

ARITHMETIC FUNCTIONS IN TORUS AND TREE NETWORKS**CROSS-REFERENCE TO RELATED APPLICATIONS**

5 The present invention claims the benefit of commonly-owned, co-pending United States Provisional Patent Application Serial Number 60/271,124 filed February 24, 2001 entitled MASSIVELY PARALLEL SUPERCOMPUTER, the whole contents and disclosure of which is expressly incorporated by reference herein as if fully set

10 forth herein. This patent application is additionally related to the following commonly-owned, co-pending United States Patent Applications filed on even date herewith, the entire contents and disclosure of each of which is expressly incorporated by reference herein as if fully set forth herein. U.S. patent application Serial No. (YOR920020027US1, YOR920020044US1 (15270)), for "Class

15 Networking Routing"; U.S. patent application Serial No. (YOR920020028US1 (15271)), for "A Global Tree Network for Computing Structures"; U.S. patent application Serial No. (YOR920020029US1 (15272)), for 'Global Interrupt and Barrier Networks'; U.S. patent application Serial No. (YOR920020030US1 (15273)), for 'Optimized Scalable Network Switch'; U.S. patent application Serial

20 No. (YOR920020031US1, YOR920020032US1 (15258)), for "Arithmetic Functions in Torus and Tree Networks"; U.S. patent application Serial No. (YOR920020033US1, YOR920020034US1 (15259)), for 'Data Capture Technique for High Speed Signaling'; U.S. patent application Serial No. (YOR920020035US1 (15260)), for 'Managing Coherence Via Put/Get Windows'; U.S. patent application

25 Serial No. (YOR920020036US1, YOR920020037US1 (15261)), for "Low Latency Memory Access And Synchronization"; U.S. patent application Serial No. (YOR920020038US1 (15276), for 'Twin-Tailed Fail-Over for Fileservers Maintaining Full Performance in the Presence of Failure'; U.S. patent application Serial No. (YOR920020039US1 (15277)), for "Fault Isolation Through No-

30 Overhead Link Level Checksums"; U.S. patent application Serial No. (YOR920020040US1 (15278)), for "Ethernet Addressing Via Physical Location for Massively Parallel Systems"; U.S. patent application Serial No. (YOR920020041US1 (15274)), for "Fault Tolerance in a Supercomputer Through Dynamic Repartitioning"; U.S. patent application Serial No. (YOR920020042US1 (15279)), for "Checkpointing Filesystem"; U.S. patent application Serial No.

35 YOR920020031US1

(YOR920020043US1 (15262)), for "Efficient Implementation of Multidimensional Fast Fourier Transform on a Distributed-Memory Parallel Multi-Node Computer"; U.S. patent application Serial No. (YOR9-20010211US2 (15275)), for "A Novel Massively Parallel Supercomputer"; and U.S. patent application Serial No.
5 (YOR920020045US1 (15263)), for "Smart Fan Modules and System".

2. BACKGROUND ART

Provisional patent application no. 60/271,124, titled "A Novel Massively Parallel SuperComputer" describes a computer comprised of many computing nodes and a
10 smaller number of I/O nodes. These nodes are connected by several networks. In particular, these nodes are interconnected by both a torus network and by a dual functional tree network. This torus network may be used in a number of ways to improve the efficiency of the computer.

15 To elaborate, on a machine which has a large enough number of nodes and with a network that has the connectivity of an M-dimensional torus, the usual way to do a global operation is by the means of shift and operate. For example, to do a global sum (MPI_SUM) over all nodes, after each computer node has done its own local partial sum, each node first sends the local sum to its plus neighbor along one
20 dimension and then adds the number it itself received from its neighbor to its own sum. Second, it passes the number it received from its minus neighbor to its plus neighbor, and again adds the number it receives to its own sum. Repeating the second step (N-1) times (where N is the number of nodes along this one dimension) followed by repeating the whole sequence over all dimensions one at a time, yields
25 the desired results on all nodes. However, for floating point numbers, because the order of the floating point sums performed at each node is different, each node will end up with a slightly different result because of roundoff effects due to the fact that the order of the floating point sums performed at each node is different. This will cause a problem if some global decision is to be made which depends on the value
30 of the global sum. In many cases, this problem is avoided by picking a special node which will first gather data from all the other nodes, do the whole computation and then broadcast the sum to all nodes. However, when the number of nodes is sufficiently large, this method is slower than the shift and operate method.

