elements distribute the results.

3) The nodes of the interconnection network require only a small amount of sophistication.

The potential drawback of such a scheme, as compared to IIC, is the latency of the combining operation.

III. INTERCONNECT-PROCESSOR COMBINING

We now consider implementations of request combining that fall into the unexplored region of the design space: interconnect-processor combining (IPC). We initially consider two flavors of combinable operations: a restricted form of Fetch&Add (F&A), or Fetch&Increment [9], [26], and the general F&A operation. In Fetch&Increment, or simply F&I, all participants add the same constant value. In the general F&A, each participant could be adding a different value.

The rationale for the simpler F&I operation is the following. If, in the process of determining the combining set, it is also possible for a participant to determine its overall position in the combining set, i.e., its position in the serial order, then each participant can compute its value without further interactions with other participants. The following example illustrates this point.

Fig. 3 shows a system with six processors (P0–P5) connected to a shared (synchronization) bus. Each processor is assigned one channel in the bus. (The channel could be a wire in an electronic bus [26] or a specific frequency in an optical bus [9].) A given processor can read all channels, but can write to only its own channel.

A processor generates a combinable request and broadcasts its intentions on the bus by putting a 1 on its channel. All other processors that wish to participate in the access write a 1 on their respective channels. At this point, all processors, by monitoring all the channels on the bus, can determine the number and the identity of the processors that are going to participate in the combined access. The combining set is established by determining the participating processors. The ordering in the combining set is statically defined by priorities assigned to the channels.

Suppose four processors $\{P0, P1, P3, P5\}$ would like to perform a F&I operation (increment by a constant C) on memory location X. At the point when the four processors have indicated their intention to access X, the priority chain $\{S0, S1, S2, S3, S4, S5\}$ is $\{110101\}$. Each of the four processors then determines how many processors ahead of it in the priority chain are also participating in the combining operation. For example, P5 determines that there are three processors ahead of it in the priority chain. P3 sees two, P1 sees one, and P0 sees zero, because it is the highest-priority processor participating in the combined access. P0 takes responsibility

(a)

Average latency and maximum bandwidth when combining is enabled and the memory location is locked while the results are distributed.

Fig. 11. (a) All requests.

for accessing X from memory. When $P0$ accesses X, the other processors also read the value from the memory bus. Each processor is restricted to adding the same constant, C, to the shared location; therefore, each processor may compute its value locally based on the number of participants preceding it in the priority chain. $P0$ receives X, $P1$ computes $X + C$, $P3$ computes $X + 2C$, $P5$ computes $X + 3C$, and memory receives $X + 4C$. One processor (or the memory controller who could also be monitoring the bus) takes responsibility for computing $X + 4C$ and updating memory.

The method described above was proposed independently for an electrical bus by Sohi, Smith, and Goodman [26], and for an optical bus by Hiedelberger, Rathi, and Stone [9]. The ease with which combining can be carried out (in special cases), prompted Freudenthal and Gottlieb to investigate the use of the *Fetch&Increment* operation in place of the more general $F\&A$ operation [4].

When broadcast is not an option, some other method must be used to determine the combining set and to carry out the prefix operation on the combining set. Unlike networks with broadcast, there is no easy way to merge the creation of the combining set and its ordering in an arbitrary interconnection network. Here we must continue to separate the creation of the combining set and the implementation of the prefix operation. Because a prefix operation has to be carried out

on the combining set, regardless of whether it is to order the elements (as would be necessary in case of $F\&I$ operations) or to distribute the results (as would be necessary in case of $F\&A$ operation), we see no potential implementation advantage of a $F\&I$ operation over a $F\&A$ operation in an arbitrary network. Therefore, we can continue our further discussion with $F\&A$ operations.

Our proposal for IPC in an arbitrary network uses the network to create the combining set; we choose to represent the combining set as a linked list, though it is conceivable that other structures could also be used. The processors then use this structure to interact with each other and to carry out the prefix operation. We expand on each of the two functions, setting up the list and distributing results, given below.

When a processor submits a $Fetch\&\Phi(X, V_i)$ request, it generates a message that is sent over the interconnect to a destination containing X. The message consists of at least four fields: an address (X), a value field (V), and two pointers, as shown in Fig. 4. The pointers indicate the head and the tail of a list of processors, i.e., a combining set, accessing the memory location specified in the message. Initially the processor generating the request is the only member of the combining set.

