

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
3 November 2011 (03.11.2011)

PCT

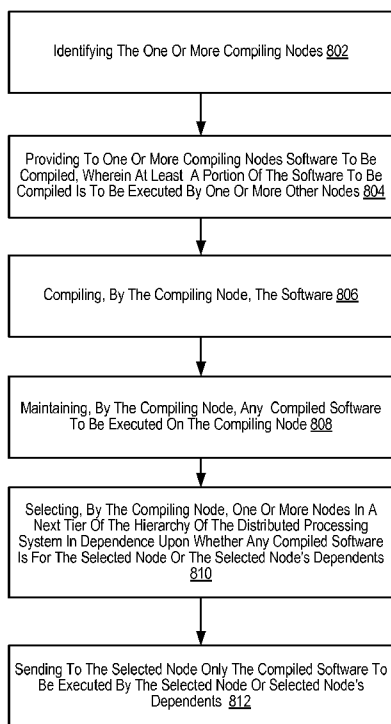
(10) International Publication Number
WO 2011/134689 A1

- (51) International Patent Classification:
G06F 9/45 (2006.01)
- (21) International Application Number:
PCT/EP2011/052105
- (22) International Filing Date:
14 February 2011 (14.02.2011)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
12/770,353 29 April 2010 (29.04.2010) US
- (71) Applicant (for all designated States except US): **INTERNATIONAL BUSINESS MACHINES CORPORATION** [US/US]; New Orchard Road, Armonk, New York 10504 (US).
- (71) Applicant (for MG only): **IBM UNITED KINGDOM LIMITED** [GB/GB]; PO Box 41, North Harbour, Portsmouth Hampshire PO6 3AU (GB).

- (72) Inventors; and
- (75) Inventors/Applicants (for US only): **RATTERMAN, Joseph** [US/US]; IBM Corporation, MD: GHD, 3605 Hwy 52 N, Rochester, Minnesota 55901-1407 (US). **ARCHER, Charles, Jens** [US/US]; IBM Corporation, MD: GHD, 3605 Hwy 52 N, Rochester, Minnesota 55901-1407 (US). **SMITH, Brian, Edward** [US/US]; IBM Corporation, MD:GHD, 3605 Hwy 52 N, Rochester, Minnesota 55901-1407 (US). **BLOCKSOME, Michael, Alan** [US/US]; IBM Corporation, MD: 30-2/GHD, 3605 Hwy 52 N, Rochester, Minnesota 55901-1407 (US).
- (74) Agent: **GASCOYNE, Belinda**; IBM United Kingdom Limited, Intellectual Property Law, Hursley Park, Winchester Hampshire SO21 2JN (GB).
- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI,

[Continued on next page]

(54) Title: COMPILING SOFTWARE FOR A HIERARCHICAL DISTRIBUTED PROCESSING SYSTEM



(57) Abstract: Compiling software for a hierarchical distributed processing system including providing to one or more compiling nodes software to be compiled, wherein at least a portion of the software to be compiled is to be executed by one or more other nodes; compiling, by the compiling node, the software; maintaining, by the compiling node, any compiled software to be executed on the compiling node; selecting, by the compiling node, one or more nodes in a next tier of the hierarchy of the distributed processing system in dependence upon whether any compiled software is for the selected node or the selected node's descendants; sending to the selected node only the compiled software to be executed by the selected node or selected node's descendant.

FIG. 7

WO 2011/134689 A1



NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

(84) Designated States (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK,

Published:

— *with international search report (Art. 21(3))*

COMPILING SOFTWARE FOR A HIERARCHICAL DISTRIBUTED PROCESSING SYSTEM

BACKGROUND OF THE INVENTION

5

Field of the Invention

The field of the invention is data processing, or, more specifically, methods, apparatus, and products for compiling software for a hierarchical distributed processing system.

10

Description of Related Art

The development of the EDVAC (Electronic Discrete Variable Automatic Computer) computer system of 1948 is often cited as the beginning of the computer era. Since that time, computer systems have evolved into extremely complicated devices. Today's computers are much more sophisticated than early systems such as the EDVAC. Computer systems typically include a combination of hardware and software components, application programs, operating systems, processors, buses, memory, input/output devices, and so on. As advances in semiconductor processing and computer architecture push the performance of the computer higher and higher, more sophisticated computer software has evolved to take advantage of the higher performance of the hardware, resulting in computer systems today that are much more powerful than just a few years ago.

Distributed computing is an area of computer technology that has experienced advances. Distributed computing generally refers to computing with multiple semi-autonomous computer systems that communicate through a data communications network. The semi-autonomous computer systems interact with one another in order to achieve a common goal. A computer program or application that executes in a distributed computing system may be referred to as a distributed program. Distributed computing may also refer to the use of distributed computing systems to solve computational problems. In distributed computing, a problem may be divided into many tasks, each of which may be solved by one of the semi-autonomous computer systems.

30

Some distributed computing systems are optimized to perform parallel computing. Parallel computing is the simultaneous execution of the same task (split up and specially adapted) on multiple processors in order to obtain results faster. Parallel computing is based on the fact that the process of solving a problem usually can be divided into smaller tasks, which may be carried out simultaneously with some coordination.

Parallel computers execute parallel algorithms. A parallel algorithm can be split up to be executed a piece at a time on many different processing devices, and then put back together again at the end to get a data processing result. Some algorithms are easy to divide up into pieces. Splitting up the job of checking all of the numbers from one to a hundred thousand to see which are primes could be done, for example, by assigning a subset of the numbers to each available processor, and then putting the list of positive results back together. In this specification, the multiple processing devices that execute the individual pieces of a parallel program are referred to as 'compute nodes.' A parallel computer is composed of compute nodes and other processing nodes as well, including, for example, input/output ('I/O') nodes, and service nodes.

Parallel algorithms are valuable because it is faster to perform some kinds of large computing tasks via a parallel algorithm than it is via a serial (non-parallel) algorithm, because of the way modern processors work. It is far more difficult to construct a computer with a single fast processor than one with many slow processors with the same throughput. There are also certain theoretical limits to the potential speed of serial processors. On the other hand, every parallel algorithm has a serial part and so parallel algorithms have a saturation point. After that point adding more processors does not yield any more throughput but only increases the overhead and cost.

Parallel algorithms are designed also to optimize one more resource the data communications requirements among the nodes of a parallel computer. There are two ways parallel processors communicate, shared memory or message passing. Shared memory processing needs additional locking for the data and imposes the overhead of additional processor and bus cycles and also serializes some portion of the algorithm.

Message passing processing uses high-speed data communications networks and message buffers, but this communication adds transfer overhead on the data communications networks as well as additional memory need for message buffers and latency in the data communications among nodes. Designs of parallel computers use specially designed data communications links so that the communication overhead will be small but it is the parallel algorithm that decides the volume of the traffic.

Many data communications network architectures are used for message passing among nodes in parallel computers. Compute nodes may be organized in a network as a 'torus' or 'mesh,' for example. Also, compute nodes may be organized in a network as a tree. A torus network connects the nodes in a three-dimensional mesh with wrap around links. Every node is connected to its six neighbors through this torus network, and each node is addressed by its x,y,z coordinate in the mesh. In such a manner, a torus network lends itself to point to point operations. In a tree network, the nodes typically are connected into a binary tree: each node has a parent, and two children (although some nodes may only have zero children or one child, depending on the hardware configuration). Although a tree network typically is inefficient in point to point communication, a tree network does provide high bandwidth and low latency for certain collective operations, message passing operations where all compute nodes participate simultaneously, such as, for example, an allgather operation. In computers that use a torus and a tree network, the two networks typically are implemented independently of one another, with separate routing circuits, separate physical links, and separate message buffers.

SUMMARY OF THE INVENTION

Methods, apparatus, and products for compiling software for a hierarchical distributed processing system including providing to one or more compiling nodes software to be compiled, wherein at least a portion of the software to be compiled is to be executed by one or more other nodes; compiling, by the compiling node, the software; maintaining, by the compiling node, any compiled software to be executed on the compiling node; selecting, by the compiling node, one or more nodes in a next tier of the hierarchy of the distributed processing system in dependence upon whether any compiled software is for the selected

node or the selected node's descendants; sending to the selected node only the compiled software to be executed by the selected node or selected node's descendant; receiving, by a selected node, compiled software; determining whether the compiled software is for the selected node or one of its descendants; if the compiled software is for the selected node,
5 maintaining the software by the selected node for execution; and if the compiled software is for one of the descendants, selecting another node in a next tier of the hierarchical distributed processing system in dependence upon a descendant for the compiled software and sending the compiled software to the selected another node.

10 The foregoing and other objects, features and advantages of the invention will be apparent from the following more particular descriptions of exemplary embodiments of the invention as illustrated in the accompanying drawings wherein like reference numbers generally represent like parts of exemplary embodiments of the invention.

15 BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 illustrates an exemplary distributed computing system for compiling software for a hierarchical distributed processing system according to embodiments of the present invention.

20 Figure 2 sets forth a block diagram of an exemplary compute node useful in a parallel computer capable of for compiling software for a hierarchical distributed processing system according to embodiments of the present invention.

25 Figure 3A illustrates an exemplary Point To Point Adapter useful in systems capable of for compiling software for a hierarchical distributed processing system according to embodiments of the present invention.

30 Figure 3B illustrates an exemplary Global Combining Network Adapter useful in systems capable of for compiling software for a hierarchical distributed processing system according to embodiments of the present invention.

Figure 4 sets forth a line drawing illustrating an exemplary data communications network optimized for point to point operations useful in systems capable of for compiling software for a hierarchical distributed processing system in accordance with embodiments of the present invention.

5

Figure 5 sets forth a line drawing illustrating an exemplary data communications network optimized for collective operations useful in systems capable of for compiling software for a hierarchical distributed processing system in accordance with embodiments of the present invention.

10

Figure 6 sets forth a further exemplary distributed computing system for compiling software for a hierarchical distributed processing system according to embodiments of the present invention in which the distributed computing system is implemented as a hybrid computing environment.

15

Figure 7 sets forth an exemplary method of for compiling software for a hierarchical distributed processing system according to embodiments of the present invention.

20

Figure 8 sets forth a flow chart illustrating another exemplary method of compiling software for a hierarchical distributed processing system according to embodiments of the present invention.

25

Figure 9 sets forth a block diagram of an exemplary use case of a system for compiling software for a hierarchical distributed processing system according to embodiments of the present invention.

DETAILED DESCRIPTION OF EXEMPLARY EMBODIMENTS

30

Exemplary methods, apparatus, and products for compiling software for a hierarchical distributed processing system in accordance with embodiments of the present invention are described with reference to the accompanying drawings, beginning with Figure 1. Figure 1 illustrates an exemplary distributed computing system for compiling software for a

hierarchical distributed processing system according to embodiments of the present invention. The system of Figure 1 includes a parallel computer (100), non-volatile memory for the computer in the form of data storage device (118), an output device for the computer in the form of printer (120), and an input/output device for the computer in the form of computer terminal (122). Parallel computer (100) in the example of Figure 1 includes a plurality of compute nodes (102).

The compute nodes (102) are coupled for data communications by several independent data communications networks including a Joint Test Action Group ('JTAG') network (104), a global combining network (106) which is optimized for collective operations, and a torus network (108) which is optimized point to point operations. The global combining network (106) is a data communications network that includes data communications links connected to the compute nodes so as to organize the compute nodes as a tree. Each data communications network is implemented with data communications links among the compute nodes (102). The data communications links provide data communications for parallel operations among the compute nodes of the parallel computer. The links between compute nodes are bi-directional links that are typically implemented using two separate directional data communications paths.

In addition, the compute nodes (102) of parallel computer are organized into at least one operational group (132) of compute nodes for collective parallel operations on parallel computer (100). An operational group of compute nodes is the set of compute nodes upon which a collective parallel operation executes. Collective operations are implemented with data communications among the compute nodes of an operational group. Collective operations are those functions that involve all the compute nodes of an operational group. A collective operation is an operation, a message-passing computer program instruction that is executed simultaneously, that is, at approximately the same time, by all the compute nodes in an operational group of compute nodes. Such an operational group may include all the compute nodes in a parallel computer (100) or a subset all the compute nodes. Collective operations are often built around point to point operations. A collective operation requires that all processes on all compute nodes within an operational group call the same collective operation with matching arguments. A 'broadcast' is an example of a collective operation

for moving data among compute nodes of an operational group. A 'reduce' operation is an example of a collective operation that executes arithmetic or logical functions on data distributed among the compute nodes of an operational group. An operational group may be implemented as, for example, an MPI 'communicator.'