In addition, as indicated above, in the computer disclosed in provisional patent application no. 60/271,124, the nodes are also connected by a dual-functional tree network that supports integer combining operations, such as integer sums and integer maximums (max) and minimums (min). The existence of a global
5 combining network opens up possibilities to efficiently implement global arithmetic operations over this network. For example, adding up floating point numbers from each of the computing nodes, and broadcasting the sum to all participating nodes. On a regular parallel supercomputer, these kinds of operations are usually done over the network that carries the normal message-passing traffic. There is usually high
10 latency associated with such kinds of global operations.

SUMMARY OF THE INVENTION

An object of this invention is to improve procedures for computing global values for global operations on a distributed parallel computer.

15

Another object of the present invention is to compute a unique global value for a global operation using the shift and operate method in a highly efficient way on distributed parallel M-torus architectures with a large number of nodes.

20 A further object of the invention is to provide a method and apparatus, working in conjunction with software algorithms and hardware implementations of class network routing, to achieve a very significant reduction in the time required for global arithmetic operations on a torus architecture.

25 Another object of this invention is to efficiently implement global arithmetic operations on a network that supports global combining operations.

A further objective of the invention is to implement global arithmetic operations to generate binary reproducible results.

30

An object of the present invention is to provide an improved procedure for conducting a global sum operation.

A further object of the invention is to provide an improved procedure for conducting a global all gather operation.

These and other objectives are attained with the below described methods and systems for performing arithmetic functions. In accordance with a first aspect of the invention, methods and apparatus are provided, working in conjunction of software algorithms and hardware implementation of class network routing, to achieve a very significant reduction in the time required for global arithmetic operations on the torus. This leads to greater scalability of applications running on large parallel machines. The invention involves three steps for improving the efficiency, accuracy and exact reproducibility of global operations:

1. Ensuring, when necessary, that all the nodes do the global operation on the data in the same order and so obtain a unique answer, independent of roundoff error.
2. Using the topology of the torus to minimize the number of hops and the bidirectional capabilities of the network to reduce the number of time steps in the data transfer operation to an absolute minimum.
3. Using class function routing (patent application no. _____ (Attorney Docket 15270)) to reduce latency in the data transfer. With the method of this invention, every single element is injected into the network only once and it will be stored and forwarded without any further software overhead.

In accordance with a second aspect of the invention, methods and systems are provided to efficiently implement global arithmetic operations on a network that supports the global combining operations. The latency of doing such global operations are greatly reduced by using these methods. In particular, with a combining tree network that supports integer maximum MAX, addition SUM, and bitwise AND, OR, and XOR, one can implement virtually all predefined global reduce operations in MPI (Message-Passing Interface Standard): MPI_SUM, MPI_MAX, MPI_MIN, MPI_LAND, MPI_BAND, MPI_LOR, MPI_BOR, MPI_LXOR, MPI_BXOR, MPI_MAXLOC, AND MPI_MINLOC plus MPI_ALLGATHER over this network. The implementations are easy and efficient, demonstrating the great

flexibility and efficiency a combining tree network brings to a large scale parallel supercomputer.

Further benefits and advantages of the invention will become apparent from a
5 consideration of the following detailed description, given with reference to the accompanying drawings, which specify and show preferred embodiments of the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

10 Figure 1 schematically represents a torus network that connects the nodes of a computer. The wrapped network connections are not shown.

Figure 2 schematically represents a tree network that connects the nodes of a computer.

15

Figure 3 illustrates a procedure for performing a global sum on a one-dimensional torus.

Figure 4 is a table identifying steps that can be used to improve the efficiency of
20 global arithmetic operations on a torus architecture.

Figure 5 illustrates the operation of global sum on a dual-functional tree network.

Figure 6 illustrates the operation of global all gathering on a dual-functional tree
25 network.

Figure 7 illustrates a 3 by 4 torus network.

Figure 8 illustrates a tree network for doing a final broadcast operation.
30

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

The present invention relates to performing arithmetic functions in a computer, and one suitable computer is disclosed in provisional patent application no. 60/271,124.

This computer is comprised of many computing nodes and a smaller number of I/O nodes; and the nodes of this computer are interconnected by both a torus network, schematically represented at 10 in Figure 1, and a dual functional tree network, schematically represented at 20 in Figure 2.

5

More specifically, one aspect of the present invention provides a method and apparatus working in conjunction with software algorithms and hardware implementation of class network routing, to achieve a very significant reduction in the time required for global arithmetic operation on the torus architecture.