When two messages destined for the same memory location meet at a node of the interconnect, the two messages are

(b)

Average latency and maximum bandwidth when combining is enabled and the memory location is locked while the results are distributed.

Fig. 11.   (b) Hot requests. Average latency vs. maximum bandwidth: Lock memory.

combined, as illustrated in Fig. 5. As a result of the combining, a forward message is sent on to the destination location X, and the head and tail fields are updated to indicate the head and tail of the new combining set. The new combining set is the union of the combining sets of the two combined messages. (If ALU's are present, then the value field of the forward message is updated to reflect the Φ operation on the value fields of the combined messages.) A link message is sent to the processor at the tail of the combining set of the first message, with instructions to create a link to the processor at the head of the combining set of the second message. (The route taken by the link message depends upon the topology of the network. For example, in an Omega network with unidirectional links, the link message would travel in the forward direction to some arbitrary destination, and reflect off that destination to go back to the appropriate processor.) Details of the combining operation at each node are shown in Fig. 5. It is important to note that *no state is saved in the network at the point where messages are combined.*

Fig. 6 illustrates the above with an example. Assume that processors A and B generate $F\&A$ requests to location X, as shown in Fig. 6. When the two requests are merged, a message is sent to processor A indicating that processor B is the head of a list of processors that are also participating in the $F\&A$ operation. At this point, processor B is the only member of the

list being added. Processor A now has a pointer to processor B; this forms the new linked list representing a combining set for location X. The fields of the forward message are the memory location (X), the head of the list (A), and the tail of the list (B). (The values fields have not been shown in the figure to prevent a clutter.)

If we assume that processors C and D also formed a combining set for location X, then we have two lists with two members each and two messages en route to memory. When these two messages merge, a message is sent to processor B, indicating that processor C is the head of a linked list of processors also participating in the $F\&A$ operation. This creates a single combining set that is the union of the two original combining sets. The message that is forwarded to memory has its fields set to the head of the first list (A) and the tail of the second list (D). We now have a single combining set represented by one message going to memory. When memory receives the request, it returns the value stored at location X to the head of the list (A) and also sends a message to the tail of the list, (D), indicating that it is responsible for providing the final result to memory. This allows the results of the $F\&A$ operations to be distributed to the processors, as we shall discuss shortly.

When a $F\&A$ request reaches the destination X, the head and tail fields of the request point to the head and the tail of

Fig. 12. Waiting time at memory.

the combining set. We have several options on how to proceed. If ALU's were present in the network nodes where combining took place, and had been used to update the value field of the forward message (as shown in Fig. 5), then the value field of the message reaching X contains the (final) result of the prefix operation applied to all the value fields of the members of the combining set (determined by the head and tail pointers), even though the individual processor (or intermediate) results of the prefix operation are not yet known. The old value of memory location X can be returned to the head, and used by the prefix operation (described later) to determine the results for the members of the combining set. Memory location X could be updated with the final value (old value of X plus the value field of the message reaching X). Further accesses to the same location, X, can proceed *while* the results of the prefix operation on the (first) combining set are being carried out.

If ALU's were not present in the network nodes where combining took place, then the final value of the prefix operation on the combining set cannot be determined when the (combined) message reaches memory location X. We must wait until the prefix operation on the combining set is complete before memory can receive the up-to-date value. What happens to further requests to location X during that time? There are two options. The first option is to simply lock the memory

location X until the final result of the prefix operation is known (the tail of the combining set is responsible for unlocking X). This makes further requests to X wait, building a second combining set while waiting for the prefix operation on the first to complete. Locking X is fairly easy to do, but as we shall see in Section IV, it has implications for performance.

The second option is to append the new combining set onto the old combining set (on which a prefix operation is in progress), thereby creating one combining set, and make the prefix operation robust enough so that it can operate on variable-size combining sets. When appending a new combining set onto an old combining set, care must be taken to avoid potential race conditions. In particular, the message from the memory to the tail of the old combining set (telling it to create a pointer to the head of the new combining set, thereby continuing the prefix operation) could arrive *after* the prefix operation on the old combining set has been completed, and after the (final) result is on its way to memory. The message must be reflected back to memory, which can then supply the up-to-date value to the new combining set, and allow it to start its prefix operation. The message handling protocol must take the possibility of such race conditions into account, and take appropriate action to prevent incorrect operation. Although at first glance appending onto a combining set on which a prefix

(a)

Fig. 13(a) shows the average latency versus maximum bandwidth when the subsequent combining sets are appended to the existing set.