5

'MPI' refers to 'Message Passing Interface,' a prior art parallel communications library, a module of computer program instructions for data communications on parallel computers. Examples of prior-art parallel communications libraries that may be improved for use with systems according to embodiments of the present invention include MPI and the 'Parallel Virtual Machine' ('PVM') library. PVM was developed by the University of Tennessee, The Oak Ridge National Laboratory, and Emory University. MPI is promulgated by the MPI Forum, an open group with representatives from many organizations that define and maintain the MPI standard. MPI at the time of this writing is a de facto standard for communication among compute nodes running a parallel program on a distributed memory parallel computer. This specification sometimes uses MPI terminology for ease of explanation, although the use of MPI as such is not a requirement or limitation of the present invention.

10

15

20

25

Some collective operations have a single originating or receiving process running on a particular compute node in an operational group. For example, in a 'broadcast' collective operation, the process on the compute node that distributes the data to all the other compute nodes is an originating process. In a 'gather' operation, for example, the process on the compute node that received all the data from the other compute nodes is a receiving process. The compute node on which such an originating or receiving process runs is referred to as a logical root.

30

Most collective operations are variations or combinations of four basic operations: broadcast, gather, scatter, and reduce. The interfaces for these collective operations are defined in the MPI standards promulgated by the MPI Forum. Algorithms for executing collective operations, however, are not defined in the MPI standards. In a broadcast operation, all processes specify the same root process, whose buffer contents will be sent.

Processes other than the root specify receive buffers. After the operation, all buffers contain the message from the root process.

5 In a scatter operation, the logical root divides data on the root into segments and distributes a different segment to each compute node in the operational group. In scatter operation, all processes typically specify the same receive count. The send arguments are only significant to the root process, whose buffer actually contains $\text{sendcount} * N$ elements of a given data type, where N is the number of processes in the given group of compute nodes. The send buffer is divided and dispersed to all processes (including the process on the logical root).
10 Each compute node is assigned a sequential identifier termed a 'rank.' After the operation, the root has sent sendcount data elements to each process in increasing rank order. Rank 0 receives the first sendcount data elements from the send buffer. Rank 1 receives the second sendcount data elements from the send buffer, and so on.

15 A gather operation is a many-to-one collective operation that is a complete reverse of the description of the scatter operation. That is, a gather is a many-to-one collective operation in which elements of a datatype are gathered from the ranked compute nodes into a receive buffer in a root node.

20 A reduce operation is also a many-to-one collective operation that includes an arithmetic or logical function performed on two data elements. All processes specify the same 'count' and the same arithmetic or logical function. After the reduction, all processes have sent count data elements from computer node send buffers to the root process. In a reduction operation, data elements from corresponding send buffer locations are combined pair-wise by
25 arithmetic or logical operations to yield a single corresponding element in the root process's receive buffer. Application specific reduction operations can be defined at runtime. Parallel communications libraries may support predefined operations. MPI, for example, provides the following pre-defined reduction operations:

30 MPI_MAX maximum
MPI_MIN minimum
MPI_SUM sum

	MPI_PROD	product
	MPI LAND	logical and
	MPI_BAND	bitwise and
	MPI_LOR	logical or
5	MPI_BOR	bitwise or
	MPI_LXOR	logical exclusive or
	MPI_BXOR	bitwise exclusive or

10 In addition to compute nodes, the parallel computer (100) includes input/output ('I/O') nodes (110, 114) coupled to compute nodes (102) through the global combining network (106). The compute nodes in the parallel computer (100) are partitioned into processing sets such that each compute node in a processing set is connected for data communications to the same I/O node. Each processing set, therefore, is composed of one I/O node and a subset of

15 compute nodes (102). The ratio between the number of compute nodes to the number of I/O nodes in the entire system typically depends on the hardware configuration for the parallel computer. For example, in some configurations, each processing set may be composed of eight compute nodes and one I/O node. In some other configurations, each processing set may be composed of sixty-four compute nodes and one I/O node. Such example are for

20 explanation only, however, and not for limitation. Each I/O nodes provide I/O services between compute nodes (102) of its processing set and a set of I/O devices. In the example of Figure 1, the I/O nodes (110, 114) are connected for data communications I/O devices (118, 120, 122) through local area network ('LAN') (130) implemented using high-speed Ethernet.

25 The parallel computer (100) of Figure 1 also includes a service node (116) coupled to the compute nodes through one of the networks (104). Service node (116) provides services common to pluralities of compute nodes, administering the configuration of compute nodes, loading programs into the compute nodes, starting program execution on the compute nodes, retrieving results of program operations on the computer nodes, and so on. Service node

30 (116) runs a service application (124) and communicates with users (128) through a service application interface (126) that runs on computer terminal (122).

In the example of Figure 1, one of the compute nodes has installed upon it a hierarchical distributed compiler (155), a module of automated computing machinery capable of compiling software for a hierarchical distributed processing system according to embodiments of the present invention. The hierarchical distributed compiler (155) includes computer program instruction for compiling, by the compiling node, the software; maintaining, by the compiling node, any compiled software to be executed on the compiling node; selecting, by the compiling node, one or more nodes in a next tier of the hierarchy of the distributed processing system in dependence upon whether any compiled software is for the selected node or the selected node's descendants; sending to the selected node only the compiled software to be executed by the selected node or selected node's descendant. Each of the other compute nodes (102) of Figure 1 are also capable of receiving compiled software; determining whether the compiled software is for that node or one of its descendants; if the compiled software is for that node, maintaining the software for execution; and if the compiled software is for one of the descendants, selecting another node in a next tier of the hierarchical distributed processing system in dependence upon a descendant for the compiled software and sending the compiled software to the selected another node.

The arrangement of nodes, networks, and I/O devices making up the exemplary system illustrated in Figure 1 are for explanation only, not for limitation of the present invention. Data processing systems capable of compiling software for a hierarchical distributed processing system according to embodiments of the present invention may include additional nodes, networks, devices, and architectures, not shown in Figure 1, as will occur to those of skill in the art. Although the parallel computer (100) in the example of Figure 1 includes sixteen compute nodes (102), readers will note that parallel computers capable of compiling software for a hierarchical distributed processing system according to embodiments of the present invention may include any number of compute nodes. In addition to Ethernet and JTAG, networks in such data processing systems may support many data communications protocols including for example TCP (Transmission Control Protocol), IP (Internet Protocol), and others as will occur to those of skill in the art. Various embodiments of the present invention may be implemented on a variety of hardware platforms in addition to those illustrated in Figure 1.

5 Compiling software for a hierarchical distributed processing system according to
embodiments of the present invention may be generally implemented on a parallel computer
that includes a plurality of compute nodes. In fact, such computers may include thousands
of such compute nodes. Each compute node is in turn itself a kind of computer composed of
one or more computer processors (or processing cores), its own computer memory, and its
own input/output adapters. For further explanation, therefore, Figure 2 sets forth a block
diagram of an exemplary compute node useful in a parallel computer capable of compiling
software for a hierarchical distributed processing system according to embodiments of the
present invention. The compute node (152) of Figure 2 includes one or more processing
10 cores (164) as well as random access memory ('RAM') (156). The processing cores (164)
are connected to RAM (156) through a high-speed memory bus (154) and through a bus
adapter (194) and an extension bus (168) to other components of the compute node (152).
Stored in RAM (156) is an application program (158), a module of computer program
instructions that carries out parallel, user-level data processing using parallel algorithms.

15 Also stored in RAM (156) is a messaging module (160), a library of computer program
instructions that carry out parallel communications among compute nodes, including point to
point operations as well as collective operations. Application program (158) executes
collective operations by calling software routines in the messaging module (160). A library
20 of parallel communications routines may be developed from scratch for use in systems
according to embodiments of the present invention, using a traditional programming
language such as the C programming language, and using traditional programming methods
to write parallel communications routines that send and receive data among nodes on two
independent data communications networks. Alternatively, existing prior art libraries may
25 be improved to operate according to embodiments of the present invention. Examples of
prior-art parallel communications libraries include the 'Message Passing Interface' ('MPI')
library and the 'Parallel Virtual Machine' ('PVM') library.

30 Also stored in RAM is a hierarchical distributed compiler (155), a module of automated
computing machinery capable of compiling software for a hierarchical distributed processing
system according to embodiments of the present invention. The hierarchical distributed
compiler (155) includes computer program instruction for compiling, by the compiling node,

the software; maintaining, by the compiling node, any compiled software to be executed on the compiling node; selecting, by the compiling node, one or more nodes in a next tier of the hierarchy of the distributed processing system in dependence upon whether any compiled software is for the selected node or the selected node's descendants; sending to the selected node only the compiled software to be executed by the selected node or selected node's descendant. Each of the other compute nodes in a parallel computer are also capable of receiving compiled software; determining whether the compiled software is for that node or one of its descendants; if the compiled software is for that node, maintaining the software for execution; and if the compiled software is for one of the descendants, selecting another node in a next tier of the hierarchical distributed processing system in dependence upon a descendant for the compiled software and sending the compiled software to the selected another node.

Also stored in RAM (156) is an operating system (162), a module of computer program instructions and routines for an application program's access to other resources of the compute node. It is typical for an application program and parallel communications library in a compute node of a parallel computer to run a single thread of execution with no user login and no security issues because the thread is entitled to complete access to all resources of the node. The quantity and complexity of tasks to be performed by an operating system on a compute node in a parallel computer therefore are smaller and less complex than those of an operating system on a serial computer with many threads running simultaneously. In addition, there is no video I/O on the compute node (152) of Figure 2, another factor that decreases the demands on the operating system. The operating system may therefore be quite lightweight by comparison with operating systems of general purpose computers, a pared down version as it were, or an operating system developed specifically for operations on a particular parallel computer. Operating systems that may usefully be improved, simplified, for use in a compute node include UNIX_{TM}, Linux_{TM}, Microsoft XP_{TM}, AIX_{TM}, IBM's i5/OS_{TM}, and others as will occur to those of skill in the art.

The exemplary compute node (152) of Figure 2 includes several communications adapters (172, 176, 180, 188) for implementing data communications with other nodes of a parallel computer. Such data communications may be carried out serially through RS-232

connections, through external buses such as Universal Serial Bus ('USB'), through data communications networks such as IP networks, and in other ways as will occur to those of skill in the art. Communications adapters implement the hardware level of data communications through which one computer sends data communications to another
5 computer, directly or through a network. Examples of communications adapters useful in systems that for compiling software for a hierarchical distributed processing system according to embodiments of the present invention include modems for wired communications, Ethernet (IEEE 802.3) adapters for wired network communications, and 802.11b adapters for wireless network communications.

10

The data communications adapters in the example of Figure 2 include a Gigabit Ethernet adapter (172) that couples example compute node (152) for data communications to a Gigabit Ethernet (174). Gigabit Ethernet is a network transmission standard, defined in the IEEE 802.3 standard, that provides a data rate of 1 billion bits per second (one gigabit).
15 Gigabit Ethernet is a variant of Ethernet that operates over multimode fiber optic cable, single mode fiber optic cable, or unshielded twisted pair.

20

The data communications adapters in the example of Figure 2 include a JTAG Slave circuit (176) that couples example compute node (152) for data communications to a JTAG Master circuit (178). JTAG is the usual name used for the IEEE 1149.1 standard entitled Standard Test Access Port and Boundary-Scan Architecture for test access ports used for testing printed circuit boards using boundary scan. JTAG is so widely adapted that, at this time, boundary scan is more or less synonymous with JTAG. JTAG is used not only for printed circuit boards, but also for conducting boundary scans of integrated circuits, and is also
25 useful as a mechanism for debugging embedded systems, providing a convenient "back door" into the system. The example compute node of Figure 2 may be all three of these: It typically includes one or more integrated circuits installed on a printed circuit board and may be implemented as an embedded system having its own processor, its own memory, and its own I/O capability. JTAG boundary scans through JTAG Slave (176) may efficiently
30 configure processor registers and memory in compute node (152) for use in compiling software for a hierarchical distributed processing system according to embodiments of the present invention.

The data communications adapters in the example of Figure 2 includes a Point To Point Adapter (180) that couples example compute node (152) for data communications to a network (108) that is optimal for point to point message passing operations such as, for example, a network configured as a three-dimensional torus or mesh. Point To Point Adapter (180) provides data communications in six directions on three communications axes, x, y, and z, through six bidirectional links: +x (181), -x (182), +y (183), -y (184), +z (185), and -z (186).

The data communications adapters in the example of Figure 2 includes a Global Combining Network Adapter (188) that couples example compute node (152) for data communications to a network (106) that is optimal for collective message passing operations on a global combining network configured, for example, as a binary tree. The Global Combining Network Adapter (188) provides data communications through three bidirectional links: two to children nodes (190) and one to a parent node (192).