10 Therefore, it leads to greater scalability of applications running on large parallel machines. As illustrated in Figure 3, the invention involves three steps in improving the efficiency and accuracy of global operations:

- 15 1. Ensuring, when necessary, that all the nodes do the global operation on the data in the same order, and so obtain a unique answer, independent of roundoff error.
2. Using the topology of the torus to minimize the number of hops and the bi-directional capabilities of the network to reduce the number of time steps in the data transfer operation to an absolute minimum.
- 20 3. Using class function routing to reduce latency in the data transfer. With the preferred method of this invention, every single element is injected into the network only once and it will be stored and forwarded without any further software overhead.

25 Each of these steps is discussed below in detail.

1. Ensuring that all nodes to the global operation (eg. MPI_SUM) in the same order:

30 When doing the one-dimensional shift and addition of the local partial sums, instead of adding numbers when they come in, each node will keep the N-1 numbers received for each direction. The global operation is performed on the numbers after

they have all been received so that the operation is done in a fixed order and results in a unique result on all nodes.

For example, as illustrated in Figure 4, if each node adds numbers as they are received, the sums computed would be $S_0+S_3+S_2+S_1$ on node0, $S_1+S_0+S_3+S_2$ on node1, $S_2+S_1+S_0+S_3$ on node3, and $S_3+S_2+S_1+S_0$ on node4. There should be roundoff differences between these sums. However, if the numbers that every node receives are kept in memory, then all the nodes could do the sum in the same order to obtain $S_0+S_1+S_2+S_3$ and there would be no roundoff difference.

10

This is repeated for all the other dimensions. At the end, all the nodes will have the same number and the final broadcast is unnecessary.

2. Minimizing the number of steps in the data transfer on an M-torus architecture:

15

On any machine where the network links between two neighboring nodes are bi-directional, we can send data in both directions in each of the steps. This will mean that the total distance each data element has to travel on the network is reduced by a factor of two. This reduces the time for doing global arithmetic on the torus also by almost a factor of two.

20

3. Reducing the latency using class function routing:

Additional performance gains can be achieved by including a store and forward class network routing operation in the network hardware, thereby eliminating the software overhead of extracting and injecting the same data element multiple times into the network. When implementing global arithmetic operations on a network capable of class routing, steps 1 to 3 illustrated in Figure 4 will simply become a single network step; i.e., each node will need to inject a number only once and every other node will automatically forward this number along to all the other nodes that need it, while keeping a copy for its own use. This will greatly reduce the latency of the global operation. Instead of paying software overhead every hop on the torus, one will pay only a single overhead cost per dimension of the machine. For

30

example, on the computer system disclosed in Provisional Application No. 60/271,124, we estimate that there will be an improvement of at least a factor of five when the CPU is running in single user mode and more than a factor of ten in multi-user mode.

5

With the three improvement steps discussed above, one can achieve at least a factor of ten improvement for global arithmetic operations on distributed parallel architectures and also greatly improve the scalability of applications on a large parallel machine.

10

In addition, as previously mentioned, in the computer system disclosed in the above-identified provisional application, the nodes are also connected by a tree network that supports data combining operations, such as integer sums and integer maximums and minimums, bitwise AND, OR and XORs. In addition, the tree network will automatically broadcast the final combined result to all participating nodes. With a computing network supporting global combining operations, many of the global communication patterns can be efficiently supported by this network. By far the simplest requirement for the combining network hardware is to support unsigned integer add and unsigned integer maximum up to certain precision. For example, the supercomputer disclosed in the above-identified provisional patent application will support at least 32 bit, 64 bit and 128 bit unsigned integer sums or maximums, plus a very long precision sum or maximum up to the 2048 bit packet size. The combining functions in the network hardware provide great flexibility in implementing high performance global arithmetic functions. A number of examples of these implementations are presented below.

25

1. Global sum of signed integers

Figure 5 shows the operation of global sum. Each participating node has an equal size array of numbers with the same number of array elements. The result of the global sum operation is that each node will have the sum of the corresponding elements of arrays from all nodes. This relates to the MPI_SUM function in the MPI (Message Passing Interface) standard.

30

It is usually necessary to use a higher precision in the network compared to each local number to maintain the precision of the final result. Let N be the number of nodes participating in the sum, M be the largest absolute value of the integer numbers to be summed, and 2^P be a large positive integer number greater than M .

- 5 To implement signed integer sum in a network that supports the unsigned operation, we only need to

(1) Add the large positive integer 2^P to all the numbers to be summed so that they now become non-negative.

10

(2) Do the unsigned integer sum over the network.