Fig. 13. (a) All requests.

operation is in progress appears to be a good idea, it also has implications for performance, as we shall see in Section IV.

From the above discussion, we can discern the amount of sophistication required in each of the interconnect nodes. The use of comparators is necessary to determine if two messages are destined for the same memory location. A small amount of additional logic is required to construct the new message that is sent back to the processors and to update the tail of the message forwarded to memory. ALU's are needed if we expect to update memory with the final value immediately.

Now we consider how a prefix operation could be carried out on the combining set, which is represented as a linked list. For the purposes of our discussion on result distribution, assume that processors can use only point-to-point messages, reflected off the memory modules if necessary, for communication. The naive method of distributing the results is to start at the head of the list and sequentially move from one node to the next in the list. Although there is no hot spot when using the naive method to distribute results, the messages are unnecessarily serialized. To eliminate this serialization for

the distribution of results, we turn to the literature on parallel applications to see how the prefix operation on the combining set can be carried out in parallel.

Several algorithms for performing a parallel prefix computation on a linked list exist in the literature [10], [13], [15], [20]. Most of the algorithms are concerned with the case of having more nodes in the linked list than processors available. In our case, the number of nodes in the linked list is equivalent to the number of processors that are participating in the combined access. Therefore, we use the all partial sums algorithm given by Hillis and Steele [10] and shown in Fig. 7. The algorithm uses recursive doubling: Each iteration of the loop performs half as many operations as the previous iteration until the entire computation is complete. This algorithm has the advantage that the result is computed in $O(\log S)$ steps, where S is the number of nodes in the list, i.e., the cardinality of the combining set. An example of a partial prefix sum computation for a list with eight nodes is shown in Fig. 8. The array *initial*[] is indexed by processor number and contains the number of the next processor in the initial linked list, whereas the array

(b)

As shown in Fig. 13(b), the average delay of hot requests is quite large; hence, the average latency of all requests increases.

Fig. 13. (b) Hot requests. Average latency when combining sets are appended.

forward[] indicates the processor that is communicated with during the current iteration of the algorithm.

It is important to note that the recursive doubling algorithm is defined for a SIMD machine, and therefore appropriate synchronization must be added for MIMD operation. Fig. 9 shows the MIMD version for result distribution in IPC. (The reader is referred to [16] for further discussion of the transformation from SIMD to MIMD.) Processors also require a limited amount of memory $(O(\log S))$ for storing the pointers to neighboring processors in the combining set.

We have not tried to exhaust the methods of IPC, nor have we addressed all of the issues involved in IPC. There are undoubtedly alternative data structures for maintaining the combining set and associated algorithms for distributing the results. However, we do provide a viable solution that deserves further investigation.

IV. EVALUATION OF IPC

To investigate the overall system performance using IPC, we developed a simulator of a multiprocessor system that performs IPC. An enhanced Omega network establishes the combining sets as described in Section III, and the processors use the MIMD version of the recursive doubling parallel prefix algorithm to distribute the results of the combined access.

The simulator consists of three distinct parts: processor elements, interconnection network, and memory modules. Table II summarizes the parameters of our system. The processor elements are memory reference generators with the additional code required to distribute the results of the combined access. Processors generate a memory reference each cycle with probability r, provided that the network can accept the request. Of these memory references, $h\%$ are directed at a single hot memory location [22], [32]. Each processor may have only one outstanding $F\&A$ request, but unlimited outstanding uniform requests. However, a processor does not generate any requests when there is an outstanding $F\&A$. Yew *et al.* [32] call this the *limited-variable* access pattern.

The interconnection network used in our simulations is an enhanced Omega network. Each node has *bidirectional*

(a)

Fig. 14(a) shows the average latency versus maximum bandwidth when the nodes of the network have ALU's.

Fig. 14.   (a) All requests.

links,[1] and a queue exists between the forward network and the reverse network. Therefore, each switch in the forward network has two inputs and three outputs, whereas each switch in the reverse network has three inputs and two outputs. We assume that each message is a single packet, and that no buffering occurs if there is no contention in the network. Two messages combine only if one is buffered, and we assume a full comparator is used to determine if an arriving message is combinable with any buffered message. Only pairwise combining is carried out at each node.