Example compute node (152) includes two arithmetic logic units ('ALUs'). ALU (166) is a component of each processing core (164), and a separate ALU (170) is dedicated to the exclusive use of Global Combining Network Adapter (188) for use in performing the arithmetic and logical functions of reduction operations. Computer program instructions of a reduction routine in parallel communications library (160) may latch an instruction for an arithmetic or logical function into instruction register (169). When the arithmetic or logical function of a reduction operation is a 'sum' or a 'logical or,' for example, Global Combining Network Adapter (188) may execute the arithmetic or logical operation by use of ALU (166) in processor (164) or, typically much faster, by use dedicated ALU (170).

The example compute node (152) of Figure 2 includes a direct memory access ('DMA') controller (195), which is computer hardware for direct memory access and a DMA engine (197), which is computer software for direct memory access. The DMA engine (197) of Figure 2 is typically stored in computer memory of the DMA controller (195). Direct memory access includes reading and writing to memory of compute nodes with reduced operational burden on the central processing units (164). A DMA transfer essentially copies

a block of memory from one location to another, typically from one compute node to another. While the CPU may initiate the DMA transfer, the CPU does not execute it.

5 For further explanation, Figure 3A illustrates an exemplary Point To Point Adapter (180) useful in systems capable of compiling software for a hierarchical distributed processing system according to embodiments of the present invention. Point To Point Adapter (180) is designed for use in a data communications network optimized for point to point operations, a network that organizes compute nodes in a three-dimensional torus or mesh. Point To Point Adapter (180) in the example of Figure 3A provides data communication along an x-axis
10 through four unidirectional data communications links, to and from the next node in the -x direction (182) and to and from the next node in the +x direction (181). Point To Point Adapter (180) also provides data communication along a y-axis through four unidirectional data communications links, to and from the next node in the -y direction (184) and to and from the next node in the +y direction (183). Point To Point Adapter (180) in Figure 3A also
15 provides data communication along a z-axis through four unidirectional data communications links, to and from the next node in the -z direction (186) and to and from the next node in the +z direction (185).

20 For further explanation, Figure 3B illustrates an exemplary Global Combining Network Adapter (188) useful in systems capable of compiling software for a hierarchical distributed processing system according to embodiments of the present invention. Global Combining Network Adapter (188) is designed for use in a network optimized for collective operations, a network that organizes compute nodes of a parallel computer in a binary tree. Global Combining Network Adapter (188) in the example of Figure 3B provides data
25 communication to and from two children nodes through four unidirectional data communications links (190). Global Combining Network Adapter (188) also provides data communication to and from a parent node through two unidirectional data communications links (192).

30 For further explanation, Figure 4 sets forth a line drawing illustrating an exemplary data communications network (108) optimized for point to point operations useful in systems capable of compiling software for a hierarchical distributed processing system in accordance

with embodiments of the present invention. In the example of Figure 4, dots represent compute nodes (102) of a parallel computer, and the dotted lines between the dots represent data communications links (103) between compute nodes. The data communications links are implemented with point to point data communications adapters similar to the one

5 illustrated for example in Figure 3A, with data communications links on three axes, x, y, and z, and to and fro in six directions +x (181), -x (182), +y (183), -y (184), +z (185), and -z (186). The links and compute nodes are organized by this data communications network optimized for point to point operations into a three dimensional mesh (105). The mesh (105) has wrap-around links on each axis that connect the outermost compute nodes in the mesh

10 (105) on opposite sides of the mesh (105). These wrap-around links form part of a torus (107). Each compute node in the torus has a location in the torus that is uniquely specified by a set of x, y, z coordinates. Readers will note that the wrap-around links in the y and z directions have been omitted for clarity, but are configured in a similar manner to the wrap-around link illustrated in the x direction. For clarity of explanation, the data

15 communications network of Figure 4 is illustrated with only 27 compute nodes, but readers will recognize that a data communications network optimized for point to point operations for use in compiling software for a hierarchical distributed processing system in accordance with embodiments of the present invention may contain only a few compute nodes or may contain thousands of compute nodes.

20 For further explanation, Figure 5 sets forth a line drawing illustrating an exemplary data communications network (106) optimized for collective operations useful in systems capable of compiling software for a hierarchical distributed processing system in accordance with embodiments of the present invention. The example data communications network of Figure

25 5 includes data communications links connected to the compute nodes so as to organize the compute nodes as a tree. In the example of Figure 5, dots represent compute nodes (102) of a parallel computer, and the dotted lines (103) between the dots represent data communications links between compute nodes. The data communications links are implemented with global combining network adapters similar to the one illustrated for

30 example in Figure 3B, with each node typically providing data communications to and from two children nodes and data communications to and from a parent node, with some exceptions. Nodes in a binary tree (106) may be characterized as a physical root node (202),

branch nodes (204), and leaf nodes (206). The root node (202) has two children but no parent. The leaf nodes (206) each has a parent, but leaf nodes have no children. The branch nodes (204) each has both a parent and two children. The links and compute nodes are thereby organized by this data communications network optimized for collective operations into a binary tree (106). For clarity of explanation, the data communications network of Figure 5 is illustrated with only 31 compute nodes, but readers will recognize that a data communications network optimized for collective operations for use in systems for compiling software for a hierarchical distributed processing system in accordance with embodiments of the present invention may contain only a few compute nodes or may contain thousands of compute nodes.

In the example of Figure 5, each node in the tree is assigned a unit identifier referred to as a 'rank' (250). A node's rank uniquely identifies the node's location in the tree network for use in both point to point and collective operations in the tree network. The ranks in this example are assigned as integers beginning with 0 assigned to the root node (202), 1 assigned to the first node in the second layer of the tree, 2 assigned to the second node in the second layer of the tree, 3 assigned to the first node in the third layer of the tree, 4 assigned to the second node in the third layer of the tree, and so on. For ease of illustration, only the ranks of the first three layers of the tree are shown here, but all compute nodes in the tree network are assigned a unique rank.

For further explanation, Figure 6 sets forth a further exemplary distributed computing system for compiling software for a hierarchical distributed processing system according to embodiments of the present invention in which the distributed computing system is implemented as a hybrid computing environment. A 'hybrid computing environment,' as the term is used in this specification, is a computing environment in that it includes computer processors operatively coupled to computer memory so as to implement data processing in the form of execution of computer program instructions stored in the memory and executed on the processors. The hybrid computing environment (600) of Figure 6 includes one compute node (603) that represents a small, separate hybrid computing environment which, when taken with other similar compute nodes (602), together make up a larger hybrid computing environment.

The example compute node (603) of Figure 6 may carry out principal user-level computer program execution, accepting administrative services, such as initial program loads and the like, from a service application executing on a service node connected to the compute node (603) through a data communications network. The example compute node may also be
5 coupled for data communications to one or more input/output (I/O) nodes that enable the compute node to gain access to data storage and other I/O functionality. The I/O nodes and service node may be connected to the example compute node (603), to other compute nodes (602) in the larger hybrid computing environment, and to I/O devices, through a local area network ('LAN') implemented using high-speed Ethernet or a data communications fabric of
10 another fabric type as will occur to those of skill in the art. I/O devices useful in a larger hybrid computing environment that includes the compute node (603) may include non-volatile memory for the computing environment in the form of data storage device, an output device for the hybrid computing environment in the form of printer, and a user I/O device in the form of computer terminal that executes a service application interface that provides to a
15 user an interface for configuring compute nodes in the hybrid computing environment and initiating execution by the compute nodes of principal user-level computer program instructions.

The compute node (603) in the example of Figure 6 is illustrated in an expanded view to aid
20 a more detailed explanation of a hybrid computing environment (600) that may be combined with other hybrid computing environments, such as the other compute nodes (602), to form a larger hybrid computing environment. The compute node (603) in the example of Figure 6 includes a host computer (610). A host computer (610) is a 'host' in the sense that it is the host computer that carries out interface functions between a compute node and other
25 components of the hybrid computing environment external to any particular compute node. That is, it is the host computer that executes initial boot procedures, power on self tests, basic I/O functions, accepts user-level program loads from service nodes, and so on.

The host computer (610) in the example of Figure 6 includes a computer processor (652)
30 operatively coupled to computer memory, Random Access Memory ('RAM') (642), through a high speed memory bus (653). The processor (652) in each host computer (610) has a set of architectural registers (654) that defines the host computer architecture.

The example compute node (603) of Figure 6 also includes one or more accelerators (604, 605). An accelerator (604) is an 'accelerator' in that each accelerator has an accelerator architecture that is optimized, with respect to the host computer architecture, for speed of execution of a particular class of computing functions. Such accelerated computing functions include, for example, vector processing, floating point operations, and others as will occur to those of skill in the art. Each accelerator (604, 605) in the example of Figure 6 includes a computer processor (648) operatively coupled to RAM (640) through a high speed memory bus (651). Stored in RAM (640, 642) of the host computer and the accelerators (604, 605) is an operating system (645). Operating systems useful in host computers and accelerators of hybrid computing environments according to embodiments of the present invention include UNIX_{TM}, Linux_{TM}, Microsoft XP_{TM}, Microsoft Vista_{TM}, Microsoft NT_{TM}, AIX_{TM}, IBM's i5/OS_{TM}, and others as will occur to those of skill in the art. There is no requirement that the operating system in the host computers should be the same operating system used on the accelerators.

The processor (648) of each accelerator (604, 605) has a set of architectural registers (650) that defines the accelerator architecture. The architectural registers (650) of the processor (648) of each accelerator are different from the architectural registers (654) of the processor (652) in the host computer (610). The architectural registers are registers that are accessible by computer program instructions that execute on each architecture, registers such as an instruction register, a program counter, memory index registers, stack pointers, and the like. With differing architectures, it would be uncommon, although possible, for a host computer and an accelerator to support the same instruction sets. As such, computer program instructions compiled for execution on the processor (648) of an accelerator (604) generally would not be expected to execute natively on the processor (652) of the host computer (610) and vice versa. Moreover, because of the typical differences in hardware architectures between host processors and accelerators, computer program instructions compiled for execution on the processor (652) of a host computer (610) generally would not be expected to execute natively on the processor (648) of an accelerator (604) even if the accelerator supported the instruction set of the host. The accelerator architecture in example of Figure 6 is optimized, with respect to the host computer architecture, for speed of execution of a particular class of computing functions. That is, for the function or functions for which the

accelerator is optimized, execution of those functions will proceed faster on the accelerator than if they were executed on the processor of the host computer.

5 Examples of hybrid computing environments include a data processing system that in turn includes one or more host computers, each having an x86 processor, and accelerators whose architectural registers implement the PowerPC instruction set. Computer program instructions compiled for execution on the x86 processors in the host computers cannot be executed natively by the PowerPC processors in the accelerators. Readers will recognize in addition that some of the example hybrid computing environments described in this
10 specification are based upon the Los Alamos National Laboratory ('LANL') supercomputer architecture developed in the LANL Roadrunner project (named for the state bird of New Mexico), the supercomputer architecture that famously first generated a 'petaflop,' a million billion floating point operations per second. The LANL supercomputer architecture includes many host computers with dual-core AMD Opteron processors coupled to many accelerators
15 with IBM Cell processors, the Opteron processors and the Cell processors having different architectures.

In the example of Figure 6, the host computer (610) and the accelerators (604, 605) are adapted to one another for data communications by a system level message passing module
20 ('SLMPM') (646) and two data communications fabrics (628, 630) of at least two different fabric types. A data communications fabric (628, 630) is a configuration of data communications hardware and software that implements a data communications coupling between a host computer and an accelerator. Examples of data communications fabric types include Peripheral Component Interconnect ('PCI'), PCI express ('PCIe'), Ethernet,
25 Infiniband, Fibre Channel, Small Computer System Interface ('SCSI'), External Serial Advanced Technology Attachment ('eSATA'), Universal Serial Bus ('USB'), and so on as will occur to those of skill in the art. In the example of Figure 6, the host computer (610) and the accelerators (604, 605) are adapted to one another for data communications by a PCIe fabric (630) through PCIe communications adapters (660) and an Ethernet fabric (628)
30 through Ethernet communications adapters (661). The use of PCIe and Ethernet is for explanation, not for limitation of the invention. Readers of skill in the art will immediately recognize that hybrid computing environments according to embodiments of the present

invention may include fabrics of other fabric types such as, for example, PCI, Infiniband, Fibre Channel, SCSI, eSATA, USB, and so on.

5 An SLMPM (646) is a module or library of computer program instructions that exposes an application programming interface ('API') to user-level applications for carrying out message-based data communications between the host computer (610) and the accelerator (604, 605). Examples of message-based data communications libraries that may be improved for use as an SLMPM according to embodiments of the present invention include:

10 • the Message Passing Interface or 'MPI,' an industry standard interface in two versions, first presented at Supercomputing 1994, not sanctioned by any major standards body,

15 • the Data Communication and Synchronization interface ('DACS') of the LANL supercomputer,

• the POSIX Threads library ('Pthreads'), an IEEE standard for distributed, multithreaded processing,

20 • the Open Multi-Processing interface ('OpenMP'), an industry-sanctioned specification for parallel programming, and

• other libraries that will occur to those of skill in the art.