(3) Subtract $(N * 2^P)$ from the result.

- 15 P is chosen so that $2^P > M$, and $(N * 2^P)$ will not overflow in the combining network.

2. Global Max and Min of signed integers

- 20 These operations are very similar to the global sum discussed above, except that the final result is not the sum of the corresponding elements but the maximum or the minimum one. They relate to `MPI_MAX` and `MPI_MIN` functions in the MPI standard with the integer inputs. The implementation of global max is very similar to the implementations of the global sum, as discussed above.

25

(1) Add a large positive integer 2^P to all numbers to make them non-negative.

(2) Do the unsigned global max over the network.

30

(3) Subtract 2^P from the result.

To do a global min, just negate all the numbers and do a global max.

3. Global sum of floating point numbers:

The operation of the global sum of floating point numbers is very similar to the earlier discussed integer sums except that now the inputs are floating point numbers.

- 5 For simplicity, we will demonstrate summing of one number from each node. To do an array sum, just repeat the steps.

The basic idea is to do two round-trips on the combining network.

- 10 (1) Find the integer maximum of the exponent of all numbers, $E_{\max}$, using the steps outlined in the discussion of the Global max.

- (2) Each node will then normalize its local number and convert it to an integer. Let the local number on node "i" be X_i , whose
15 exponent is E_i . Using the notation defined in the description of the Global sum, this conversion corresponds to the calculation,

$$A_i = 2^P + 2^{[P - (E_{\max} - E) - 1]} * X_i,$$

[Eq. 1]

20

Where A_i is an unsigned integer. A global unsigned integer sum can then be preformed on the network using the combining hardware. Once the final sum A has arrived at each node, the true sum S can be obtained on each node locally by calculating

25

$$S = (A - N * 2^P) / 2^{(P-1)}.$$

Again, P is chosen so that $N * 2^{(P+1)}$ will not overflow in the combining network.

- 30 It should be noted that the step done in equation (1) above is achieved with the best possible precision by using a microprocessor's floating point unit to convert negative numbers to positive and then by using its integer unit to do proper shifting

One important feature of this floating point sum algorithm is that because the actual sum is done through an integer sum, there is no dependence on how the order of the sum is carried out. Each participating node will get the exact same number after the global sum. No additional broadcast from a special node is necessary, which is
5 usually the case when the floating point global sum is implemented through the normal message passing network.

Those skilled in the art will recognize that even if the network hardware supports only unsigned integer sums, when integers are represented in 2's complementary
10 format, correct sums will be obtained as long as no overflow occurs on any final and intermediate results and the carry bit of the sum over any two numbers are dropped by the hardware. The simplification of the operational steps to the global integer and floating point sums comes within the scope of the invention, as well as when the network hardware directly supports signed integer sums with correct overflow
15 handling.

For example, when the hardware only supports unsigned integer sums and drops all carry bits from unsigned integer overflow, such as implemented on the supercomputer, disclosed in provisional patent applications no. 60/271,124 a
20 simplified signed integer sum steps could be:

- (1) sign extend each integer to a higher precision to ensure no overflow of any results would occur; i.e., pad 0 to all the extended high order bits for positive integers and zero, pad 1, to all the extended bit for negative integers.
- 25 (2) do the sum over the network. The final result will have the correct sign.

The above can also be applied to the summing step of floating point sums.

With a similar modification from the description of the Global Sum of integers to
30 the description of the Global max, floating point max and min can also easily be obtained.

There is also a special case for floating point max of non-negative numbers, the operation can be accomplished in one round trip instead of two. For numbers using the IEEE 754 Standard for Floating Point Binary Arithmetic format, as in most of the modern microprocessors, no additional local operations are required. With
5 proper byte ordering, each node can just put the numbers on the combining network. For other floating point formats, like those used in some Digital Signal Processors, some local manipulation of the exponent field may be required. The same single round-trip can also be achieved for the min of negative numbers by doing a global max on their absolute values.

10

4. Global all gather operation using integer global sum

The global all gather operation is illustrated In Figure 6. Each node contributes one or more numbers. The final result is that these number are put into an array with
15 their location corresponding to where they came from. For example, numbers from node 1 appear first in the final array, followed by numbers from node 2, ..., etc. This operation corresponds to the MPI_ALLGATHER function in the MPI standard.

This function can be easily implemented in a one pass operation on a combining
20 network supporting integer sums. Using the fact that adding zero to a number yields the same number, each node simply needs to assemble an array whose size equals the final array, and then it will put its numbers in the corresponding place and put zero in all other places corresponding to numbers from all other nodes. After an integer sum of arrays from all nodes over the combining network, each node will
25 have the final array with all the numbers sorted into their places.

