



(19) **United States**

(12) **Patent Application Publication**  
**Frank et al.**

(10) **Pub. No.: US 2017/0160984 A1**

(43) **Pub. Date: Jun. 8, 2017**

(54) **MEMORY FABRIC OPERATIONS AND COHERENCY USING FAULT TOLERANT OBJECTS**

(52) **U.S. Cl.**  
CPC ..... **G06F 3/0653** (2013.01); **G06F 3/0604** (2013.01); **G06F 3/064** (2013.01); **G06F 3/0673** (2013.01)

(71) Applicant: **Ultrata, LLC**, Vienna, VA (US)

(57) **ABSTRACT**

(72) Inventors: **Steven Frank**, Boulder, CO (US);  
**Larry Reback**, Vienna, VA (US)

According to one embodiment, a hardware-based processing node of a plurality of hardware-based processing nodes in an object memory fabric can comprise a memory module storing and managing a plurality of memory objects in a hierarchy of the object memory fabric. Each memory object can be created natively within the memory module, accessed using a single memory reference instruction without Input/Output (I/O) instructions, and managed by the memory module at a single memory layer. The object memory fabric can distribute and track the memory objects across the hierarchy of the object memory fabric and the plurality of hardware-based processing nodes on a per-object basis. Distributing the memory objects across the hierarchy of the object memory fabric and the plurality of hardware-based processing nodes can comprise storing, on a per-object basis, each memory object on two or more nodes of the plurality of hardware-based processing nodes of the object memory fabric.

(21) Appl. No.: **15/371,440**


(22) Filed: **Dec. 7, 2016**

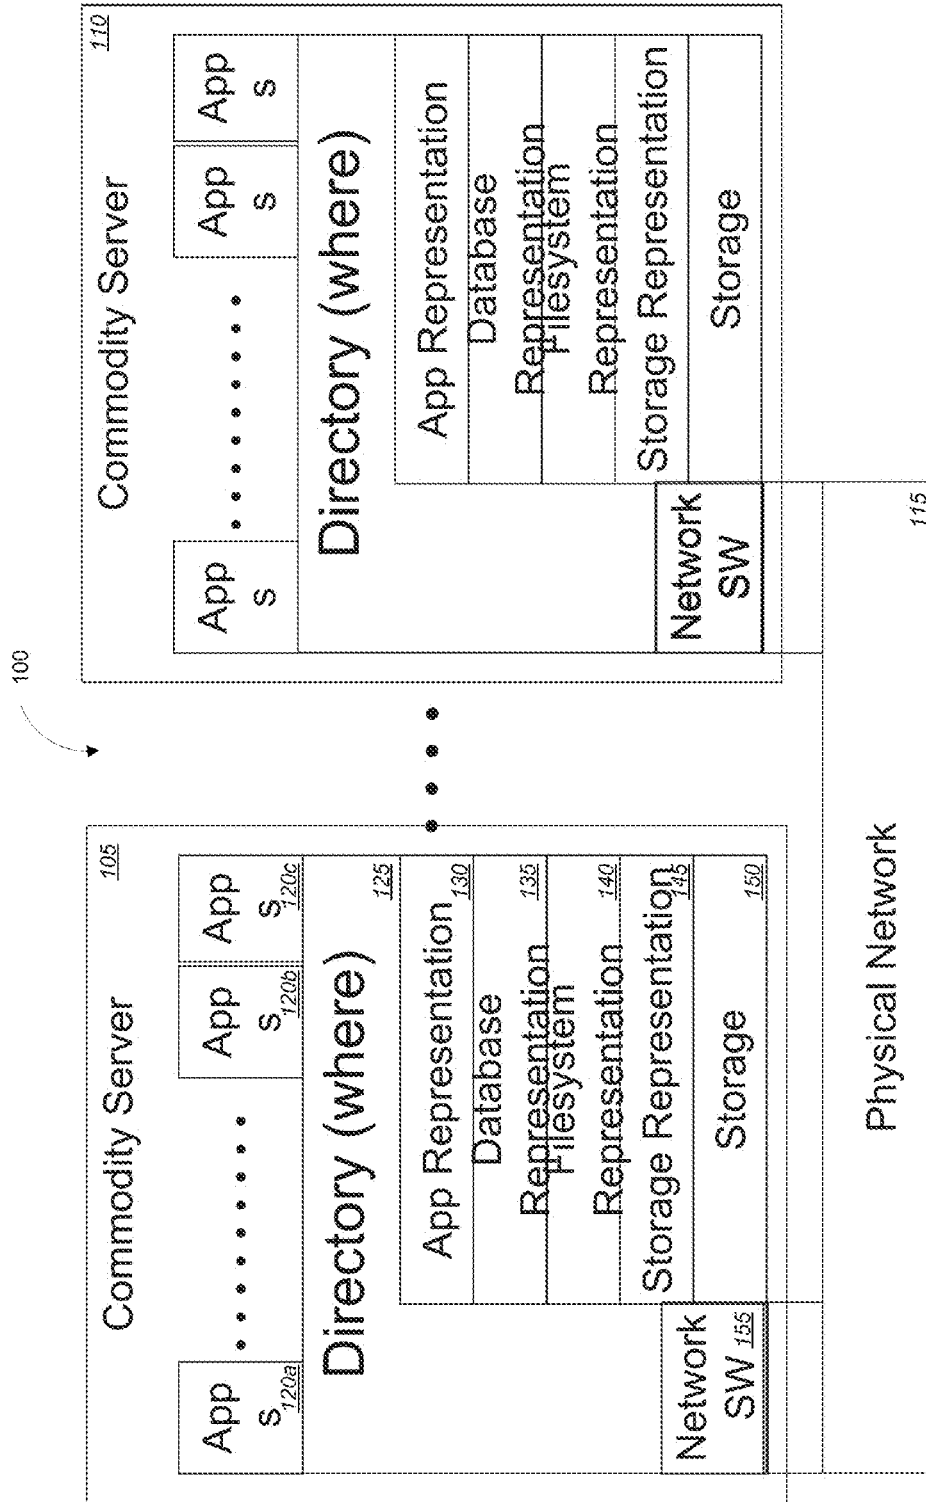
**Related U.S. Application Data**

(60) Provisional application No. 62/264,652, filed on Dec. 8, 2015.

**Publication Classification**

(51) **Int. Cl.**  
**G06F 3/06** (2006.01)