25 In this example, to support message-based data communications between the host computer (610) and the accelerator (604), both the host computer (610) and the accelerator (604) have an SLMPM (646) so that message-based communications can both originate and be received on both sides of any coupling for data communications.

30 The SLMPM (646) in this example operates generally for data processing in a hybrid computing environment (600) by monitoring data communications performance for a plurality of data communications modes between the host computer (610) and the

accelerators (604, 605), receiving a request (668) to transmit data according to a data communications mode from the host computer to an accelerator, determining whether to transmit the data according to the requested data communications mode, and if the data is not to be transmitted according to the requested data communications mode: selecting
5 another data communications mode and transmitting the data according to the selected data communications mode. In the example of Figure 6, the monitored performance is illustrated as monitored performance data (674) stored by the SLMPM (646) in RAM (642) of the host computer (610) during operation of the compute node (603).

10 A data communications mode specifies a data communications fabric type, a data communications link, and a data communications protocol (678). A data communications link (656) is data communications connection between a host computer and an accelerator. In the example of Figure 6, a link (656) between the host computer (610) and the accelerator (604) may include the PCIe connection (638) or the Ethernet connection (631, 632) through
15 the Ethernet network (606). A link (656) between the host computer (610) and the accelerator (605) in the example of Figure 6, may include the PCIe connection (636) or the Ethernet connection (631, 634) through the Ethernet network (606). Although only one link for each fabric type is illustrated between the host computer and the accelerator in the example of Figure 6, readers of skill in the art will immediately recognize that there may any
20 number of links for each fabric type.

A data communications protocol is a set of standard rules for data representation, signaling, authentication and error detection required to send information from a host computer (610) to an accelerator (604). In the example of Figure 6, the SLMPM (646) may select one of
25 several protocols (678) for data communications between the host computer (610) and the accelerator. Examples of such protocols (678) include shared memory transfers ('SMT') (680) executed with a send and receive operations (681), and direct memory access ('DMA') (682) executed with PUT and GET operations (683).

30 Shared memory transfer is a data communications protocol for passing data between a host computer and an accelerator into shared memory space (658) allocated for such a purpose such that only one instance of the data resides in memory at any time. Consider the

following as an example shared memory transfer between the host computer (610) and the accelerator (604) of Figure 6. An application (669) requests (668) a transmission of data (676) from the host computer (610) to the accelerator (604) in accordance with the SMT (680) protocol. Such a request (668) may include a memory address allocated for such shared memory. In this example, the shared memory segment (658) is illustrated in a memory location on the accelerator (604), but readers will recognize that shared memory segments may be located on the accelerator (604), on the host computer (610), on both the host computer and the accelerator, or even off the local compute node (603) entirely – so long as the segment is accessible as needed by the host and the accelerator. To carry out a shared memory transfer, the SLMPM (646) on the host computer (610) establishes a data communications connection with the SLMPM (646) executing on the accelerator (604) by a handshaking procedure similar to that in the TCP protocol. The SLMPM (646) then creates a message (670) that includes a header and a payload data and inserts the message into a message transmit queue for a particular link of a particular fabric. In creating the message, the SLMPM inserts, in the header of the message, an identification of the accelerator and an identification of a process executing on the accelerator. The SLMPM also inserts the memory address from the request (668) into the message, either in the header or as part of the payload data. The SLMPM also inserts the data (676) to be transmitted in the message (670) as part of the message payload data. The message is then transmitted by a communications adapter (660, 661) across a fabric (628, 630) to the SLMPM executing on the accelerator (604) where the SLMPM stores the payload data, the data (676) that was transmitted, in shared memory space (658) in RAM (640) in accordance with the memory address in the message.

Direct memory access ('DMA') is a data communications protocol for passing data between a host computer and an accelerator with reduced operational burden on the computer processor (652). A DMA transfer essentially effects a copy of a block of memory from one location to another, typically from a host computer to an accelerator or vice versa. Either or both a host computer and accelerator may include DMA controller and DMA engine, an aggregation of computer hardware and software for direct memory access. Direct memory access includes reading and writing to memory of accelerators and host computers with reduced operational burden on their processors. A DMA engine of an accelerator, for

example, may write to or read from memory allocated for DMA purposes, while the processor of the accelerator executes computer program instructions, or otherwise continues to operate. That is, a computer processor may issue an instruction to execute a DMA transfer, but the DMA engine, not the processor, carries out the transfer.

5

In the example of Figure 6, only the accelerator (604) includes a DMA controller (685) and DMA engine (684) while the host computer does not. In this embodiment the processor (652) on the host computer initiates a DMA transfer of data from the host to the accelerator by sending a message according to the SMT protocol to the accelerator, instructing the accelerator to perform a remote 'GET' operation. The configuration illustrated in the example of Figure 6 in which the accelerator (604) is the only device containing a DMA engine is for explanation only, not for limitation. Readers of skill in the art will immediately recognize that in many embodiments, both a host computer and an accelerator may include a DMA controller and DMA engine, while in yet other embodiments only a host computer includes a DMA controller and DMA engine.

10

15

To implement a DMA protocol in the hybrid computing environment of Figure 6 some memory region is allocated for access by the DMA engine. Allocating such memory may be carried out independently from other accelerators or host computers, or may be initiated by and completed in cooperation with another accelerator or host computer. Shared memory regions, allocated according to the SMA protocol, for example, may be memory regions made available to a DMA engine. That is, the initial setup and implementation of DMA data communications in the hybrid computing environment (600) of Figure 6 may be carried out, at least in part, through shared memory transfers or another out-of-band data communications protocol, out-of-band with respect to a DMA engine. Allocation of memory to implement DMA transfers is relatively high in latency, but once allocated, the DMA protocol provides for high bandwidth data communications that requires less processor utilization than many other data communications protocols.

20

25

30

A direct 'PUT' operation is a mode of transmitting data from a DMA engine on an origin device to a DMA engine on a target device. A direct 'PUT' operation allows data to be transmitted and stored on the target device with little involvement from the target device's

processor. To effect minimal involvement from the target device's processor in the direct 'PUT' operation, the origin DMA engine transfers the data to be stored on the target device along with a specific identification of a storage location on the target device. The origin DMA knows the specific storage location on the target device because the specific storage location for storing the data on the target device has been previously provided by the target DMA engine to the origin DMA engine.

A remote 'GET' operation, sometimes denominated an 'rGET,' is another mode of transmitting data from a DMA engine on an origin device to a DMA engine on a target device. A remote 'GET' operation allows data to be transmitted and stored on the target device with little involvement from the origin device's processor. To effect minimal involvement from the origin device's processor in the remote 'GET' operation, the origin DMA engine stores the data in an storage location accessible by the target DMA engine, notifies the target DMA engine, directly or out-of-band through a shared memory transmission, of the storage location and the size of the data ready to be transmitted, and the target DMA engine retrieves the data from storage location.

Monitoring data communications performance for a plurality of data communications modes may include monitoring a number of requests (668) in a message transmit request queue (662-165) for a data communications link (656). In the example of Figure 6, each message transmit request queue (662-165) is associated with one particular data communications link (656). Each queue (662-165) includes entries for messages (670) that include data (676) to be transmitted by the communications adapters (660, 661) along a data communications link (656) associated with queue.

Monitoring data communications performance for a plurality of data communications modes may also include monitoring utilization of a shared memory space (658). In the example of Figure 6, shared memory space (658) is allocated in RAM (640) of the accelerator.

Utilization is the proportion of the allocated shared memory space to which data has been stored for sending to a target device and has not yet been read or received by the target device, monitored by tracking the writes and reads to and from the allocated shared memory. In the hybrid computing environment (600) of Figure 6, shared memory space, any memory

in fact, is limited. As such, a shared memory space (658) may be filled during execution of an application program (669) such that transmission of data from the host computer (610) to an accelerator may be slowed, or even stopped, due to space limitations in the shared memory space.

5

In some embodiments of the present invention, the hybrid computing environment (600) of Figure 6 may be configured to operate as a parallel computing environment in which two or more instances the application program (669) executes on two or more host computers (610) in the parallel computing environment. In such embodiments, monitoring data
10 communications performance across data communications modes may also include aggregating data communications performance information (674) across a plurality of instances of the application program (669) executing on two or more host computers in a parallel computing environment. The aggregated performance information (674) may be used to calculate average communications latencies for data communications modes, average
15 number of requests in data communications links of a particular fabric type, average shared memory utilization among the plurality of host computers and accelerators in the parallel computing environment, and so on as will occur to those of skill in the art. Any combination of such measures may be used by the SLMPM for both determining whether to transmit the data according to requested data communications mode and selecting another data
20 communications mode for transmitting the data if the data is not to be transmitted according to the requested data communications mode.

20

The SLMPM (646) of Figure 6 receives, from an application program (669) on the host computer (610), a request (668) to transmit data (676) according to a data communications
25 mode from the host computer (610) to the accelerator (604). Such data (676) may include computer program instructions compiled for execution by the accelerator (604), such as an executable file of an accelerator application program, work piece data for an accelerator application program, files necessary for execution of an accelerator application program, such as libraries, databases, drivers, and the like. Receiving a request (668) to transmit data
30 (676) according to a data communications mode may include receiving a request to transmit data by a specified fabric type, receiving a request to transmit data through a specified data

30

communications link from the host computer to the accelerator, or receiving a request to transmit data from the host computer to the accelerator according to a protocol.

5 A request (668) to transmit data (676) according to a data communications mode may be implemented as a user-level application function call through an API to the SLMPM (646), a call that expressly specifies a data communications mode according to protocol, fabric type, and link. A request implemented as a function call may specify a protocol according to the operation of the function call itself. A `dacs_put()` function call, for example, may represent a call through an API exposed by an SLMPM implemented as a DACS library to transmit data
10 in the default mode of a DMA 'PUT' operation. Such a call, from the perspective of the calling application and the programmer who wrote the calling application, represents a request to the SLMPM library to transmit data according to the default mode, known to the programmer to be default mode associated with the express API call. The called function, in this example `dacs_put()`, may be coded in embodiments with multiple fabric types, protocols,
15 and links, to make its own determination whether to transmit the data according to the requested data communications mode, that is, according to the default mode of the called function. In a further example, a `dacs_send()` instruction may represent a call through an API exposed by an SLMPM implemented as a DACS library to transmit data in the default mode of an SMT 'send' operation, where the called function `dacs_send()` is again coded in
20 embodiments with multiple fabric types, protocols, and links, to make its own determination whether to transmit the data according to the requested mode.

An identification of a particular accelerator in a function call may effectively specify a fabric type. Such a function call may include as a call parameters an identification of a particular
25 accelerator. An identification of a particular accelerator by use of a PCIe ID, for example, effectively specifies a PCI fabric type. In another, similar, example, an identification of a particular accelerator by use of a media access control ('MAC') address of an Ethernet adapter effectively specifies the Ethernet fabric type. Instead of implementing the accelerator ID of the function call from an application executing on the host in such a way as
30 to specify a fabric type, the function call may only include a globally unique identification of the particular accelerator as a parameter of the call, thereby specifying only a link from the host computer to the accelerator, not a fabric type. In this case, the function called may

implement a default fabric type for use with a particular protocol. If the function called in the SLMPM is configured with PCIe as a default fabric type for use with the DMA protocol, for example, and the SLMPM receives a request to transmit data to the accelerator (604) according to the DMA protocol, a DMA PUT or DMA remote GET operation, the function
5 called explicitly specifies the default fabric type for DMA, the PCIe fabric type.

In hybrid computing environments in which only one link of each fabric type adapts a single host computer to a single accelerator, the identification of a particular accelerator in a parameter of a function call, may also effectively specify a link. In hybrid computing
10 environments where more than one link of each fabric type adapts a host computer and an accelerator, such as two PCIe links connecting the host computer (610) to the accelerator (604), the SLMPM function called may implement a default link for the accelerator identified in the parameter of the function call for the fabric type specified by the identification of the accelerator.

The SLMPM (646) in the example of Figure 6 also determines, in dependence upon the monitored performance (674), whether to transmit the data (676) according to the requested data communications mode. Determining whether to transmit the data (676) according to the requested data communications mode may include determining whether to transmit data
15 20 by a requested fabric type, whether to transmit data through a requested data communications link, or whether to transmit data according to a requested protocol.

In hybrid computing environments according to embodiments of the present invention, where monitoring data communications performance across data communications modes
25 includes monitoring a number of requests in a message transmit request queue (662-165) for a data communications link, determining whether to transmit the data (676) according to the requested data communications mode may be carried out by determining whether the number of requests in the message transmit request queue exceeds a predetermined threshold. In hybrid computing environments according to embodiments of the present
30 invention, where monitoring data communications performance for a plurality of data communications modes includes monitoring utilization of a shared memory space, determining whether to transmit the data (676) according to the requested data

communications mode may be carried out by determining whether the utilization of the shared memory space exceeds a predetermined threshold.