We simulated a system with 256 processors connected to 256 memory modules. In all of our simulations, we vary h from 0% to 32% and r from 20% to 100%. Fig. 10 shows the average latency and the maximum bandwidth when no combining is performed. As previously shown in [22] and [32],

there is a point of saturation after which bandwidth ceases to increase and latency increases.

In our first experiment, we implemented IPC, without ALU's in the network nodes, and locking the memory location until the prefix operation on the current combining set is complete. Fig. 11(a) presents the latencies of all requests, and saturation bandwidths, for varying hot rates. In comparing Figs. 10 and 11 (note the different scales), we see that the latencies have decreased for small values of h, but increased for large values, and that the saturation bandwidth has also decreased for all values of h. The reason for this is twofold. Since we lock the memory location, we prevent the second combining set from starting its prefix operation. Consequently, the hot requests in the second combining set are stalled for a long time waiting for the prefix operation on the first combining set to complete. This is illustrated in Fig. 12. Fig. 11(b) presents the latency of hot requests, and Fig. 12 presents the average amount of time a hot request spends waiting for the lock on memory to be released. Comparing Figs. 11(b) and 12, we can see that the time waiting for the

(b)

The ALU's eliminate the need to lock memory for result distribution, greatly reducing the average latency of hot requests, as shown in Fig. 14(b).

Fig. 14. (b) Hot requests. Average latency with ALU's in interconnect nodes.

lock to be released can be a significant portion of the overall hot request latency.

The next experiment that we carried out involved appending requests as they arrive at memory onto the end of a combining set on which a prefix operation is already in progress. In this case, requests do not have to wait for the prefix operation on the previous combining set to complete; rather, they join in to form a larger combining set, and join in the existing prefix operation. Our expectation was that by reducing the waiting time, we would decrease the hot request latency. Unfortunately, our experimental results, presented in Fig. 13, indicate otherwise. Fig. 13(a) shows the saturation bandwidth and latency of all requests, and Fig. 13(b) shows the latency of hot requests, for varying hot rates. The reason for the disappointing results for this experiment is that while appending requests to a combining set decreases the time spent by the requests waiting at memory, it actually increases the time to carry out the prefix operation, because the hot requests can get appended one-by-one, and the prefix operation cannot complete as long as the requests are being appended. (No participant has its

value ready until the entire prefix operation is complete.) An extreme case is when the hot requests are appended one-by-one, and the parallel prefix degenerates to serial prefix, but one in which no processor can use its value until the entire prefix operation is complete. Moreover, because there is no deterministic pattern with which the requests are appended, there is no uniform trend in the results. Based upon the results of this experiment, we feel that appending new arrivals onto an existing combining set, on which a prefix operation is already in progress, is something that should be done with caution. More study is needed in this area.

Our final experiment with IPC uses ALU's in the network nodes. In this case, because the final result of the prefix operation is available before its intermediate results are, there is no need to make succeeding hot requests wait while a prefix operation is in progress on a previous combining set. Prefix operations could be in operation on multiple combining sets simultaneously. Fig. 14(a) presents the latencies and saturation bandwidths of all requests, and Fig. 14(b) presents the same for hot requests for varying hot rates. When we compare Figs.

10 and 14, we see that IPC is quite effective in reducing the degradation due to hot spots. IPC has reduced the latency for all requests, improved overall network bandwidth, and reduced the latency for hot requests. Interestingly, the latency for hot requests is lower the higher the hot rate (Fig. 14(b)). This is because with a higher hot rate, a larger combining set is established, and therefore more of the hot requests can be serviced in parallel. Overall, the proposed method for IPC is not as effective as the Ultracomputer style of combining (equivalent results for the Ultracomputer-style of combining can be found in the paper by Pfister and Norton [22]); however, our results suggest that it is an option worth considering.

V. SUMMARY AND CONCLUSION

Unconstrained yet synchronized access to shared memory locations is achieved by combining requests. We formulated a taxonomy for the various techniques of request combining by separating the combining operation into two parts: establishing a combining set and distributing the results. This taxonomy divides the request combining design space into four regions defined by which system component (processor elements or interconnection network) performs each of the tasks of request combining. The classification of the existing implementations of request combining has enabled us to obtain four primary issues that an implementation of request combining must address: the need for *a priori* knowledge of the requests to combine, the complexity of the interconnect nodes, restrictions placed on the routing of messages in the interconnect, and the latency to complete the combined access. We have shown that the current methods of request combining occupy only three of the four regions of the design space and do not address the first three issues satisfactorily. We presented an implementation that lies in the as yet unexplored region, known as interconnect-processor combining (IPC), in which the interconnect establishes the combining set and the processor elements distribute the results. We show that implementations in this area of the design space have the following advantages:

- no need for *a priori* knowledge of the combinable location,
- no restrictions on the routing of messages, and
- low sophistication of interconnect nodes.