**FIG. 1**  
**(Prior Art)**




FIG. 2




FIG. 3




FIG. 4




FIG. 5




FIG. 6

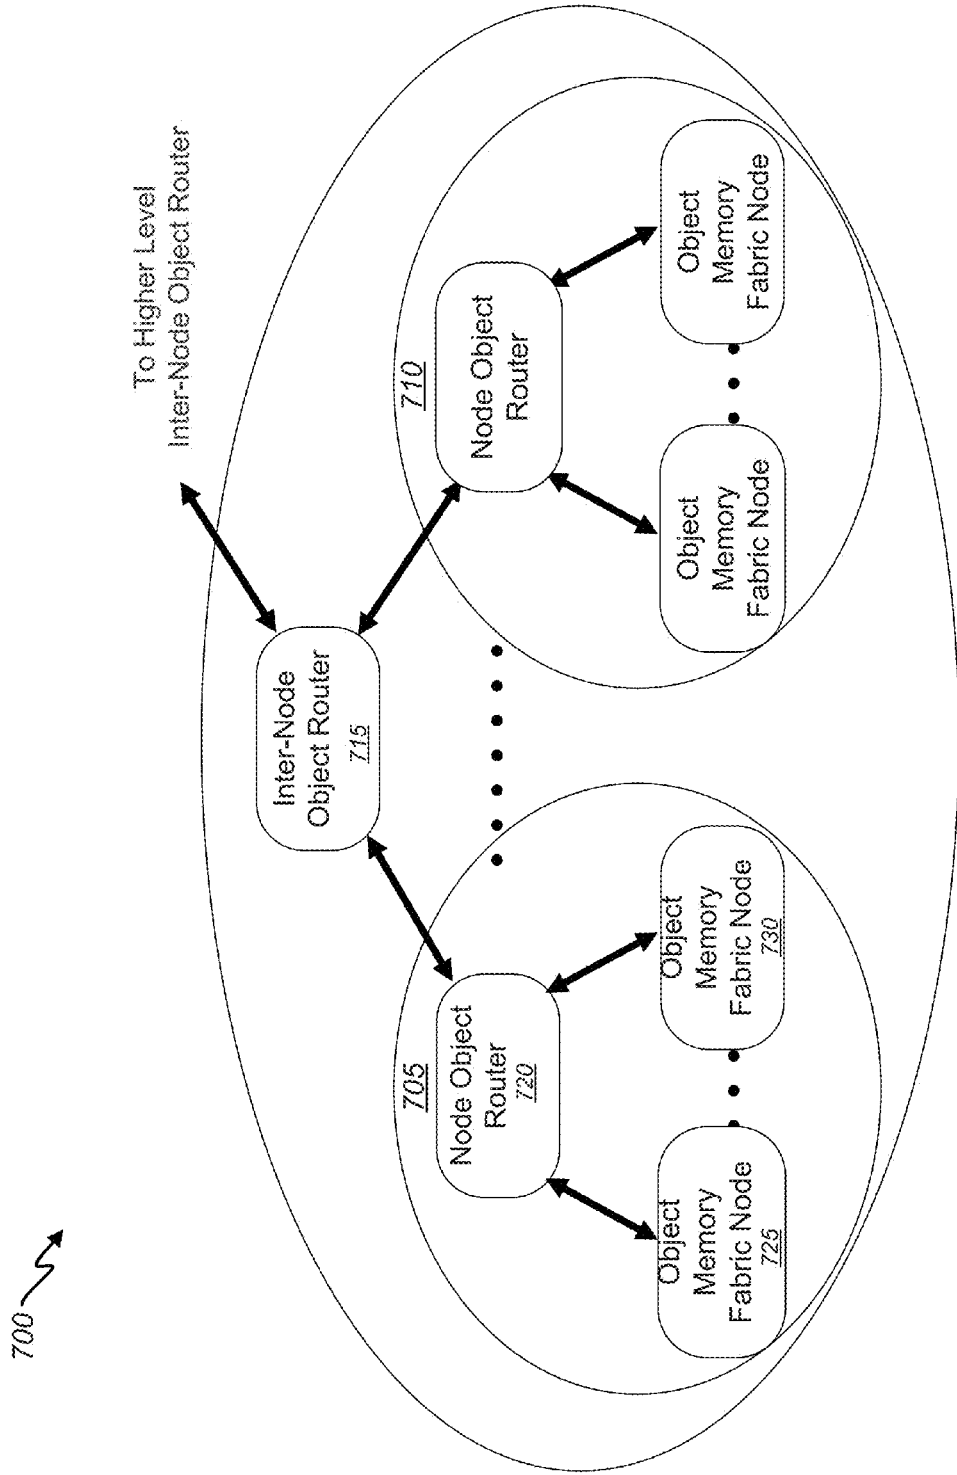


FIG. 7






FIG. 8




FIG. 9




FIG. 10




FIG. 11




FIG. 12




FIG. 13




FIG. 14




FIG. 15A



1550 ↗




FIG. 15B

1600 ↗

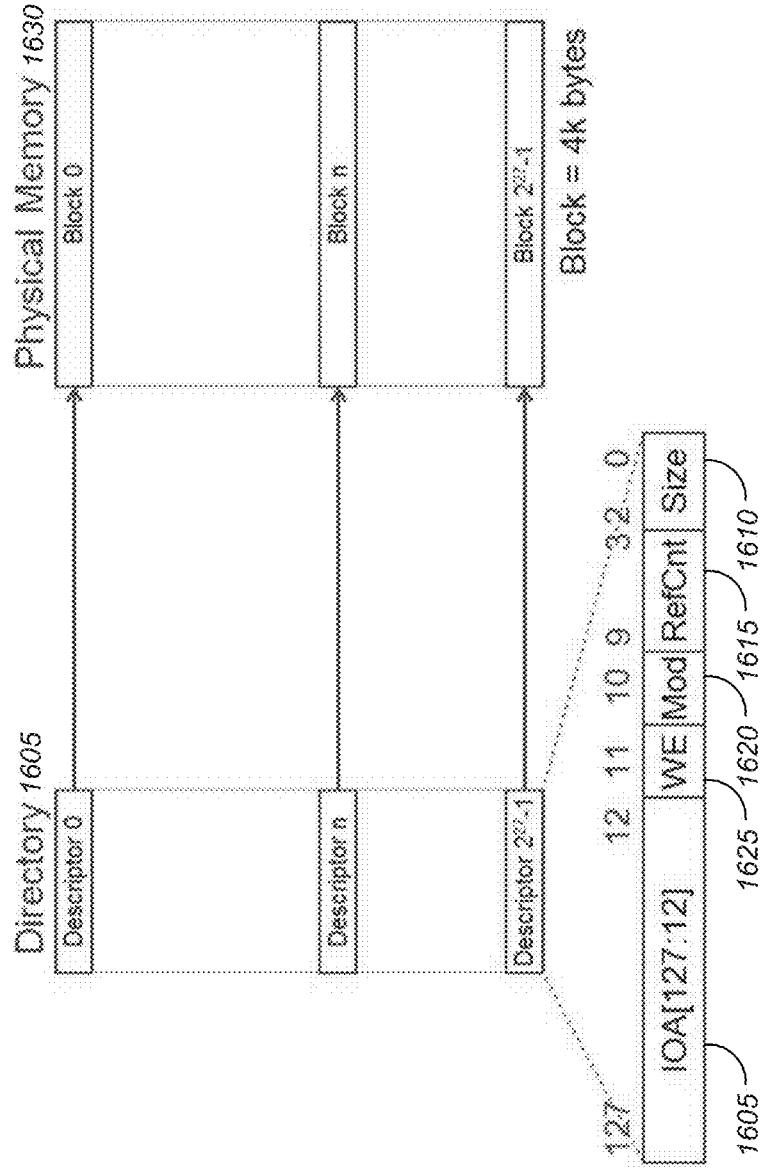


FIG. 16




FIG. 17A

1750 ↗




FIG. 17B