5. Global min_loc and max_loc, using integer global max

These functions correspond to MPI_MINLOC and MPI_MAXLOC in the MPI
30 standard. Besides finding the global minimum or maximum, an index is appended to each of the numbers so that one could find out which node has the global minimum or maximum, for example.

On a combining network that supports integer global max, these functions are straight forward to implement. We will illustrate global max_loc as an example. Let node "j", $j=1,...,N$, have number X_j and index K_j . Let M be a large integer number, $M > \max(K_j)$, the node "j" only needs put two numbers:

5

$$X_j$$

$$M - K_j$$

10 as a single unit onto the combining network for global integer max. At the end of the operation, each node would receive:

$$X$$

15

$$M - K$$

Where $X = \max(X_j)$ is the maximum value of all X_j 's, and K is the index number that corresponds to the maximum X . If there is more than one number equal to the maximum X , then K is the lowest index number.

20

Global min_loc can be achieved similarly by changing X_j to $P - X_j$ in the above where P is a large positive integer number and $P > \max(X_j)$.

25 The idea of appending the index number behind the number in the global max or mix operation also applies to floating pointing numbers. With steps similar to those described above in the discussion of the procedure for performing the global sum of floating point numbers.

6. Other operations:

30

On the supercomputer system described in the provisional patent application no. 60/271,124, additional global bitwise AND, OR, and XOR operations are also supported on the combining network. This allows for very easy implementation of

global bitwise reduction operations, such as MPI_BAND, MPI_BOR and MPI_BXOR in the MPI standard. Basically, each node just needs to put the operand for the global operation onto the network, and the global operations are handled automatically by the hardware.

5

In addition, logical operations MPI LAND, MPI_LOR and MPI_LXOR can also be implemented by just using one bit in the bitwise operations.

Finally, each of the global operations also implies a global barrier operation. This is because the network will not proceed until all operands are injected into the network. Therefore, efficient MPI_BARRIER operations can also be implemented using any one of the global arithmetic operations, such as the global bitwise AND.

10

7. Operations using both the torus and tree networks:

15

Depending on the relative bandwidths of the torus and tree networks, and on the overhead to do the necessary conversions between floating and fixed point representations, it may be more efficient to use both the torus and tree networks simultaneously to do global floating point reduction operations. In such a case, the torus is used to do the reduction operation, and the tree is used to broadcast the results to all nodes. Prior art for doing reductions on a torus are known. However, in prior art, the broadcast phase is also done on the torus. For example, in a 3 by 4 torus (or mesh) as illustrated at 30 in Figure 7, reductions are done down rows, and then down columns by nodes at the ends of the rows. In particular, in a sum reduction, Figure 7 depicts node Q20 inserting a packet and sending it to node Q21. Q21 processes this packet by adding its corresponding elements to those of the incoming packet and then sending a packet containing the sum to Q22. Q22 processes this packet by adding its corresponding elements to those of the incoming packet and then sending a packet containing the sum to Q23. These operations are repeated for each row. Node Q23 sums its local values to the corresponding values in the packet from Q22 and sends the resulting packet to Q13. Node Q13 sums its local values to those of the packets from Q12 and Q23, and sends the resulting sum to Q03. Q03 sums its local values to the corresponding values of the packet from

20
25
30

Q13. Q03 now has the global sum. In prior art, this global sum is sent over the torus network to all the other nodes (rather than on the tree as shown in the figure). The extension to more nodes and to a higher dimensional torus is within the ability of those skilled in the art and within the scope of the present invention. For
5 reductions over a large number of values, multiple packets are used in a pipelined fashion.

However, the final broadcast operation can be done faster and more efficiently by using the tree, rather than the torus, network. This is illustrated in Figure 8.
10 Performance can be optimized by having one processor handle the reduction operations on the torus and a second processor handle the reception of the packets broadcast on the tree. Performance can further be optimized by reducing the number of hops in the reduction step. For example, packets could be sent (and summed) to the middle of the rows, rather than to the end of the rows.