5 If the data is not to be transmitted according to the requested data communications mode, the SLMPM (646) selects, in dependence upon the monitored performance, another data communications mode for transmitting the data and transmits the data (676) according to the selected data communications mode. Selecting another data communications mode for transmitting the data may include selecting, in dependence upon the monitored performance, another data communications fabric type by which to transmit the data, selecting a data
10 communications link through which to transmit the data, and selecting another data communications protocol. Consider as an example, that the requested data communications mode is a DMA transmission using a PUT operation through link (638) of the PCIe fabric (630) to the accelerator (604). If the monitored data performance (674) indicates that the number of requests in transmit message request queue (662) associated with the link (638)
15 exceeds a predetermined threshold, the SLMPM may select another fabric type, the Ethernet fabric (628), and link (631, 632) through which to transmit the data (676). Also consider that the monitored performance (676) indicates that current utilization of the shared memory space (658) is less than a predetermined threshold while the number of outstanding DMA transmissions in the queue (662) exceeds a predetermined threshold. In such a case, the
20 SLMPM (646) may also select another protocol, such as a shared memory transfer, by which to transmit the data (674).

Selecting, by the SLMPM, another data communications mode for transmitting the data (672) may also include selecting a data communications protocol (678) in dependence upon
25 data communications message size (672). Selecting a data communications protocol (678) in dependence upon data communications message size (672) may be carried out by determining whether a size of a message exceeds a predetermined threshold. For larger messages (670), the DMA protocol may be a preferred protocol as processor utilization in making a DMA transfer of a larger message (670) is typically less than the processor
30 utilization in making a shared memory transfer of a message of the same size.

As mentioned above, the SLMPM may also transmit the data according to the selected data communications mode. Transmit the data according to the selected data communications mode may include transmitting the data by the selected data communications fabric type, transmitting the data through the selected data communications link, or transmitting the data according to the selected protocol. The SLMPM (646) may effect a transmission of the data according to the selected data communications mode by instructing, through a device driver, the communications adapter for the data communications fabric type of the selected data communications mode to transmit the message (670) according to a protocol of the selected data communications mode, where the message includes in a message header, an identification of the accelerator, and in the message payload, the data (676) to be transmitted.

In the example of Figure 6, both the host computer (610) and the accelerator (604) one of the compute nodes have installed upon it a hierarchical distributed compiler (155). The hierarchical distributed compiler (155) is shown in both the host computer and accelerator for completeness. In fact, hierarchical distributed compiler (155) according to embodiments of the present invention may be installed on either a host computer or one or more accelerators, or both a host computer and one or more accelerators as will occur to those of skill in the art. The hierarchical distributed compiler (155) of Figure 6 is a module of automated computing machinery capable of compiling software for a hierarchical distributed processing system according to embodiments of the present invention. The hierarchical distributed compiler (155) includes computer program instruction for compiling, by the compiling node, the software; maintaining, by the compiling node, any compiled software to be executed on the compiling node; selecting, by the compiling node, one or more nodes in a next tier of the hierarchy of the distributed processing system in dependence upon whether any compiled software is for the selected node or the selected node's descendants; sending to the selected node only the compiled software to be executed by the selected node or selected node's descendant. Each of the other compute nodes (602) of Figure 16 are also capable of receiving compiled software; determining whether the compiled software is for that node or one of its descendants; if the compiled software is for that node, maintaining the software for execution; and if the compiled software is for one of the descendants, selecting another node in a next tier of the hierarchical distributed processing system in dependence

upon a descendant for the compiled software and sending the compiled software to the selected another node.

5 For further explanation, Figure 7 sets forth an exemplary method of compiling software for a hierarchical distributed processing system according to embodiments of the present invention. Compiling is the process of transforming source code written in a computer language, often a high level programming language, into another typically executable computer language, typically having a binary form and sometimes called object code. Compiling according to embodiments of the present invention is typically carried out with a hierarchical distributed compiler installed on compiling nodes according to the present invention. Such a hierarchical distributed compiler is typically able to compile software for use on a number of different kinds of target computers. As such, a hierarchical distributed compiler may compile portions of uncompiled source software to target object code for use on computers of disparate types.

15 A hierarchical distributed processing system may be implemented in a number of ways such as, for example, implemented in a tree structure. Such a tree structure may be k-ary, that is of any order, binary, or in any other form of tree structure that will occur to those of skill in the art. Alternatively, hierarchical distributed processing systems according to the present invention may be implemented in other forms that are not considered tree structures as will occur to those of skill in the art. The method of Figure 7 may be carried out in a distributed computing system similar to the example distributed computing systems described above: the example parallel computers of Figures 1-5, the example hybrid computing environment of Figure 6, and others as will occur to those of skill in the art.

25 The method of Figure 7 includes identifying (802) the one or more compiling nodes. As mentioned above, the one or more compiling nodes transform source code written in a computer language, typically a high level programming language, into another often executable computer language, typically having a binary form and sometimes called object code. The compiling nodes compile portions of the software to be executed on the compiling nodes themselves, as well as other portions of the software to be executed on another node in the hierarchical distributed network.

Identifying (802) the one or more compiling nodes according to the method of Figure 7 may be carried out by selecting one or more nodes that are computationally optimized for compiling. Selecting one or more nodes that are computationally optimized for compiling may include identifying nodes in dependence upon their I/O capabilities, processing
5 capabilities, and memory capabilities. Often a particular balance of these capabilities is optimal for compiling. Such an optimal balance may be different for compiling different kinds of software and as such, selecting one or more nodes that are computationally optimized for compiling may include selecting one or more nodes in dependence upon the particular software to be compiled.

10 Identifying (802) the one or more compiling nodes according to the method of Figure 7 may also be carried out by selecting one or more nodes that are optimized for compiling because of their location in the topology of the hierarchical distributed processing system. In a tree structured hierarchical data processing system, for example, a node that is a root node, or a
15 node having many descendants may be located in the topology such that node is optimized for compiling software for its descendants. A descendant as that term is used in this specification is a node located in tiers of the hierarchical data processing system below the compiling node and on a branch of the hierarchical processing system that includes the compiling node.

20 The method of Figure 7 also includes providing (804) to one or more compiling nodes software to be compiled, wherein at least a portion of the software to be compiled is to be executed by one or more other nodes. Providing (804) to one or more compiling nodes software to be compiled may be carried out by sending the software to be compiled to the
25 compiling node in a message, downloading the software to the compiling node, installing the software on the compiling node by a systems administrator, or any other way of providing to one or more compiling nodes software to be compiled that will occur to those of skill in the art.

30 The method of Figure 7 also includes compiling (806), by the compiling node, the software. Compiling (806), by the compiling node, the software may be carried out by identifying portions of the software to be executed on the compiling node and transforming the code

written in a computer language of the uncompiled software to an executable computer language for execution on the compiling node. Compiling (806), by the compiling node, the software may be carried out by identifying portions of the software to be executed on another node, identifying the target execution environment of the another node, and
5 transforming the software written in the computer language of the uncompiled software to an executable computer language for execution on the another node.

The method of Figure 7 also includes maintaining (808), by the compiling node, any compiled software to be executed on the compiling node. Maintaining (808), by the
10 compiling node, any compiled software to be executed on the compiling node may be carried out by storing for execution any compiled software to be executed on the compiling node.

The method of Figure 7 also includes selecting (810), by the compiling node, one or more nodes in a next tier of the hierarchy of the distributed processing system in dependence upon
15 whether any compiled software is for the selected node or the selected node's descendants and sending (812) to the selected node only the compiled software to be executed by the selected node or selected node's descendants. Selecting (810), by the compiling node, one or more nodes in a next tier of the hierarchy of the distributed processing system in
20 dependence upon whether any compiled software is for the selected node or the selected node's descendants may be carried out by traversing a representation of the topology of the hierarchical distributed processing system to identify the location of nodes on which compiled software is to be executed, determining a branch of the hierarchical data processing system upon which those nodes to execute the compiled software reside,
25 indentifying which child node of the compiling node is also on that branch of the hierarchical data processing system.

Sending (812) to the selected node only the compiled software to be executed by the selected node or selected node's descendants may be carried out by creating a message including the portions the compiled software to be executed by the selected node or selected node's
30 descendants and sending the message to the selected node. Sending (812) to the selected node only the compiled software to be executed by the selected node or selected node's descendants also may be carried out by notifying the selected node of the location of the

compiled software for download to the node to execute the compiled software, or any other way of sending (812) to the selected node only the compiled software to be executed by the selected node or selected node's descendants that will occur to those of skill in the art.

5 The selected node may or may not execute the compiled software. That is, the compiled software may be executed by a node in a tier below the selected node. For further explanation, therefore, Figure 8 sets forth a flow chart illustrating another exemplary method of compiling software for a hierarchical distributed processing system according to
10 embodiments of the present invention. The method of Figure 8 is similar to the method of Figure 7 in that the method of Figure 8 includes identifying (802) the one or more compiling nodes; providing (804) to one or more compiling node software to be compiled, wherein at least a portion of the software to be compiled is to be executed by one or more other nodes; compiling (806), by the compiling node, the software; maintaining (808), by the compiling node, any compiled software to be executed on the compiling node; selecting (810), by the
15 compiling node, one or more nodes in a next tier of the hierarchy of the distributed processing system in dependence upon whether any compiled software is for the selected node or the selected node's descendants; and sending (812) to the selected node only the compiled software to be executed by the selected node or selected node's descendants.

20 The method of Figure 8 also includes additional steps carried out by nodes below the compiling node in the hierarchical distributed processing system. The method of Figure 8 includes receiving (814), by a selected node, compiled software. The compiled software may be received in a message, identified for download by the compiling node, or received in other ways as will occur to those of skill in the art.

25 The method of Figure 8 also includes determining (816) whether the compiled software is for the selected node or one of its descendants. Determining (816) whether the compiled software is for the selected node or one of its descendants may be carried out by receiving from the compiling node an identification of the nodes to execute specific portions of the
30 compiled software and determining whether the identified node is the selected node or whether the identified node is one of its descendants.

If the compiled software is for the selected node, the method of Figure 8 includes maintaining (818) the software by the selected node for execution. Maintaining (818) the software by the selected node for execution may be carried out by storing for execution any compiled software to be executed on the selected node.

5

If the compiled software is for one of the descendants, the method of Figure 8 includes selecting (820) another node in a next tier of the hierarchical distributed processing system in dependence upon a descendant for the compiled software and sending (822) the compiled software to the selected another node. Selecting (820) another node in a next tier of the hierarchical distributed processing system in dependence upon a descendant for the compiled software may be carried out by traversing a representation of the topology of the hierarchical distributed processing system to identify the location of nodes on which compiled software is to be executed, determining a branch of the hierarchical data processing system upon which those nodes to execute the compiled software reside, identifying which child node of the selected node node is also on that branch of the hierarchical data processing system.

10

15

Sending (822) the compiled software to the selected another node may be carried out by creating a message including the portions the compiled software to be executed by the selected another node or selected another node's descendants and sending the message to the selected another node. Sending (822) the compiled software to the selected another node also may be carried out by notifying the selected another node of the location of the compiled software for download to the node to execute the compiled software, or any other way of sending (822) the compiled software to the selected another node that will occur to those of skill in the art.

20

25

For further explanation, Figure 9 sets forth a block diagram of an exemplary use case of a system for compiling software for a hierarchical distributed processing system according to embodiments of the present invention. In the example of Figure 9 a compiling laptop (702) has uncompiled software (722). The uncompiled software (722) has portions of software (724) for execution by computer (704), portions of software (726 and 728) for execution by one or more Graphics Processing Unites ("GPUs") (704), portions of software (728, 730, 732, 734) for execution by compute nodes in the parallel computer (712), portions of

30

software (736) for execution by the hybrid computing environments front end node (714), portions of software (738, 740, and 742) for execution by host computers (718) in the hybrid computing system, and portions of software (744, 746, 748, and 750) for execution by accelerators (720) of the hybrid computing system.

5

In the example of Figure 9, the compiling laptop (702) compiles the portion of software (724) for computer (704) and sends that portion of compiled software to computer (704) for execution. In the example of Figure 9, the compiling laptop (702) also compiles the portions of software (726 and 728) for specific GPUs (708) and sends those portions of compiled software to computer (704) who in turn sends the portion of compiled software to the specific GPU that will execute that portion of compiled software.