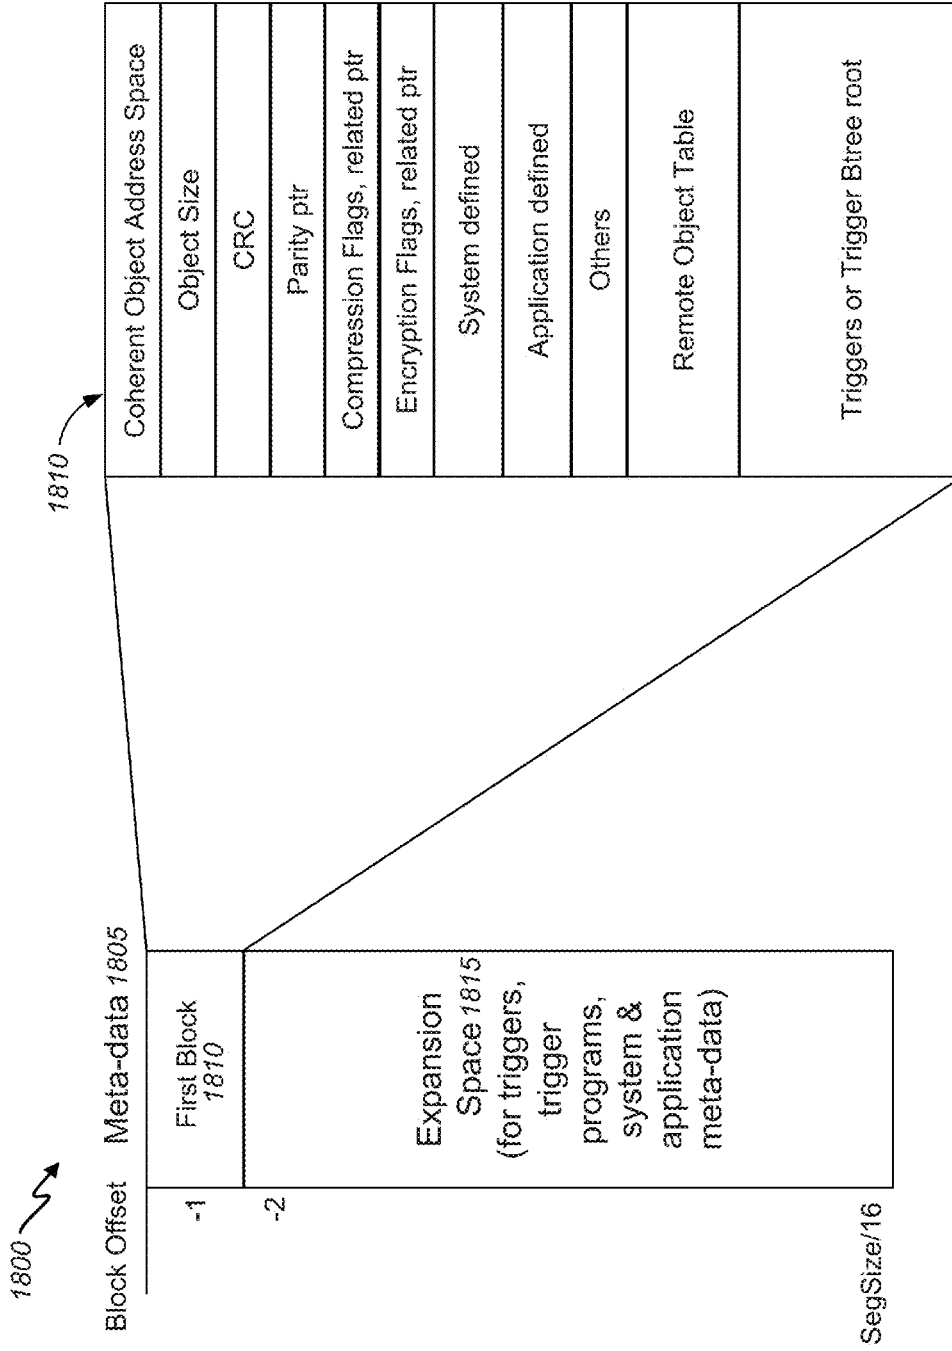


FIG. 18

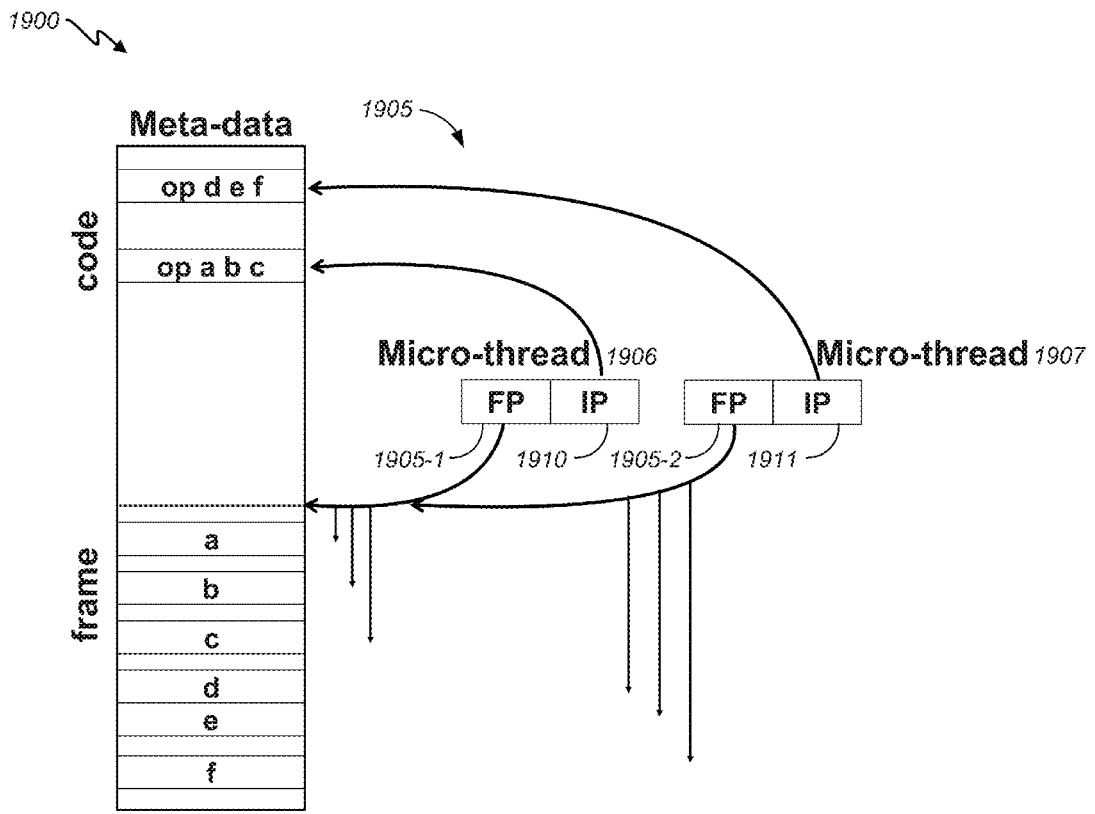


FIG. 19




FIG. 20

2100 ↗

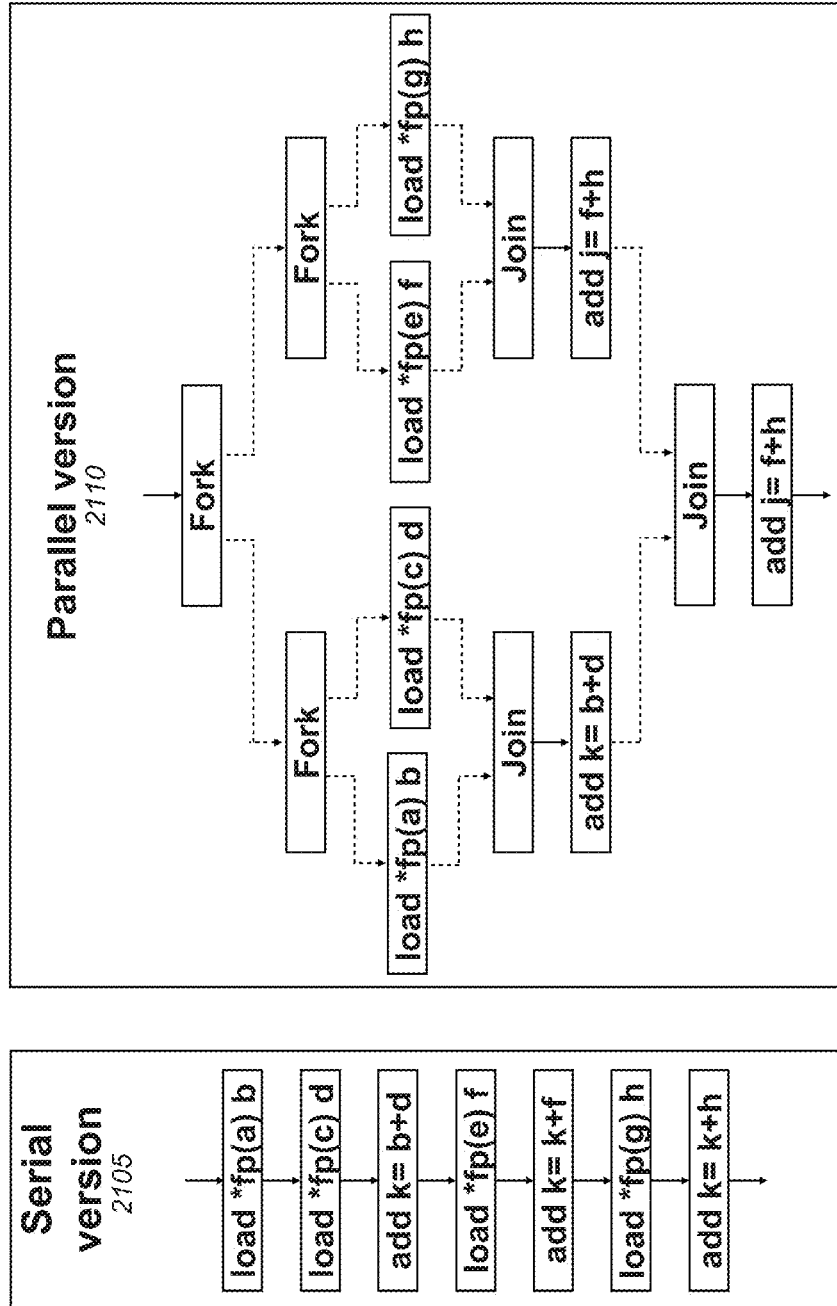


FIG. 21






FIG. 22A




FIG. 22B

STREAMS – MEMORY FABRIC ROUTER EXAMPLE




FIG. 23

# ULTRATA PRODUCT 1 FAMILY BIG DATA SOLUTION

Memory and Acceleration Engine within IMF, Application Engine on Intel Processor