15 In a 3-dimensional torus, the straightforward extension of the above results in a single node in each z plane summing their values up the z dimension. This has the disadvantage of requiring those nodes to process three incoming packets. For example, node Q03z has to receive packets from Q02z, Q13z, and Q03(z+1). If the
20 processor is not fast enough this will become the bottleneck in the operation. To optimize performance, we modify the communications pattern so that no node is required to process more than 2 incoming packets on the torus. This is illustrated in Figure 8. In this Figure, node Q03z forwards its packets to node Q00z for summing down the z-dimension. In addition, node Q00z does not send its packets to node
25 Q01z but rather receives a packet from node Q00(z+1) and sums its local values with the corresponding values of its two incoming packets. Finally, node Q000 broadcasts the final sum over the tree network.

While it is apparent that the invention herein disclosed is well calculated to fulfill
30 the objects stated above, it will be appreciated that numerous modifications and embodiments may be devised by those skilled in the art, and it is intended that the appended claims cover all such modifications and embodiments as fall within the true spirit and scope of the present invention.

WHAT IS CLAIMED IS:

- 1 1. A method of performing arithmetic functions, using a shift and operate
2 procedure, in a computer system having a distributed parallel torus architecture with
3 a multitude of interconnected nodes, the method comprising the steps:
4
5 providing each of a group of the nodes with the same set of data values;
6
7 performing a global arithmetic operation, wherein each of the nodes performs the
8 arithmetic operation on all of the data values to obtain a final value; and
9
10 ensuring that all of the nodes of the group perform said global operation on the data
11 values in the same order to ensure binary reproducible results.
- 1 2. A method according to Claim 1, wherein the ensuring step includes the step of,
2 each node performing the global arithmetic operation after the node is provided with
3 all of the data values.
- 1 3. A method according to Claim 2, wherein the providing step includes the step of
2 each node of the group receiving the data values from other nodes of the group.
- 1 4. A method according to Claim 1, wherein the nodes are connected together by bi-
2 directional links, and the providing step includes the step of sending the data values
3 to the nodes in two directions over said links.
- 1 5. A method according to Claim 1, wherein the providing step includes the step of
2 each one of the nodes injecting one of the data values into the network only once.
- 1 6. A method according to Claim 5, wherein the injecting step includes the step of,
2 nodes of the group of nodes, other than said each one of the nodes, forwarding said
3 one of the data values to other nodes of the group to reduce the latency of the global
4 operation.

1 7. A system for performing arithmetic functions, using a shift and operate
2 procedure, in a computer system having a distributed parallel torus architecture with
3 a multitude of interconnected nodes, the system comprising:
4
5 a group of the nodes provided with the same set of data values;
6
7 means for performing a global arithmetic operation, wherein each of the nodes
8 performs the arithmetic operation on all of the data values to obtain a final value;
9 and
10
11 means for ensuring that all of the nodes of the group perform said global operation
12 on the data values in the same order to ensure binary reproducible results.

1 8. A system according to Claim 7, wherein the ensuring means includes means for
2 performing the global arithmetic operation at each node after the node is provided
3 with all of the data values.

1 9. A system according to Claim 7, wherein each node of the group receives the data
2 values from other nodes of the group.

1 10. A system according to Claim 7, wherein the nodes are connected together by bi-
2 directional links, and the providing means includes means for sending the data
3 values to the nodes in two directions over said links.

1 11. A system according to Claim 7, wherein each one of the nodes injects one of the
2 data values into the network only once.

1 12. A system according to Claim 7, wherein nodes of the group of nodes, other than
2 said each one of the nodes, forward said one of the data values to other nodes of the
3 group to reduce the latency of the global operation.

1 13. A program storage device readable by machine, tangibly embodying a program
2 of instructions executable by the machine to perform method steps for performing

3 arithmetic functions, using a shift and operate procedure, in a computer system
4 having a distributed parallel torus architecture with a multitude of interconnected
5 nodes, the method steps comprising:
6
7 providing each of a group of the nodes with the same set of data values;
8
9 performing a global arithmetic operation, wherein each of the nodes performs the
10 arithmetic operation on all of the data values to obtain a final value; and
11
12 ensuring that all of the nodes of the group perform said global operation on the data
13 values in the same order to ensure binary reproducible results.

1 14. A program storage device according to Claim 13, wherein the ensuring step
2 includes the step of, each node performing the global arithmetic operation after the
3 node is provided with all of the data values.

1 15. A program storage device according to Claim 14, wherein the providing step
2 includes the step of each node of the group receiving the data values from other
3 nodes of the group.

1 16. A program storage device according to Claim 13, wherein the nodes are
2 connected together by bi-directional links, and the providing step includes the step
3 of sending the data values to the nodes in two directions over said links.

1 17. A program storage device according to Claim 13, wherein the providing step
2 includes the step of each one of the nodes injecting one of the data values into the
3 network only once.