10

In the example of Figure 9, the compiling laptop (702) compiles the portion of software (728, 730, 732, and 734) for specific compute nodes (712) of the parallel computer and sends those portions of compiled software to I/O Node (710) of the parallel computer who in turns sends the portions of the compiled software to the root node of the compute nodes (712). The root node of the compute nodes then identifies which child node has descendants that will execute the portions of the compiled software and sends to each child only the portions that will be executed by that child or its descendants. Each child then receives those portions determines whether it will execute the portion or whether one of its descendants will execute the portion and sends only those portions for its descendants to a child on the same branch as the descendant. In this manner, tier by tier, the portions of compiled software are sent to the specific compute node to execute that compiled software.

15

20

25

In the example of Figure 9, the compiling laptop (702) compiles the portion of software (736) for the hybrid computing environment front end node (714) and sends that portion of compiled software to hybrid computing environment front end node (714) for execution. In the example of Figure 9, the compiling laptop (702) also compiles the portions of software (726, 738, 740, and 742) for specific host computers (718) and sends those portions of compiled software to hybrid computing environment front end node (714) who in turn sends the portion of compiled software to the specific host computers that will execute that portion of compiled software. In the example of Figure 9, the compiling laptop (702) also compiles

30

the portions of software (744, 746, 748, and 750) for specific accelerators (720) of host computers (718) and sends those portions of compiled software to hybrid computing environment front end node (714) who in turn sends the portion of compiled software to the specific host computers for those accelerators who in turn send the portion of compiled software to the specific accelerator that will execute that portion of compiled software.

In the examples above, compiling software for a hierarchical distributed processing system has been generally discussed with a single compiling node. This is for explanation and not for limitation. In fact, in many embodiments of the present invention, more than one node may compile software for a hierarchical distributed processing system in accordance with the present invention.

As will be appreciated by one skilled in the art, aspects of the present invention may be embodied as a system, method or computer program product. Accordingly, aspects of the present invention may take the form of an entirely hardware embodiment, an entirely software embodiment (including firmware, resident software, micro-code, etc.) or an embodiment combining software and hardware aspects that may all generally be referred to herein as a "circuit," "module" or "system." Furthermore, aspects of the present invention may take the form of a computer program product embodied in one or more computer readable medium(s) having computer readable program code embodied thereon.

Any combination of one or more computer readable medium(s) may be utilized. The computer readable medium may be a computer readable signal medium or a computer readable storage medium. A computer readable storage medium may be, for example, but not limited to, an electronic, magnetic, optical, electromagnetic, infrared, or semiconductor system, apparatus, or device, or any suitable combination of the foregoing. More specific examples (a non-exhaustive list) of the computer readable storage medium would include the following: an electrical connection having one or more wires, a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), an optical fiber, a portable compact disc read-only memory (CD-ROM), an optical storage device, a magnetic storage device, or any suitable combination of the foregoing. In the context of this document, a

computer readable storage medium may be any tangible medium that can contain, or store a program for use by or in connection with an instruction execution system, apparatus, or device.

5 A computer readable signal medium may include a propagated data signal with computer readable program code embodied therein, for example, in baseband or as part of a carrier wave. Such a propagated signal may take any of a variety of forms, including, but not limited to, electro-magnetic, optical, or any suitable combination thereof. A computer readable signal medium may be any computer readable medium that is not a computer
10 readable storage medium and that can communicate, propagate, or transport a program for use by or in connection with an instruction execution system, apparatus, or device.

Program code embodied on a computer readable medium may be transmitted using any appropriate medium, including but not limited to wireless, wireline, optical fiber cable, RF,
15 etc., or any suitable combination of the foregoing.

Computer program code for carrying out operations for aspects of the present invention may be written in any combination of one or more programming languages, including an object oriented programming language such as Java, Smalltalk, C++ or the like and conventional
20 procedural programming languages, such as the "C" programming language or similar programming languages. The program code may execute entirely on the user's computer, partly on the user's computer, as a stand-alone software package, partly on the user's computer and partly on a remote computer or entirely on the remote computer or server. In the latter scenario, the remote computer may be connected to the user's computer through
25 any type of network, including a local area network (LAN) or a wide area network (WAN), or the connection may be made to an external computer (for example, through the Internet using an Internet Service Provider).

Aspects of the present invention are described above with reference to flowchart illustrations
30 and/or block diagrams of methods, apparatus (systems) and computer program products according to embodiments of the invention. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart

illustrations and/or block diagrams, can be implemented by computer program instructions. These computer program instructions may be provided to a processor of a general purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor of the
5 computer or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks.

These computer program instructions may also be stored in a computer readable medium that can direct a computer, other programmable data processing apparatus, or other devices
10 to function in a particular manner, such that the instructions stored in the computer readable medium produce an article of manufacture including instructions which implement the function/act specified in the flowchart and/or block diagram block or blocks.

The computer program instructions may also be loaded onto a computer, other
15 programmable data processing apparatus, or other devices to cause a series of operational steps to be performed on the computer, other programmable apparatus or other devices to produce a computer implemented process such that the instructions which execute on the computer or other programmable apparatus provide processes for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks.

The flowchart and block diagrams in the Figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods and computer program products according to various embodiments of the present invention. In this regard, each block in the flowchart or block diagrams may represent a module, segment, or portion of code, which
25 comprises one or more executable instructions for implementing the specified logical function(s). It should also be noted that, in some alternative implementations, the functions noted in the block may occur out of the order noted in the figures. For example, two blocks shown in succession may, in fact, be executed substantially concurrently, or the blocks may sometimes be executed in the reverse order, depending upon the functionality involved. It
30 will also be noted that each block of the block diagrams and/or flowchart illustration, and combinations of blocks in the block diagrams and/or flowchart illustration, can be

implemented by special purpose hardware-based systems that perform the specified functions or acts, or combinations of special purpose hardware and computer instructions.

5 The descriptions in this specification are for purposes of illustration only and are not to be construed in a limiting sense. The scope of the present invention is limited only by the language of the following claims.

CLAIMS

1. A method of compiling software for a hierarchical distributed processing system, the method comprising:

5 providing to one or more compiling nodes software to be compiled, wherein at least a portion of the software to be compiled is to be executed by one or more other nodes;

compiling, by the compiling node, the software;

maintaining, by the compiling node, any compiled software to be executed on the compiling node; and

10 selecting, by the compiling node, one or more nodes in a next tier of the hierarchy of the distributed processing system in dependence upon whether any compiled software is for the selected node or the selected node's descendants; and

sending to the selected node only the compiled software to be executed by the selected node or selected node's descendants.

15 2. The method of claim 1 further comprising:

receiving, by a selected node, compiled software;

determining whether the compiled software is for the selected node or one of its descendants;

20 if the compiled software is for the selected node, maintaining the software by the selected node for execution; and

if the compiled software is for one of the descendants, selecting another node in a next tier of the hierarchical distributed processing system in dependence upon a descendant for the compiled software and sending the compiled software to the selected another node.

25 3. The method of claim 1 or 2 further comprising identifying the one or more compiling nodes.

30 4. The method of claim 3 wherein identifying the one or more compiling nodes further comprising selecting one or more nodes that are computationally optimized for compiling.

5. The method of claim 3 or 4 wherein identifying the one or more compiling nodes further comprising selecting one or more nodes that are optimized for compiling because of their location in the topology of the hierarchical distributed processing system.

5 6. The method of any preceding claim wherein the hierarchical distributed processing system further comprises a parallel computer that includes:

a plurality of compute nodes;

a first data communications network coupling the compute nodes for data communications and optimized for point to point data communications; and

10 a second data communications network that includes data communications links coupling the compute nodes so as to organize the compute nodes as a tree, each compute node having a separate arithmetic logic unit ('ALU') dedicated to parallel operations.

7. The method of any preceding claim wherein the hierarchical distributed processing system further comprises a hybrid computing environment, the hybrid computing environment comprising a plurality of compute nodes, each compute node comprising:

a host computer having a host computer architecture; and

an accelerator having an accelerator architecture, the accelerator architecture optimized, with respect to the host computer architecture, for speed of execution of a

20 particular class of computing functions, the host computer and the accelerator adapted to one another for data communications by a system level message passing module.

8. An apparatus for compiling software for a hierarchical distributed processing system, the apparatus comprising a computer processor and a computer memory operatively coupled to the computer processor, the computer memory having disposed within it computer program instructions for

providing to one or more compiling nodes software to be compiled, wherein at least a portion of the software to be compiled is to be executed by one or more other nodes;

compiling, by the compiling node, the software;

30 maintaining, by the compiling node, any compiled software to be executed on the compiling node; and

selecting, by the compiling node, one or more nodes in a next tier of the hierarchy of the distributed processing system in dependence upon whether any compiled software is for the selected node or the selected node's descendants; and

5 sending to the selected node only the compiled software to be executed by the selected node or selected node's descendants.

9. The apparatus of claim 8 wherein the computer memory also has disposed within it computer program instructions for

receiving, by a selected node, compiled software;

10 determining whether the compiled software is for the selected node or one of its descendants;

if the compiled software is for the selected node, maintaining the software by the selected node for execution; and

15 if the compiled software is for one of the descendants, selecting another node in a next tier of the hierarchical distributed processing system in dependence upon a descendant for the compiled software and sending the compiled software to the selected another node.

10. The apparatus of claim 8 or 9 wherein the computer memory also has disposed within it computer program instructions for identifying the one or more compiling nodes.

20 11. The apparatus of claim 10 wherein computer program instructions for identifying the one or more compiling nodes further comprise computer program instructions for selecting one or more nodes that are computationally optimized for compiling.

25 12. The apparatus of claim 10 or 11 wherein computer program instructions for identifying the one or more compiling nodes further comprise computer program instructions for selecting one or more nodes that are optimized for compiling because of their location in the topology of the hierarchical distributed processing system.

30 13. The apparatus of any of claims 8 to 12 wherein the hierarchical distributed processing system further comprises a parallel computer that includes:

a plurality of compute nodes;

a first data communications network coupling the compute nodes for data communications and optimized for point to point data communications; and

a second data communications network that includes data communications links coupling the compute nodes so as to organize the compute nodes as a tree, each compute node having a separate arithmetic logic unit ('ALU') dedicated to parallel operations.

14. The apparatus of any of claims 8 to 13 wherein the hierarchical distributed processing system further comprises a hybrid computing environment, the hybrid computing environment comprising a plurality of compute nodes, each compute node comprising:

a host computer having a host computer architecture; and

an accelerator having an accelerator architecture, the accelerator architecture optimized, with respect to the host computer architecture, for speed of execution of a particular class of computing functions, the host computer and the accelerator adapted to one another for data communications by a system level message passing module.

15. A computer program product for compiling software for a hierarchical distributed processing system, the computer program product disposed in a computer readable storage medium, the computer program product comprising computer program instructions for

providing to one or more compiling nodes software to be compiled, wherein at least a portion of the software to be compiled is to be executed by one or more other nodes;

compiling, by the compiling node, the software;

maintaining, by the compiling node, any compiled software to be executed on the compiling node; and

selecting, by the compiling node, one or more nodes in a next tier of the hierarchy of the distributed processing system in dependence upon whether any compiled software is for the selected node or the selected node's descendants; and

sending to the selected node only the compiled software to be executed by the selected node or selected node's descendants.

16. The computer program produce of claim 15 further comprising computer program instructions for

receiving, by a selected node, compiled software;

determining whether the compiled software is for the selected node or one of its descendants;

if the compiled software is for the selected node, maintaining the software by the selected node for execution; and

5 if the compiled software is for one of the descendants, selecting another node in a next tier of the hierarchical distributed processing system in dependence upon a descendant for the compiled software and sending the compiled software to the selected another node.

10 17. The computer program produce of claim 15 or 16 further comprising computer program instructions for identifying the one or more compiling nodes.

15 18. The computer program produce of claim 17 wherein computer program instructions for identifying the one or more compiling nodes further comprise computer program instructions for selecting one or more nodes that are computationally optimized for compiling.

20 19. The computer program produce of claim 17 or 18 wherein computer program instructions for identifying the one or more compiling nodes further comprise computer program instructions for selecting one or more nodes that are optimized for compiling because of their location in the topology of the hierarchical distributed processing system.

25 20. The computer program produce of any of claims 15 to 19 wherein the hierarchical distributed processing system further comprises a parallel computer that includes:

a plurality of compute nodes;

25 a first data communications network coupling the compute nodes for data communications and optimized for point to point data communications; and

a second data communications network that includes data communications links coupling the compute nodes so as to organize the compute nodes as a tree, each compute node having a separate arithmetic logic unit ('ALU') dedicated to parallel operations.