FIG. 24

# ULTRATA PRODUCT 2 FAMILY BIG DATA SOLUTION

Memory Fabric, Acceleration Engine & Application Engine within IMF




FIG. 25

### MEMORY FABRIC SERVER VIEW




FIG. 26

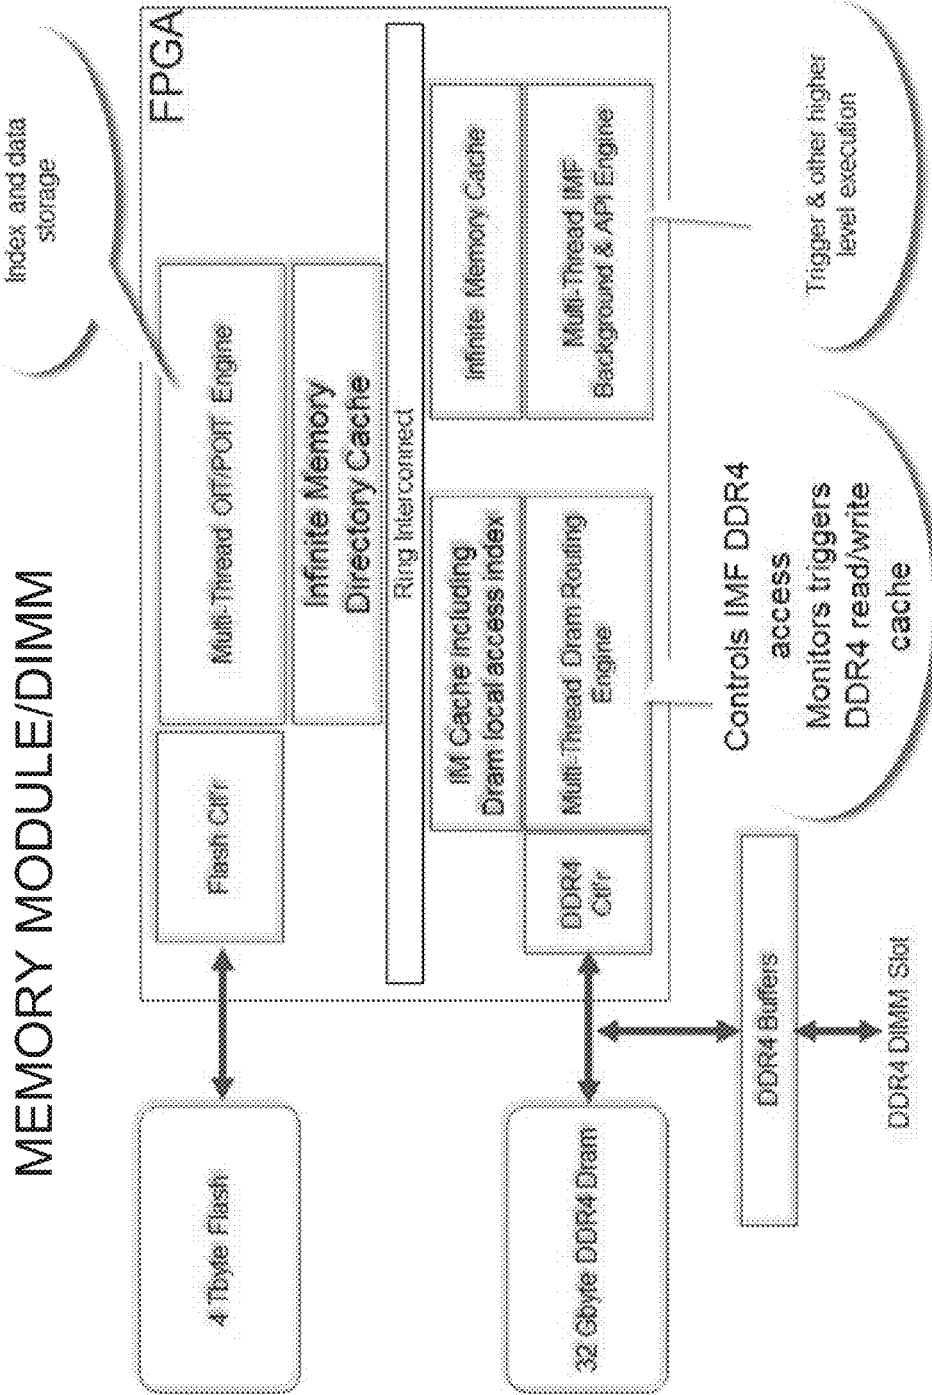


FIG. 27

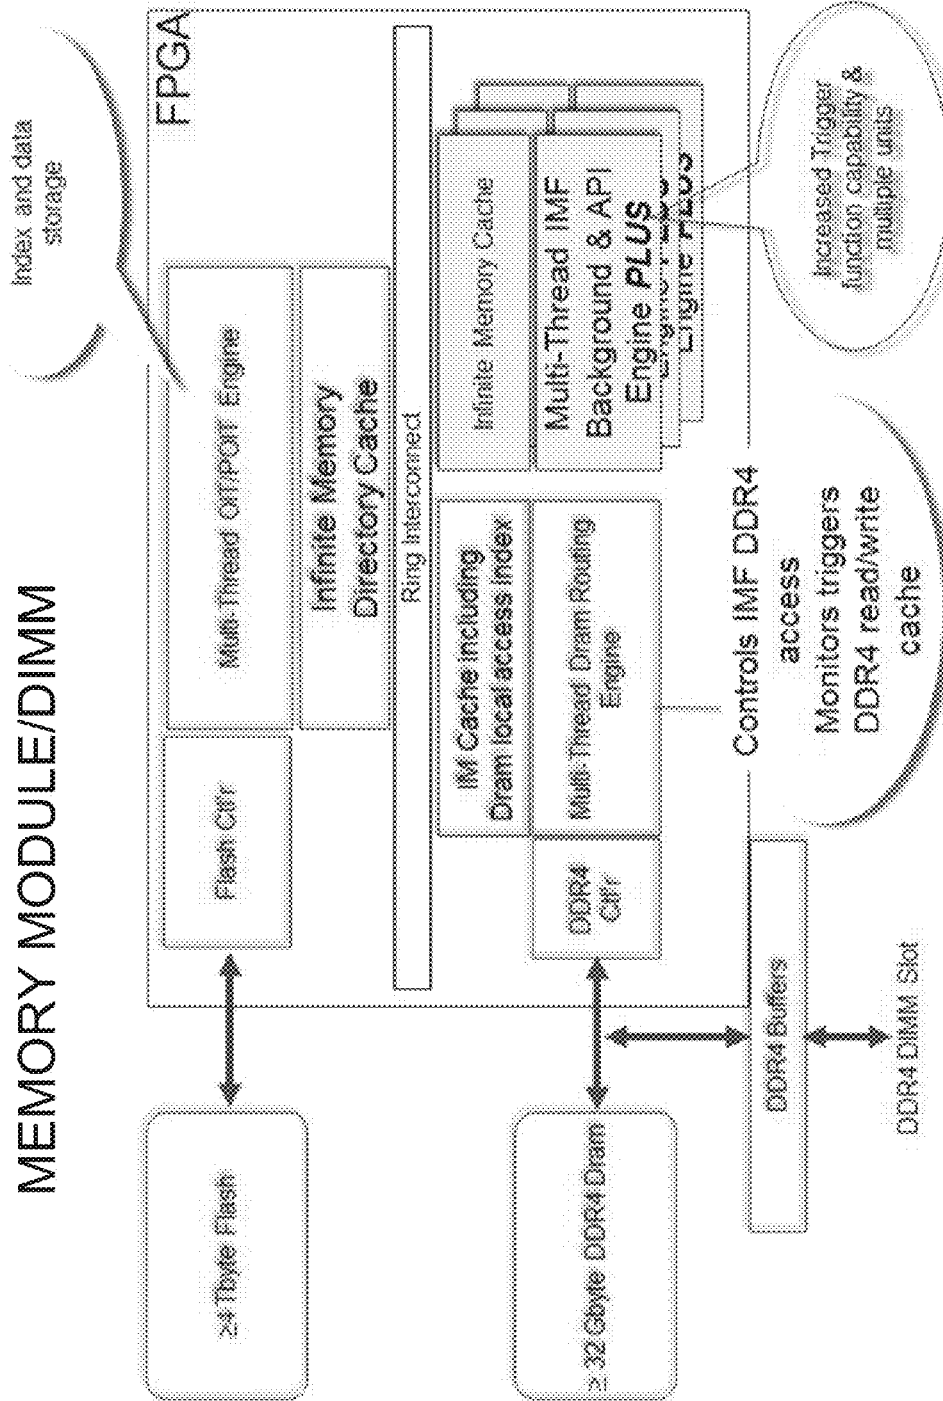
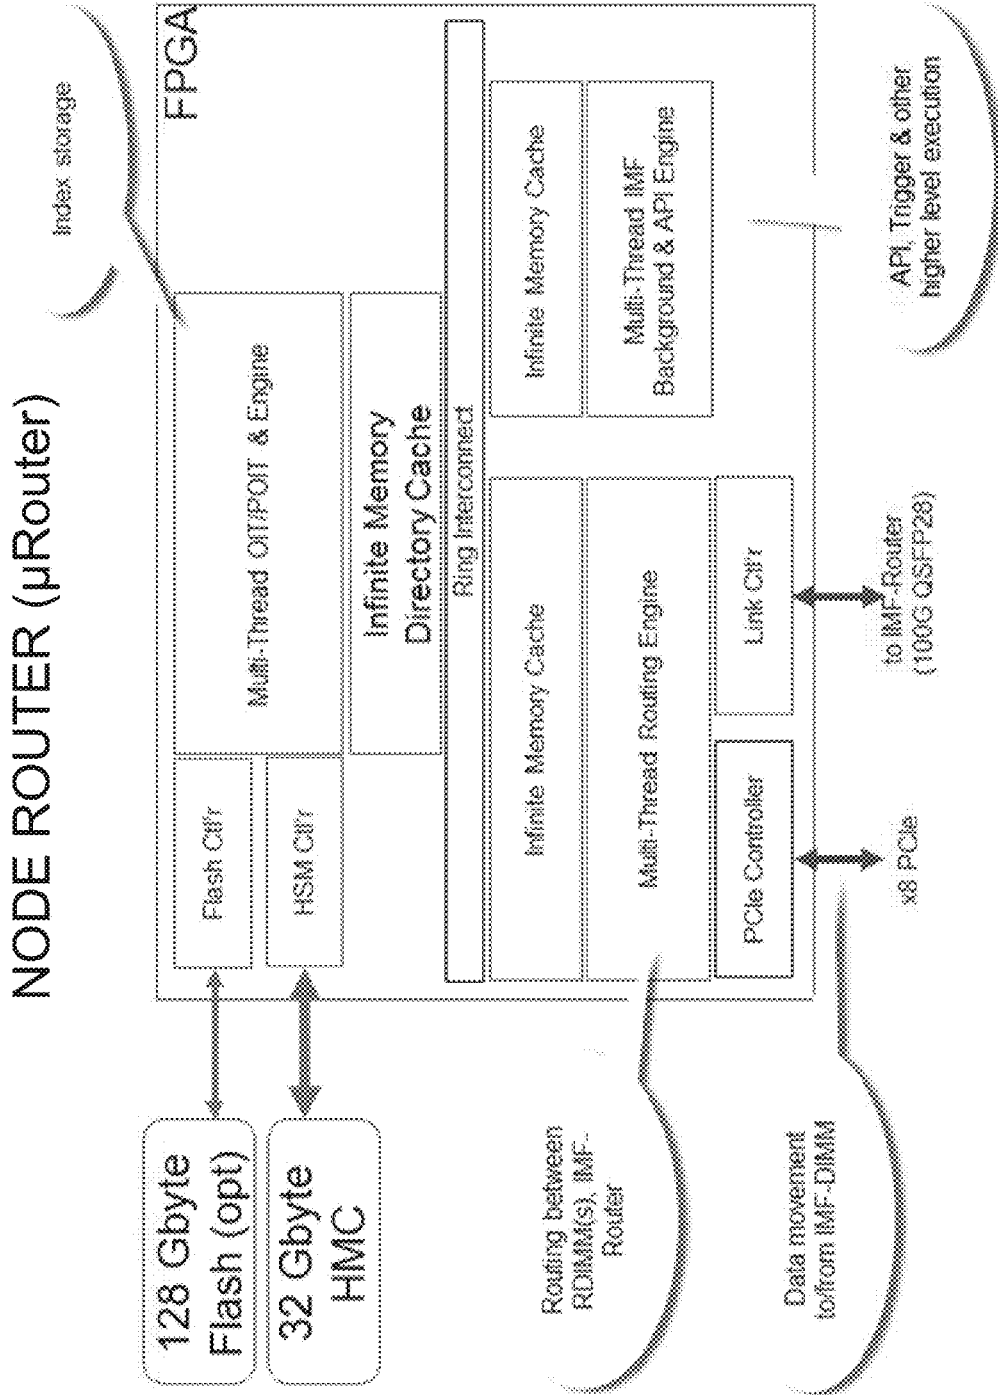