We also carried out several experiments to assess the benefits for IPC and observed that IPC is an effective technique by which to alleviate the degradation caused by serial access of memory, but it must be used with care. Some implementations of IPC could degrade the latency of hot requests intolerably, thereby reducing network bandwidth, even though they reduce the latency of uniform requests.

Although we feel that the flexibility and the performance benefits for IPC make it an attractive design choice, more research is needed in this area before a definitive answer can be obtained. Many more alternatives of IPC need to be investigated, including the choice of the data structure used to maintain the combining set, as well as the algorithm used to carry out the prefix operation on the combining set. (One

step in this direction is a recent thesis by Johnson where he investigates the use of IPC to build a tree to implement a scalable cache coherence scheme [12].) More direct comparisons between the different forms of combining, using real application workloads, and different network topologies, also need to be done so that we can get a better picture of the cost-performance benefits of the various techniques for request combining.

REFERENCES

[1] G. S. Almasi and A. Gottlieb, *Highly Parallel Computing*. Redwood City, CA: Benjamin/Cummings, 1989.

[2] K. Arvind and R. A. Iannucci, "A critique of multiprocessing von neumann style," in *Proc. 10th Int. Symp. Comput. Architecture*, 1983, pp. 426–436.

[3] G. E. Blelloch, "Scans as primitive parallel operations," *IEEE Trans. Comput.*, vol. C-38, pp. 1526–1538, Nov. 1989.

[4] E. Freudenthal and A. Gottlieb, "Process coordination with fetch-and-increment," in *Proc. ASPLOS-IV*, 1991, pp. 260–268.

[5] J. R. Goodman, M. K. Vernon, and P. J. Woest, "A set of efficient synchronization primitives for a large-scale shared-memory multiprocessor," in *Proc. ASPLOS-III*, 1989, pp. 64–73.

[6] A. Gottlieb, R. Grishman, C. P. Kruskal, K. P. McAuliffe, L. Rudolph, and M. Snir, "The NYU ultracomputer: Designing a MIMD, shared memory parallel machine," *IEEE Trans. Comput.*, vol. C-32, pp. 175–189, Feb. 1983.

[7] E. Hagersten, A. Landin, and S. Haridi, "DDM: A cache-only memory architecture," *IEEE Comput.*, pp. 44–54, Sept. 1992.

[8] M. C. Harrison, "Synchronous combining of fetch-and-add operation," *Ultracomputer Note 71*, Apr. 1984.

[9] P. Heidelberger, B. D. Rathi, and H. S. Stone, "A device for performing efficient task distribution with a bus connection," IBM Res., Tech. Disclosure YO889-0053, vol. 20, Jan. 1989.

[10] W. D. Hillis and G. L. Steele, "Data parallel algorithms," *Commun. ACM*, pp. 1170–1183, Dec. 1986.

[11] W. T. Hsu and P.-C. Yew, "An effective synchronization network for hot-spot accesses," *ACM Trans. Comput. Syst.*, vol. 10, pp. 167–189, Aug. 1992.

[12] R. E. Johnson, "Extending the scalable coherent interface for large-scale shared-memory multiprocessors," Ph.D. dissertation, Tech. Rep. 1136, Dept. Comput. Sci., University of Wisconsin—Madison, USA, Feb. 1993.

[13] P. M. Kogge and H. S. Stone, "A parallel algorithm for the efficient solution of a general class of recurrence equations," *IEEE Trans. Comput.*, vol. C-22, pp. 786–793, Aug. 1973.

[14] C. P. Kruskal, L. Rudolph, and M. Snir, "Efficient synchronization on multiprocessors with shared memory," *ACM Trans. Programming Languages and Syst.*, vol. 10, pp. 579–601, Oct. 1988.

[15] R. E. Ladner and M. J. Fischer, "Parallel prefix computation," *J. ACM*, vol. 27, pp. 831–838, Oct. 1980.