1 18. A program storage device according to Claim 17, wherein the injecting step
2 includes the step of, nodes of the group of nodes, other than said each one of the
3 nodes, forwarding said one of the data values to other nodes of the group to reduce
4 the latency of the global operation.

1 19. A method of performing an arithmetic function in a computer system having a
2 multitude of nodes interconnected by a global tree network that supports integer
3 combining operations, the method comprising the steps of:
4
5 providing each of a group of nodes with first values;
6
7 processing each of the first values, according to a first defined process, to obtain a
8 respective second value from each of the first values, wherein all of the second
9 values are integer values; and
10
11 performing a global integer combine operation, using said second values, over the
12 network.

1 20. A method according to Claim 19, wherein the performing step includes the step
2 of performing a global unsigned integer sum over the network.

1 21. A method according to Claim 19, wherein the performing step includes the step
2 of performing a global maximum operation over the network, and using results of
3 said global maximum operation to identify the maximum of the first values.

1 22. A method according to Claim 19, wherein the performing step includes the step
2 of performing a global maximum operation over the network, and using the results
3 of said global maximum operation to identify the minimum of the first values.

1 23. A system for performing an arithmetic function in a computer system having a
2 multitude of nodes interconnected by a global tree network that supports integer
3 combining operations, the method comprising the steps of:
4
5 a group of nodes, each of the nodes of the group being provided with first values;
6
7 a processor to process each of the first values, according to a first defined process, to
8 obtain a respective second value from each of the first values, wherein all of the
9 second values are integer values; and

10 means for performing a global integer combine operation, using said second values,
11 over the network.

1 24. A system according to Claim 23, wherein the means for performing includes
2 means for performing a global unsigned integer sum over the network.

1 25. A system according to Claim 23, wherein the means for performing step includes
2 means for performing a global maximum operation over the network, and for using
3 results of said global maximum operation to identify the maximum of the first
4 values.

1 26. A system according to Claim 23, wherein the means for performing step
2 includes means for performing a global maximum operation over the network, and
3 for using the results of said global maximum operation to identify the minimum of
4 the first values.

1 27. A program storage device readable by machine, tangibly embodying a program
2 of instructions executable by the machine to perform method steps for performing an
3 arithmetic function in a computer system having a multitude of nodes interconnected
4 by a global tree network that supports integer combining operations, the method
5 steps comprising:

6
7 providing each of a group of nodes with first values;

8
9 processing each of the first values, according to a first defined process, to obtain a
10 respective second value from each of the first values, wherein all of the second
11 values are integer values; and

12
13 performing a global integer combine operating, using said second values, over the
14 network.

1 28. A program storage device according to Claim 27, wherein the performing step
2 includes the step of performing a global unsigned integer sum over the network.

1 29. A program storage device according to Claim 27, wherein the performing step
2 includes the step of performing a global maximum operation over the network, and
3 using results of said global maximum operation to identify the maximum of the first
4 values.

1 30. A program storage device according to Claim 27, wherein the performing step
2 includes the step of performing a global maximum operation over the network, and
3 using the results of said global maximum operation to identify the minimum of the
4 first values.

1 31. A method of performing a global operation on a computer system having a
2 multitude of nodes interconnected by a global tree network that supports integer
3 combining operations, the method comprising:
4

5 providing each node with one or more numbers of any type;
6

7 assembling the numbers of the nodes into an array, the array having a given number
8 of positions, said assembling step including the steps of
9

10 i) each node putting one or more of the numbers of the node into
11 one or more of the positions of the array, and putting zero values
12 into all of the other positions of the array, and
13

14 ii) using the global tree network to sum all the numbers put into each
15 position in the array.

1 32. A method according to Claim 31, wherein:
2

3 the given number of position of the array are arranged in a defined sequence; and
4

5 the assembling step includes the further step of, each node establishing an associated
6 array also having said given number of positions arranged in the defined sequence,
7 and putting the one or more numbers of the node into one or more of the positions of

8 the associated array, and putting zero values in all of the other positions of the
9 associated array.

1 33. A system for performing a global operation on a computer system having a
2 multitude of nodes interconnected by a global tree network that supports integer
3 combining operations, the system comprising:

4
5 a group of nodes, each of the group of nodes having one or more numbers;

6
7 means for assembling the numbers of the nodes into an array, the array having a
8 given number of positions, said assembling means including

9
10 i) means for putting one or more of the numbers of the nodes of the
11 group into one or more of the positions of the array, and for
12 putting zero values into all of the other positions of the array, and

13
14 ii) means for using the global tree network to sum all the numbers
15 put into each position in the array.