FIG. 28





**FIG. 29**

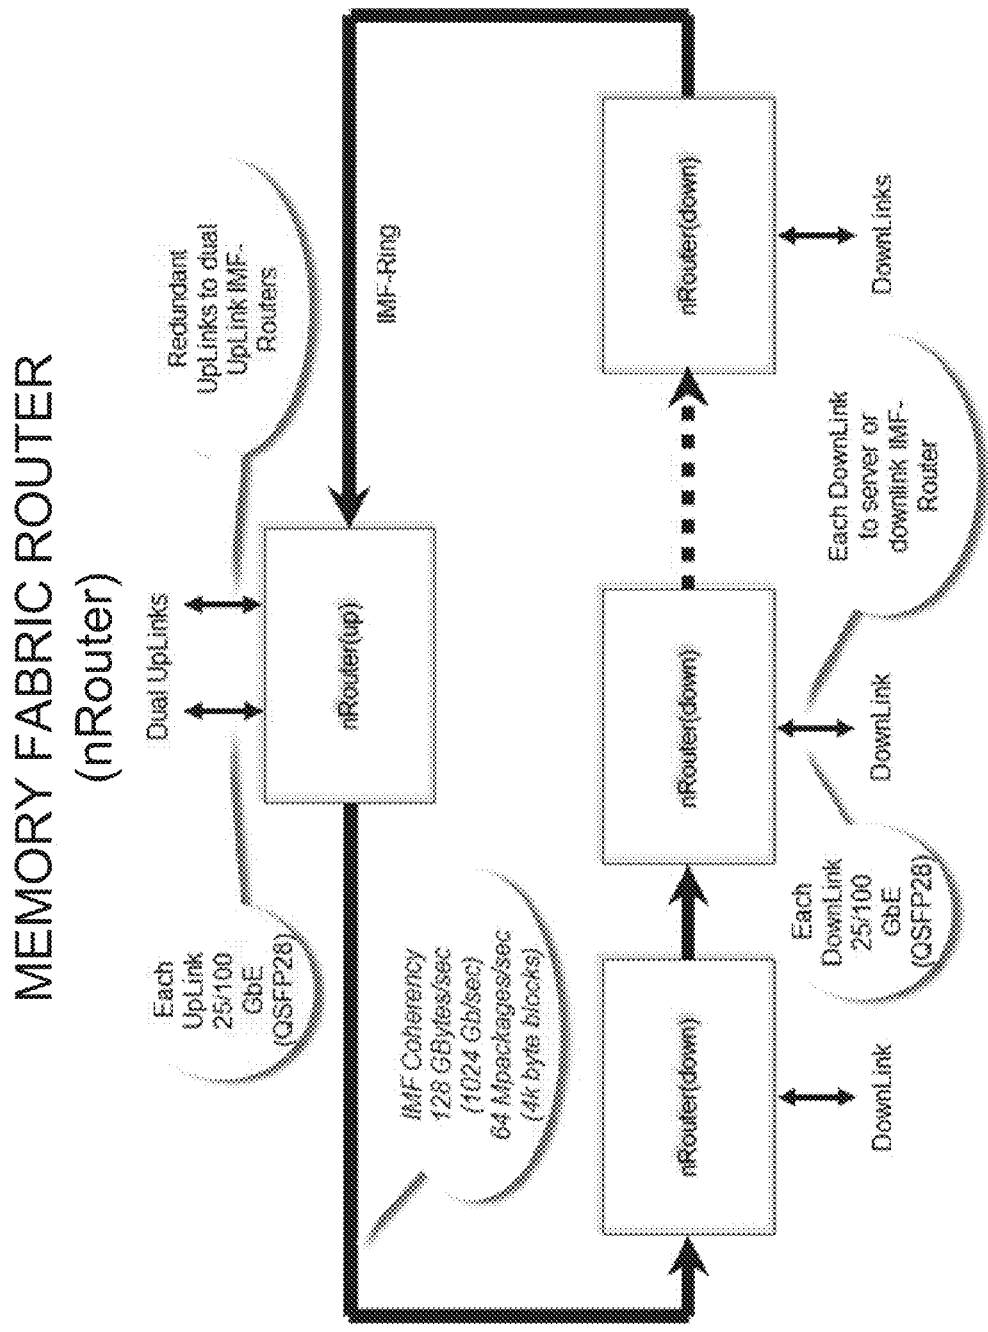


FIG. 30

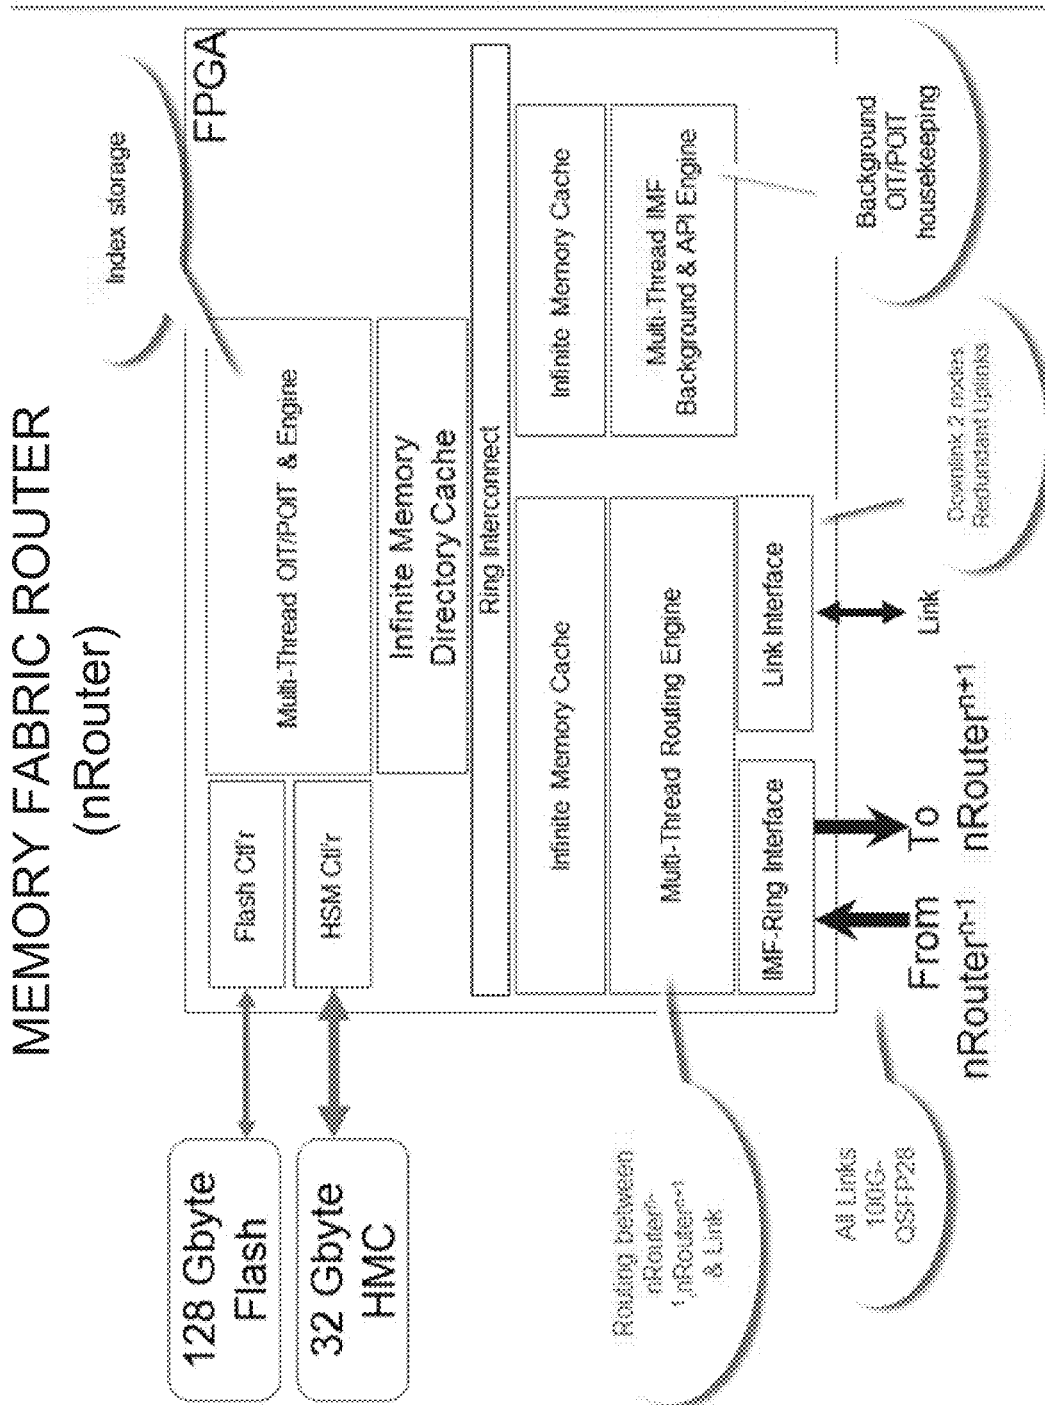