[16] A. R. Lebeck, "Request combining in multiprocessors with arbitrary interconnection networks," M.S. thesis, Dept. of Comput. Sci., Univ. of Wisconsin—Madison, USA, 1991.

[17] C. E. Leiserson, "The network architecture of the connection machine CM-5," *Proc. ACM Symp. Parallel Algorithms and Architectures*, July 1992.

[18] D. Lenoski, J. Laudon, K. Gharachorloo, A. Gupta, and J. Hennessy, "The directory-based cache coherence protocol for the DASH multiprocessor," *Proc. 17th Ann. Symp. Comput. Architecture*, May 1990.

[19] G. J. Lipovski and P. Vaughan, "A fetch-and-op implementation for parallel computers," in *Proc. 15th Ann. Symp. Comput. Architecture*, 1988, pp. 384–392.

[20] B. D. Lubachevsky and A. G. Greenberg, "Simple, efficient asynchronous parallel prefix algorithms," *Proc. 1987 Int. Conf. Parallel Processing*, 1987, pp. 66–69.

[21] G. F. Pfister, W. C. Brantley, D. A. George, S. L. Harvey, W. J. Kleinfelder, K. P. McAuliffe, E. A. Melton, V. A. Norton, and J. Weiss, "The IBM Research Parallel Processor Prototype (RP3): Introduction and architecture," *Proc. 1985 Int. Conf. Parallel Processing*, 1985, pp. 764–771.

[22] G. F. Pfister and V. A. Norton, "Hot-spot contention and combining in multistage interconnection networks," *IEEE Trans. Comput.*, vol. C-34, pp. 943–948, Oct. 1985.

[23] S. L. Scott, "Toward the design of large-scale shared-memory multiprocessors," Tech. Rep. 1100, Ph.D. dissertation, Dept. of Comput. Sci., Univ. of Wisconsin—Madison, WI, USA, July 1992.

[24] A. J. Smith, "Cache memories," *ACM Computing Surv.* , vol. 14, pp. 473–530, Sept. 1982.

[25] B. Smith, "Architecture and applications of the HEP multiprocessor computer system," *Proc. Int. Soc. Opt. Eng.*, 1992, pp. 241–248.

[26] G. S. Sohi, J. E. Smith, and J. R. Goodman, "Restricted *Fetch&Φ* operations for parallel processing," in *Proc. 3rd Int. Conf. Supercomputing* , 1989, pp. 410–416.

[27] P. Stenstrom, J. Truman, and A. Gupta, "Comparative performance evaluation of cache-coherent NUMA and COMA architectures," in *Proc. 19th Ann. Symp. Comput. Architecture*, 1992, pp. 80–91.

[28] H. Sullivan and L. Cohn, "Shared memory computer method and apparatus," U.S. Patent 4 707 781, Nov. 1987.

[29] N. F. Tzeng, "Design of a novel combining structure for shared-memory multiprocessors," *Proc. 1989 Int. Conf. Parallel Processing*, 1989, pp. I-1–I-8.

[30] A. W. Wilson, "Hierarchical cache/bus architecture for shared memory multiprocessors," in *Proc. 14th Ann. Symp. Comput. Architecture*, 1987, pp. 244–252.

[31] P.-C. Yew and P. Tang, "Software combining algorithms for distributing hot-spot addressing," *J. Parallel Distrib. Computing*, Oct. 1990.

[32] P.-C. Yew, N.-F. Tzeng, and D. H. Lawrie, "Distributing hot-spot addressing in large scale multiprocessors," *IEEE Trans. Comput.*, vol. C-36, pp. 388–395, Apr. 1987.

A. R. Lebeck received the B.S. degree in electrical engineering in 1989, and the M.S. degree in computer science in 1991, from the University of Wisconsin—Madison, USA.

He is currently a Ph.D. candidate in the Department of Computer Sciences, University of Wisconsin—Madison, USA. His research interests include memory system design for uniprocessors and multiprocessors, efficient simulation of memory systems, and synchronization techniques for multiprocessors.

Mr. Lebeck is a member of ACM.

G. S. Sohi received the B.E. degree in electrical engineering from the Birla Institute of Technology and Science, Pilani, India, in 1981, and the M.S. and Ph.D. degrees in electrical engineering from the University of Illinois—Urbana-Champaign, USA, in 1983 and 1985, respectively. Since September 1985, he has been with the Computer Sciences Department, University of Wisconsin—Madison, USA, where he is currently an Associate Professor.