1 34. A system according to Claim 33, wherein:

2
3 the given number of position of the array are arranged in a defined sequence; and

4
5 the means for assembling further includes means for establishing a respective one
6 array associated with each of the nodes of the group and also having said given
7 number of positions arranged in the defined sequence, for putting the one or more
8 numbers of each of the nodes into one or more of the positions of the associated
9 array, and for putting zero values in all of the other positions of the associated array.

1 35. A program storage device readable by machine, tangibly embodying a program
2 of instructions executable by the machine to perform method steps for performing a
3 global operation on a computer system having a multitude of nodes interconnected

4 by a global tree network that supports integer combining operations, the method
5 steps comprising:
6 providing each node with one or more numbers;
7
8 assembling the numbers of the nodes into an array, the array having a given number
9 of positions, said assembling step including the steps of
10
11 a. each node putting one or more of the numbers of the node into one or
12 more of the positions of the array, and putting zero values into all of
13 the other positions of the array, and
14
15 b. using the global tree network to sum all the numbers put into each
16 position in the array.

1 36. A method according to Claim 35, wherein:
2 the given number of position of the array are arranged in a defined sequence; and
3
4 the assembling step includes the further step of, each node establishing an associated
5 array also having said given number of positions arranged in the defined sequence,
6 and putting the one or more numbers of the node into one or more of the positions of
7 the associated array, and putting zero values in all of the other positions of the
8 associated array.

1 37. A method of performing an arithmetic function in a computer systems having a
2 multitude interconnected by a global tree network that supports integers combining
3 operations, the method comprising the step of:
4
5 each of the nodes contributing a set of first values; and
6
7 performing a global integer combine operation, using said first values, over the
8 network.

1 38. A method according to Claim 37, wherein the performing step includes the step
2 of using results of the global integers combine operations to identify a characteristic
3 of the first values

1 39. A system for performing an arithmetic function is a computer system having a
2 multitude of nodes interconnected by a global tree network that supports integers
3 combining operations, the system comprising:
4
5 a group of nodes, each of the nodes of the group consisting of set of first values; and
6
7 a processor to perform a global integer combine operations, using the first values,
8 over the network.

1 40. A system according to Claim 39, wherein the processor includes means for
2 using results of the global integers combine operation to identify a characteristic of
3 the first values.

1 41. A method of operating a parallel processing computer system having a
2 multitude of nodes interconnected by both a global tree network and a torus
3 network, the method comprising:
4
5 using the computer systems to perform defined operations; and
6
7 using both the torus and tree networks to cooperate on reduction operations.

1 42. A method according to Claim 41, wherein the step of using both the torus and
2 tree networks includes the step of doing so by having one processor handle torus
3 operations and another processor handle the tree operations.

1 43. A method according to Claim 41, wherein the step of using both the torus and
2 tree networks includes the step of doing so by arranging the torus communications
3 so that, in a three-dimensional torus, no node on the torus receives more than two
4 packets to combine.

1 44. A program storage device readable by machine, tangibly embodying a program
2 of instructions executable by the machine to perform method steps for operating a
3 parallel processing computer system having a multitude of nodes interconnected by
4 both a global tree network and a torus network, the method step comprising:
5
6 using the computer systems to perform defined operations; and
7
8 using both the torus and tree networks to cooperate on reduction operations.

1 45. A program storage device according to Claim 44, wherein the step of using both
2 the torus and tree networks includes the step of doing so by having one processor
3 handle torus operations and another processor handle the tree operations.

1 46. A program storage device according to Claim 44, wherein the step of using both
2 the torus and tree networks includes the step of doing so by arranging the torus
3 communications so that, in a three-dimensional torus, no node on the torus receives
4 more than two packets to combine.

1 47. A parallel processing computer system comprising:
2
3 a multitude of nodes;
4
5 a global tree network also interconnecting the nodes;
6
7 a torus network also interconnecting the nodes; and
8
9 means for using both the torus and tree networks to cooperate on reduction
10 operations.

1 48. A computer system according to Claim 47, wherein the means for using both the
2 torus and tree networks include one processor to handle torus operations and another
3 processor to handle the tree operations.

1 49. A computer system according to Claim 47, wherein the means for using both the
2 torus and tree networks includes means for arranging the torus communications so
3 that, is a three-dimensional torus, no node on the torus receives more than two
4 packets to combine.