FIG. 31

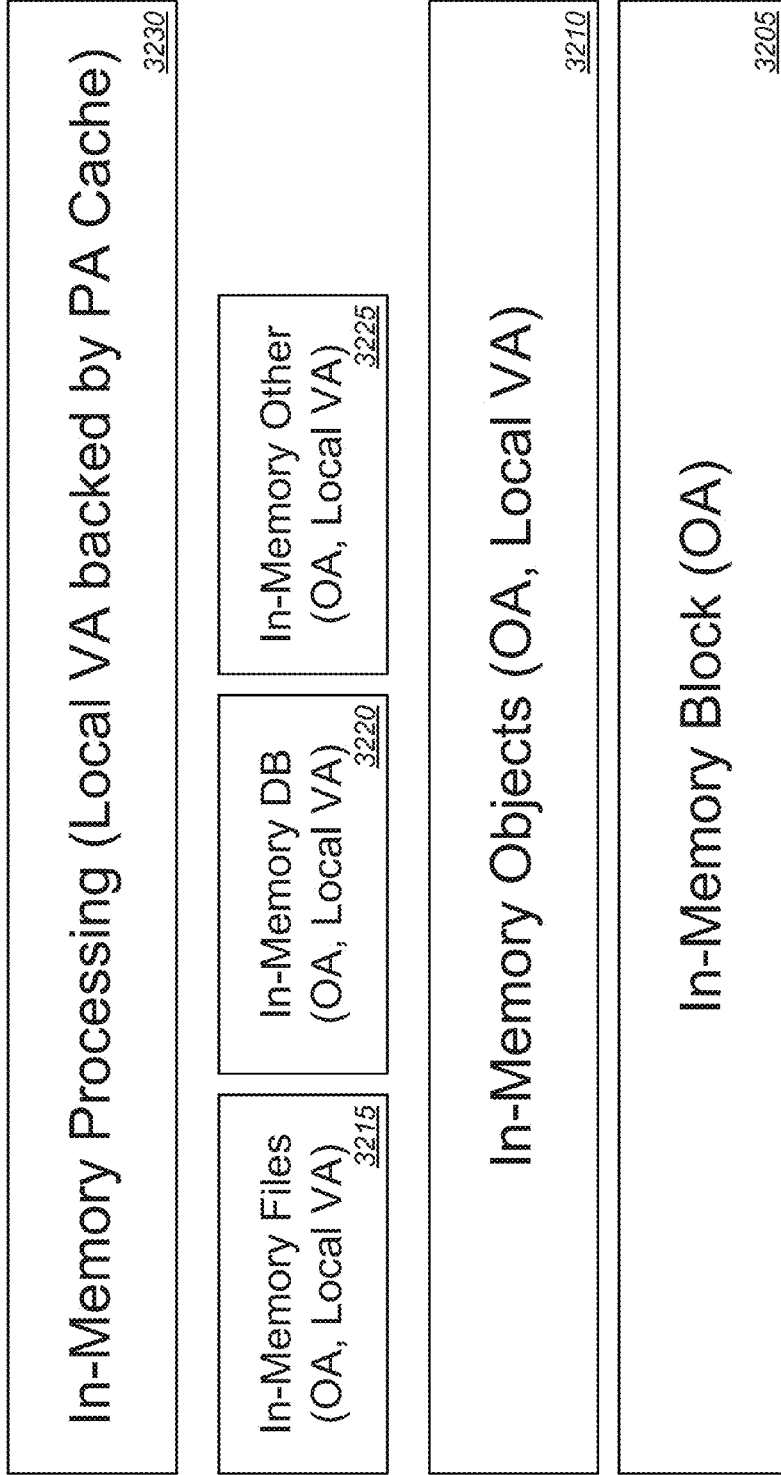


FIG. 32

3300 ↘

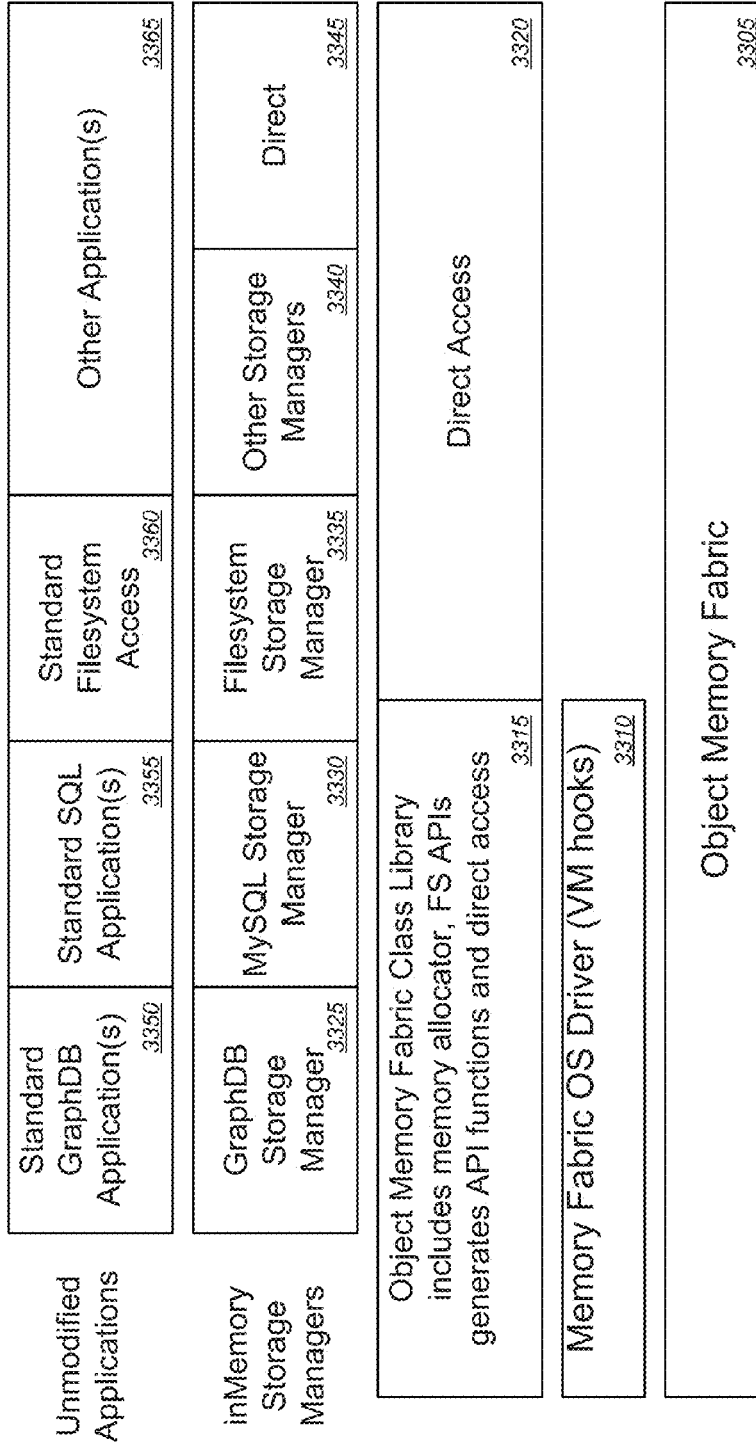