30

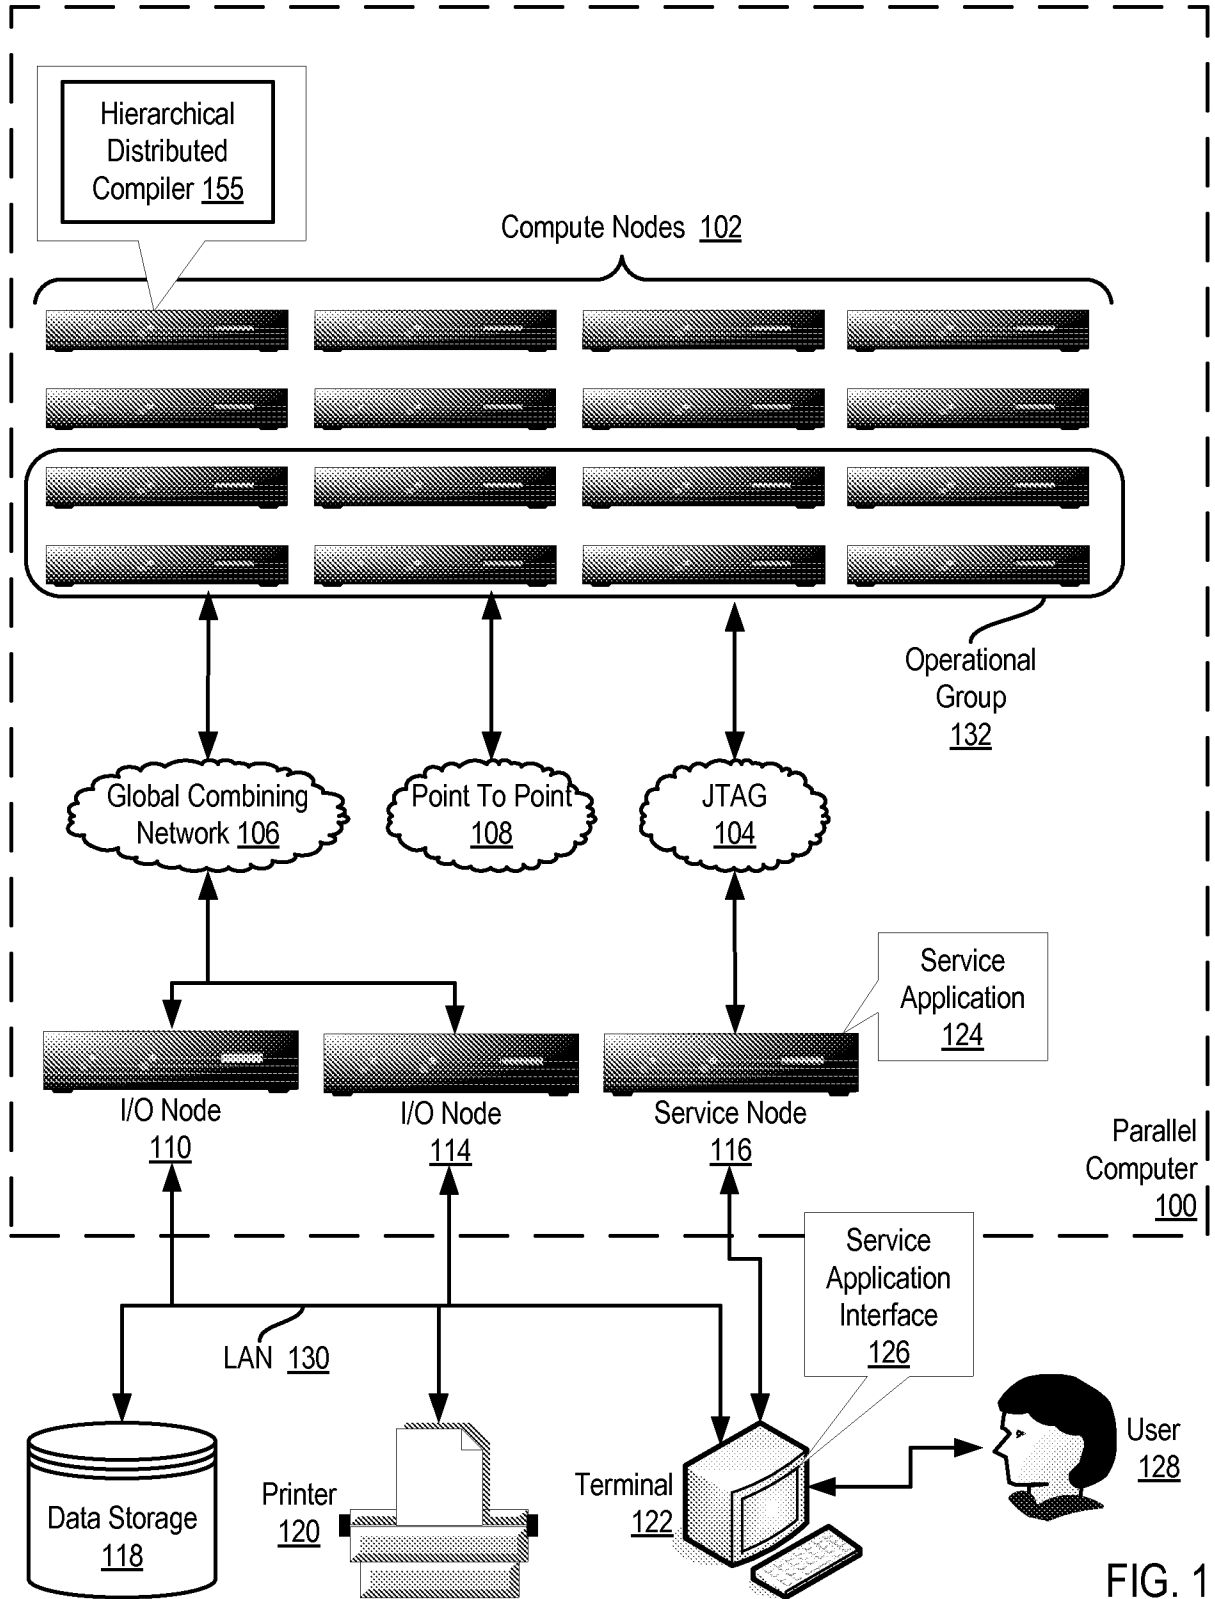


FIG. 1




FIG. 2

3/9




FIG. 3A




FIG. 3B




FIG. 4

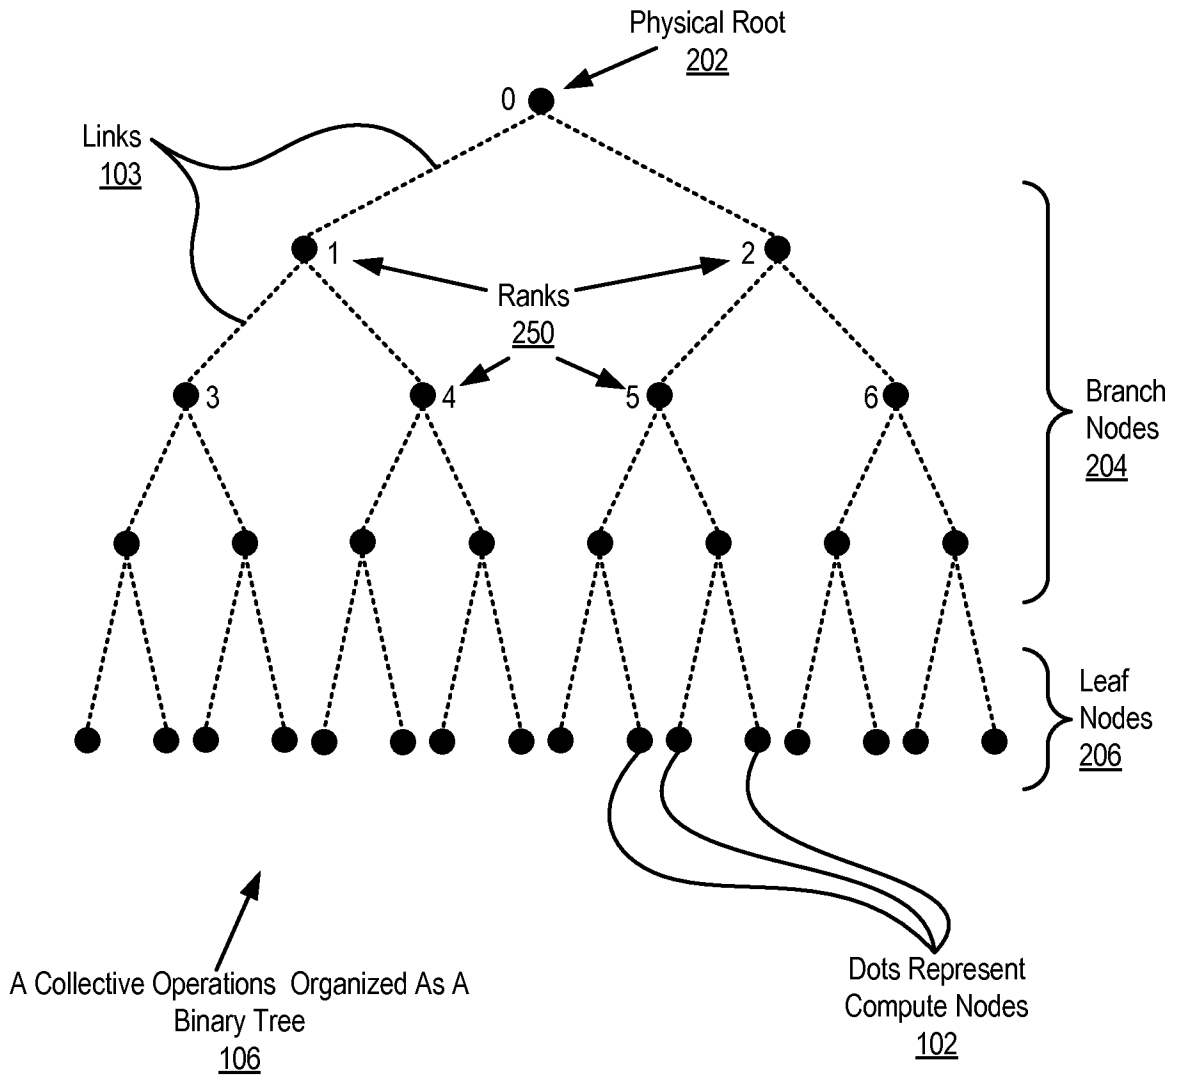


FIG. 5

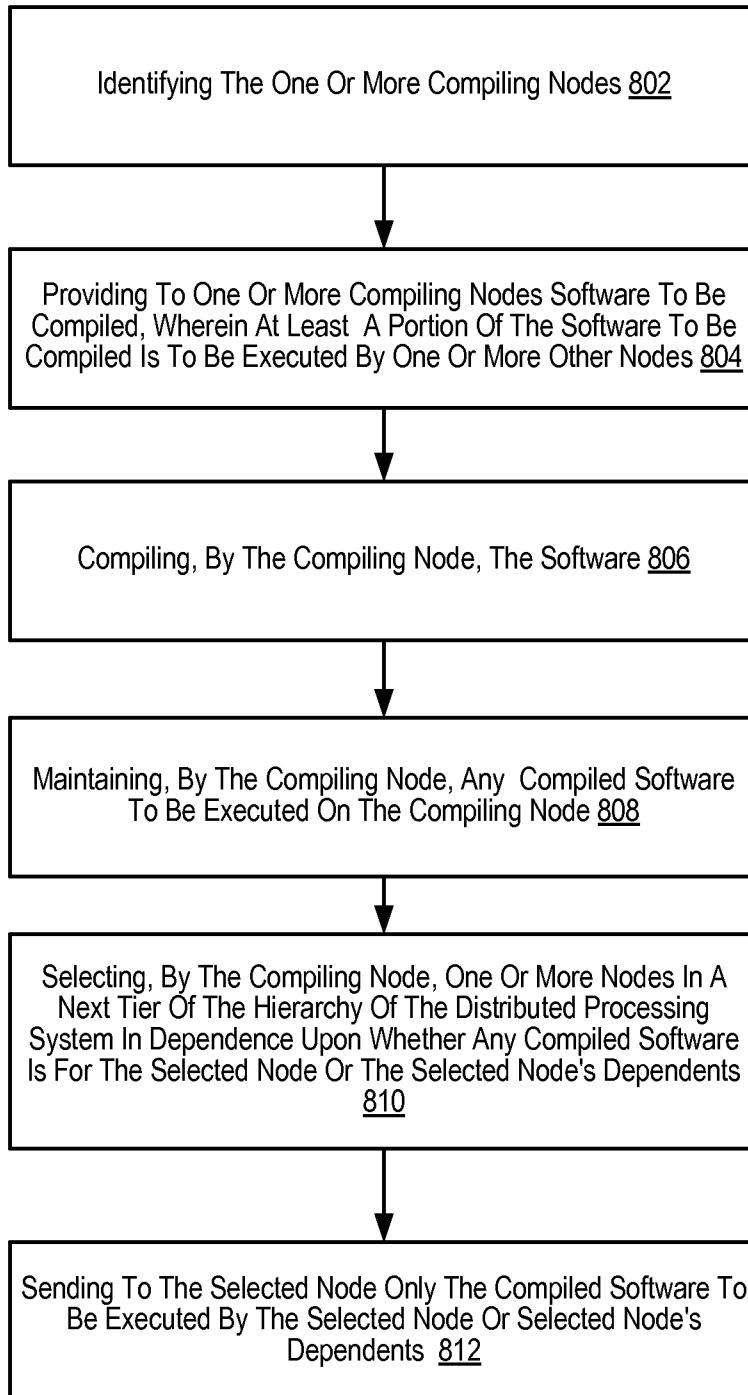


FIG. 7

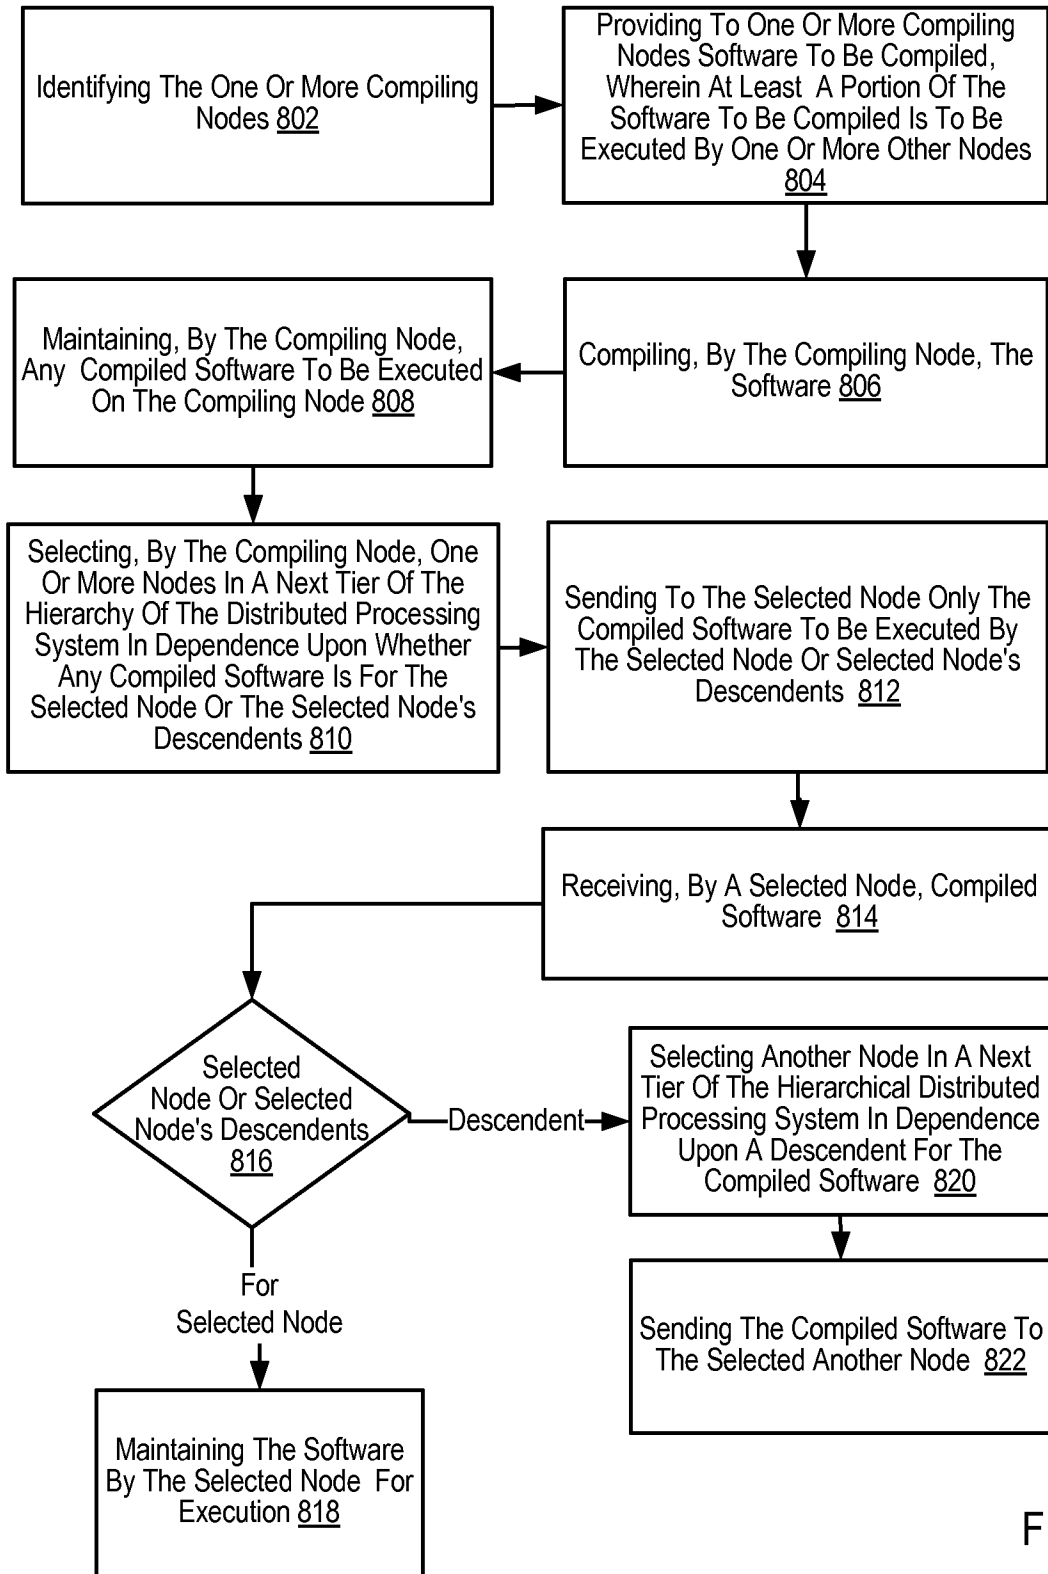


FIG. 8

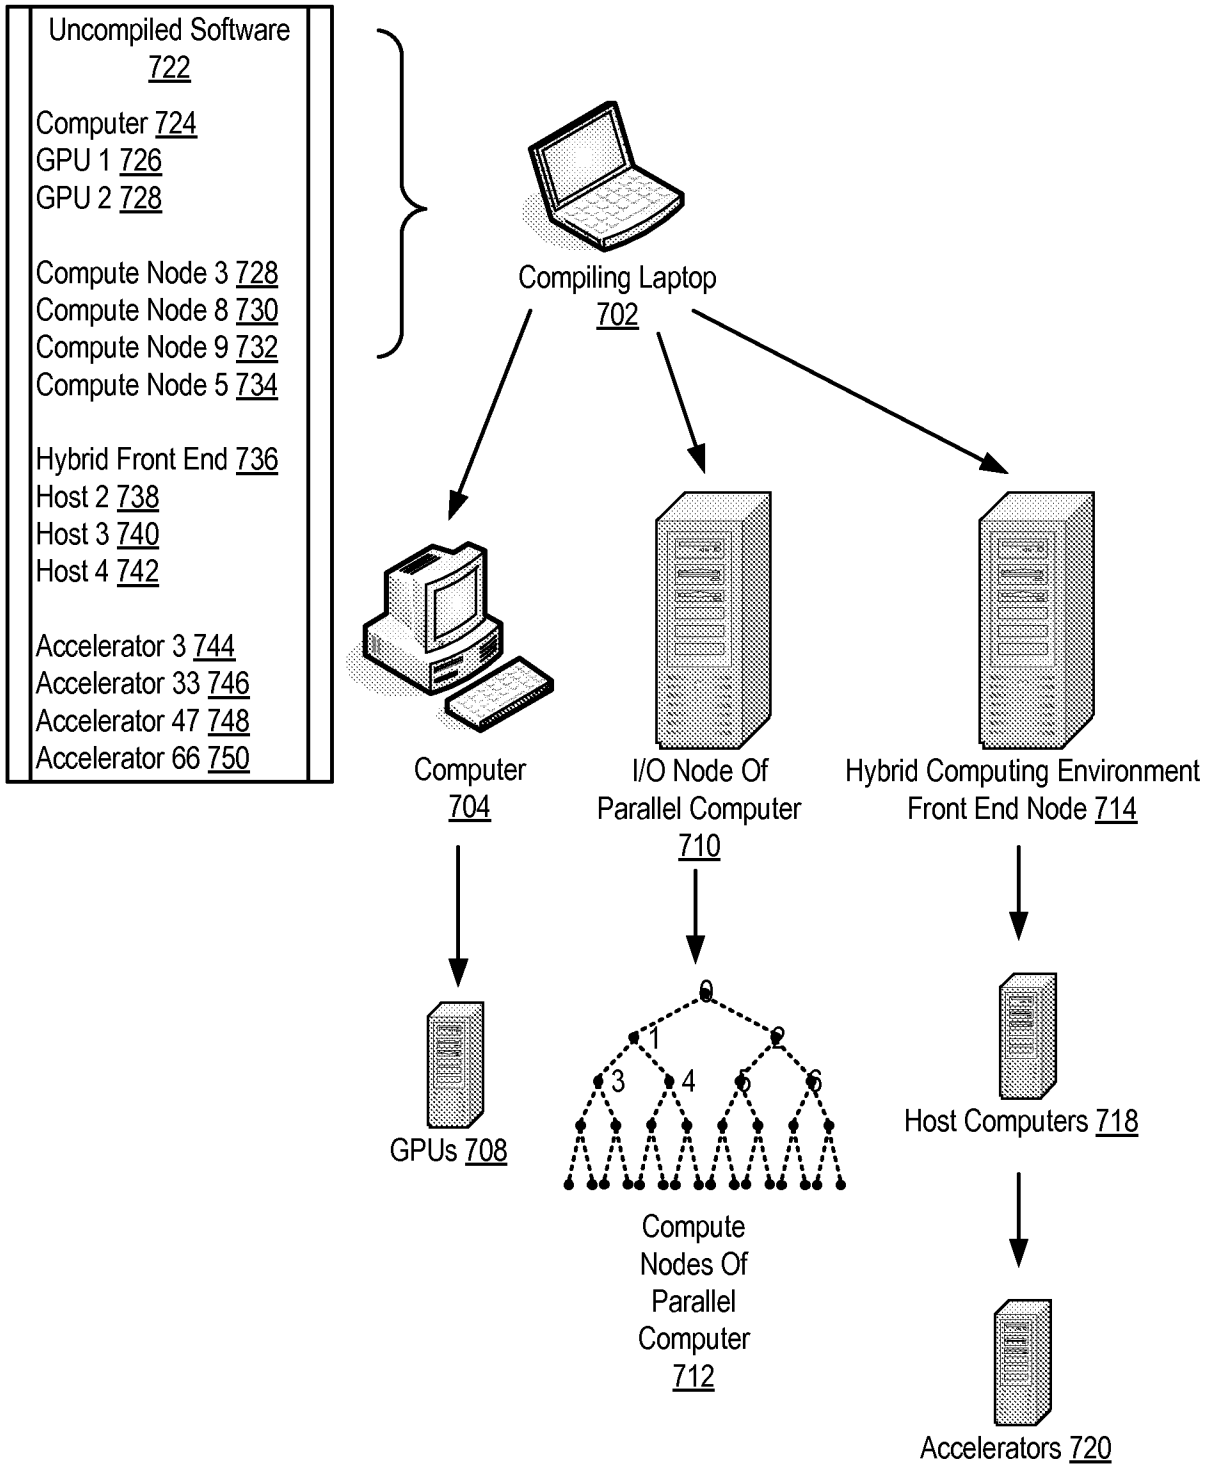


FIG. 9

INTERNATIONAL SEARCH REPORT

International application No
PCT/EP2011/052105

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06F9/45
ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED
Minimum documentation searched (classification system followed by classification symbols)
G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)
EPO-Internal

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2004/083462 A1 (GSCHWIND MICHAEL KARL [US] ET AL) 29 April 2004 (2004-04-29) paragraph [0022] paragraph [0029]	1-20
A	US 2008/256219 A1 (ZHANG WEIJIA [US] ET AL) 16 October 2008 (2008-10-16) paragraph [0032] paragraph [0035] paragraph [0039]	1-20

Further documents are listed in the continuation of Box C.

See patent family annex.

* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

- "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- "&" document member of the same patent family

Date of the actual completion of the international search

11 May 2011

Date of mailing of the international search report

19/05/2011

Name and mailing address of the ISA/
European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

Ebert, Werner

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/EP2011/052105

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2004083462	A1	29-04-2004 CN 1499369 A	26-05-2004
US 2008256219	A1	16-10-2008 NONE	