FIG. 33

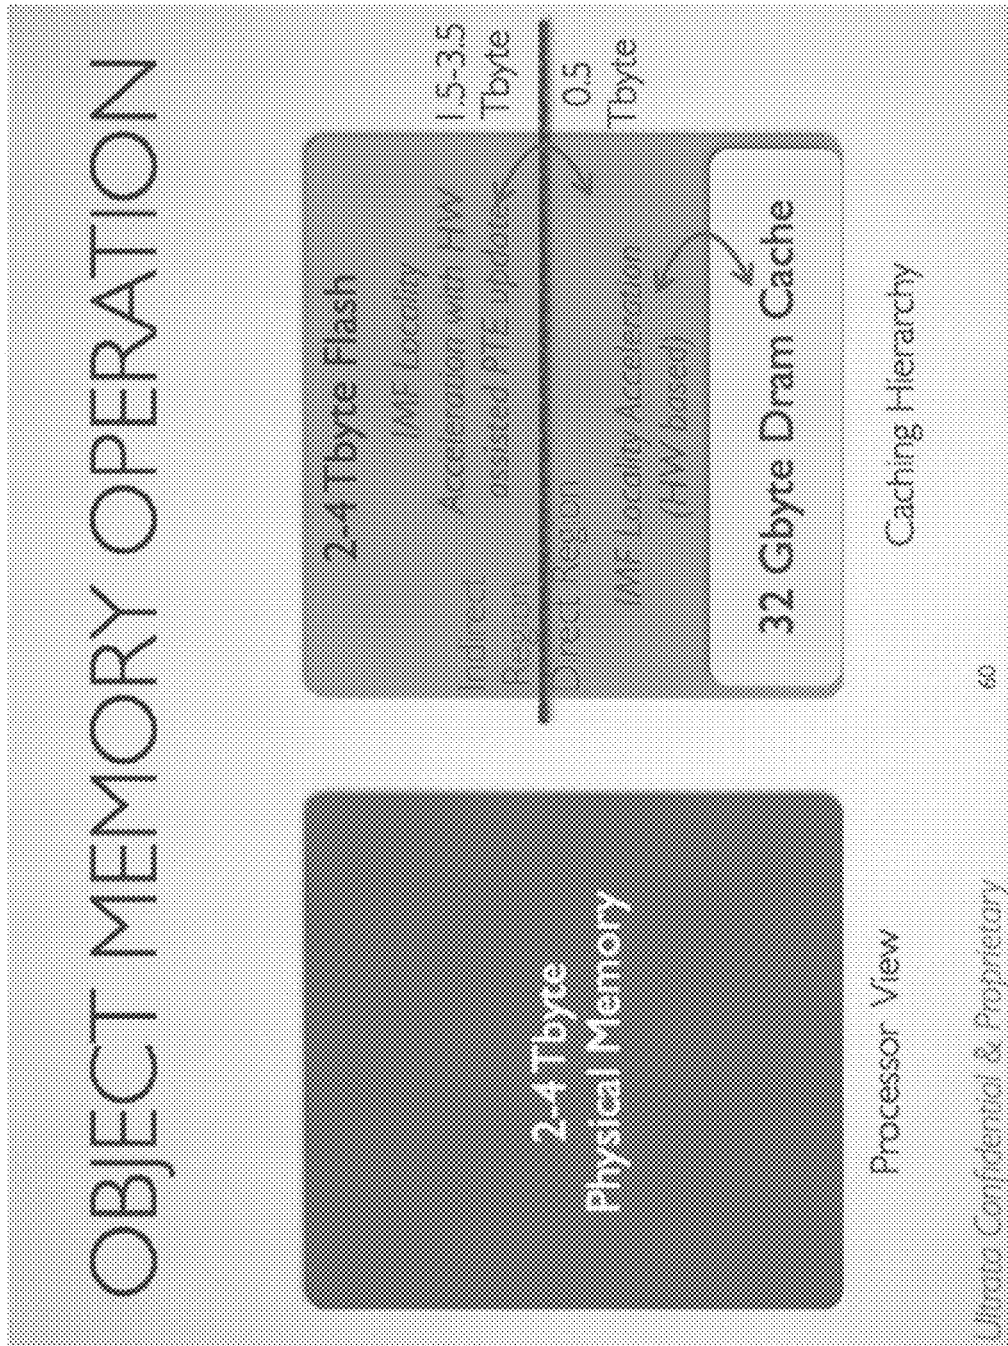
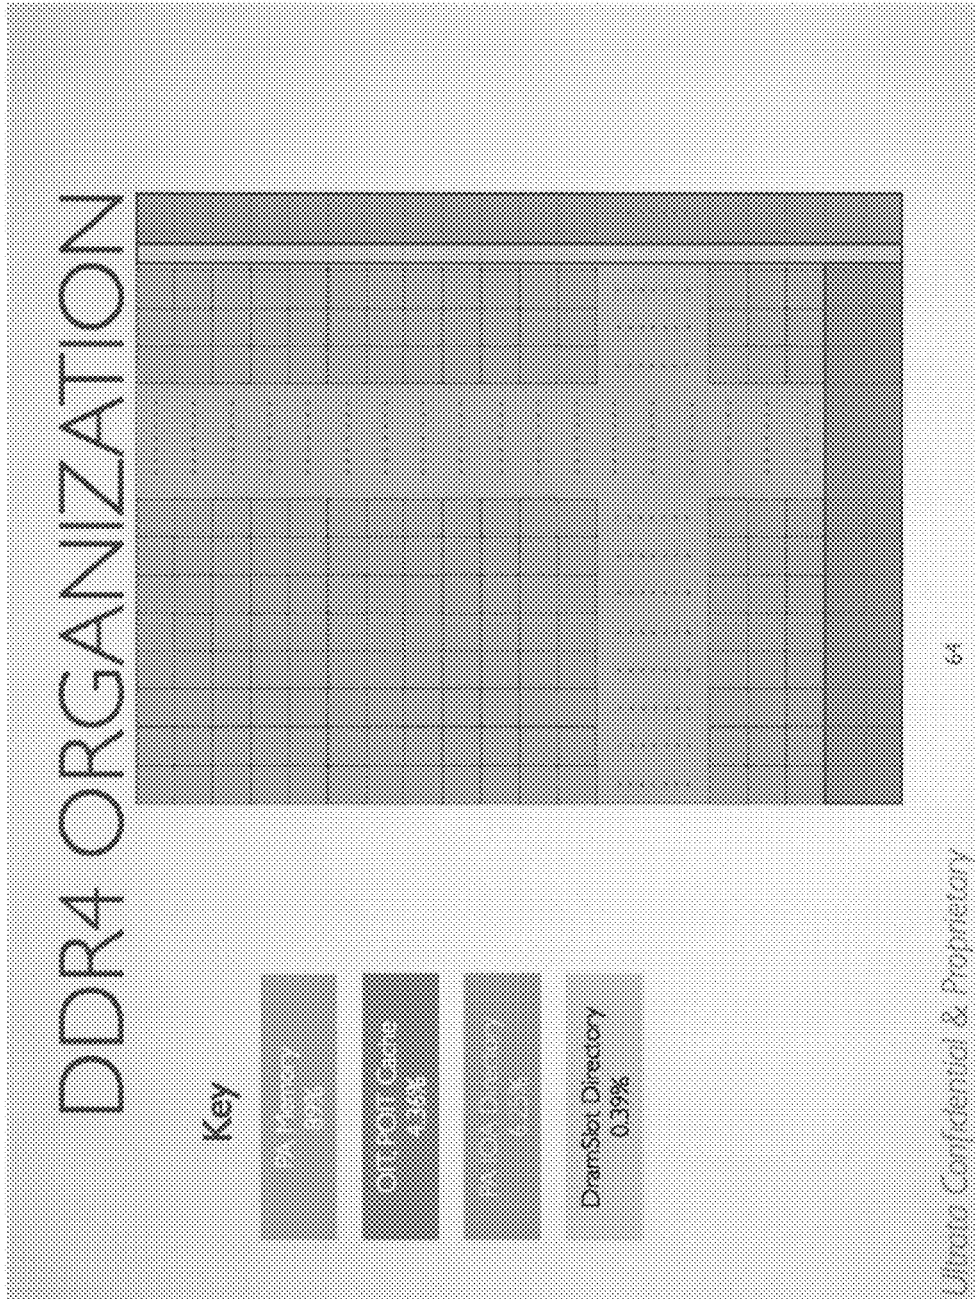


FIG. 34



**FIG. 35**

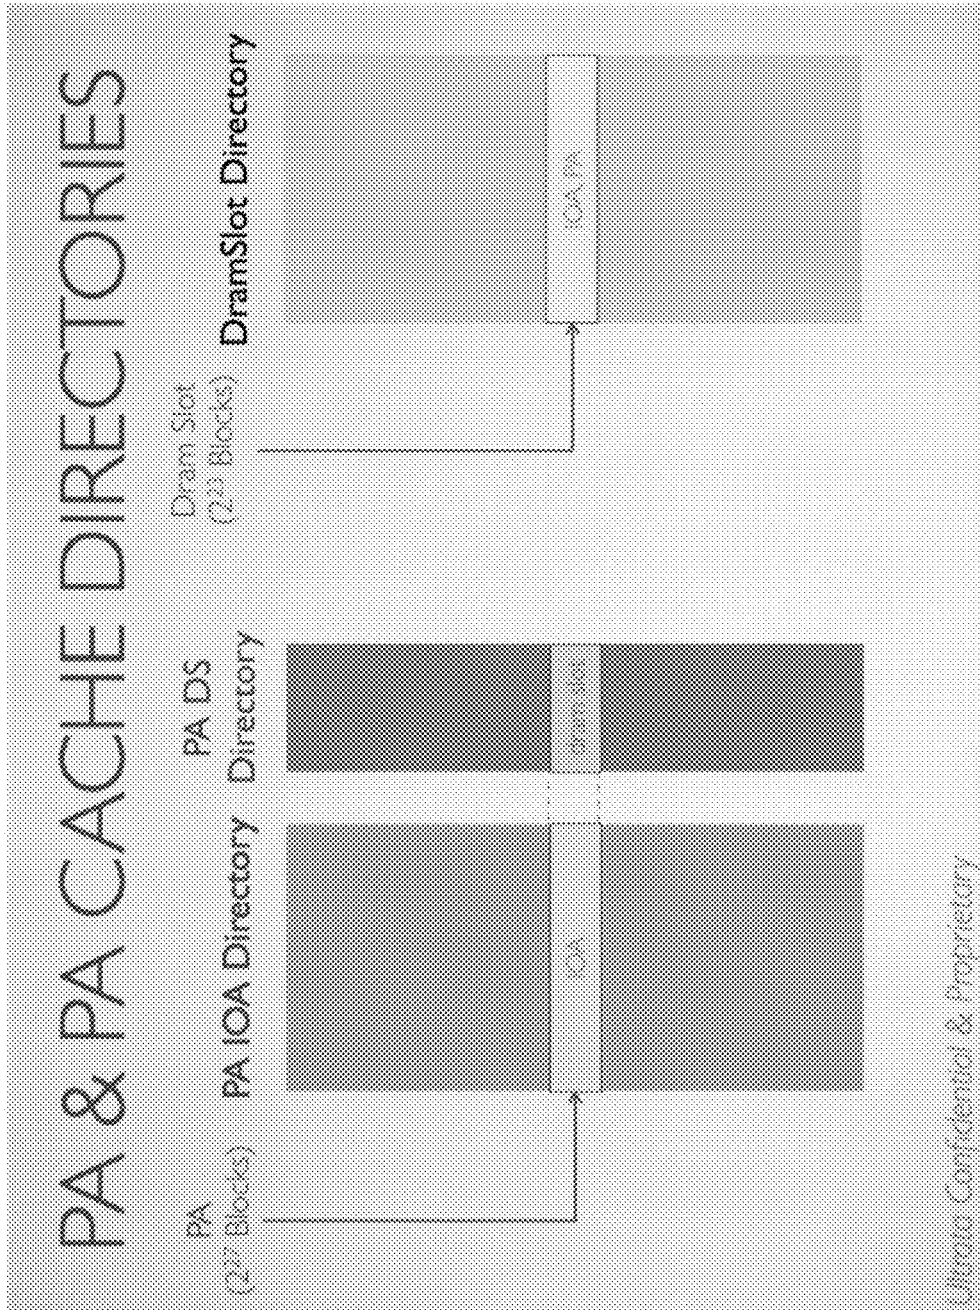


FIG. 36



**MEMORY FABRIC OPERATIONS AND  
COHERENCY USING FAULT TOLERANT  
OBJECTS**

CROSS-REFERENCE TO RELATED  
APPLICATIONS

**[0001]** The present application claims benefit under 35 USC 119(e) of U.S. Provisional Application No. 62/264,652, filed on Dec. 8, 2015 by Frank et al and entitled “Infinite Memory Fabric Operations, Coherency, and Interfaces,” of which the entire disclosure is incorporated herein by reference for all purposes.

**[0002]** The present application is also related to the following co-pending and commonly assigned U.S. Patent Applications:

**[0003]** U.S. patent application Ser. No. 15/001,320, filed on Jan. 20, 2016, by Frank and entitled “Object Based Memory Fabric;”

**[0004]** U.S. patent application Ser. No. 15/001,332, filed on Jan. 20, 2016, by Frank and entitled “Trans-Cloud Object Based Memory;”

**[0005]** U.S. patent application Ser. No. 15/001,340, filed on Jan. 20, 2016, by Frank and entitled “Universal Single Level Object Memory Address Space;”

**[0006]** U.S. patent application Ser. No. 15/001,343, filed on Jan. 20, 2016, by Frank and entitled “Object Memory Fabric Performance Acceleration;”

**[0007]** U.S. patent application Ser. No. 15/001,451, filed on Jan. 20, 2016, by Frank and entitled “Distributed Index for Fault Tolerant Object Memory Fabric;”

**[0008]** U.S. patent application Ser. No. 15/001,494, filed on Jan. 20, 2016, by Frank and entitled “Implementation of an Object Memory Centric Cloud;”

**[0009]** U.S. patent application Ser. No. 15/001,524, filed on Jan. 20, 2016, by Frank and entitled “Managing Metadata in an Object Memory Fabric;”

**[0010]** U.S. patent application Ser. No. 15/001,652, filed on Jan. 20, 2016, by Frank and entitled “Utilization of a Distributed Index to Provide Object Memory Fabric Coherency;”

**[0011]** U.S. patent application Ser. No. 15/001,366, filed on Jan. 20, 2016, by Frank and entitled “Object Memory Data Flow Instruction Execution;”

**[0012]** U.S. patent application Ser. No. 15/001,490, filed on Jan. 20, 2016, by Frank and entitled “Object Memory Data Flow Triggers;”

**[0013]** U.S. patent application Ser. No. 15/001,526, filed on Jan. 20, 2016, by Frank and entitled “Object Memory Instruction Set;”

**[0014]** U.S. patent application Ser. No. 15/168,965 filed on May 31, 2016 by Frank and entitled “Infinite Memory Fabric Streams and APIs;”

**[0015]** U.S. patent application Ser. No. 15/169,580 filed on May 31, 2016 by Frank and entitled “Infinite Memory Fabric Hardware Implementation with Memory;”

**[0016]** U.S. patent application Ser. No. 15/169,585 filed on May 31, 2016 by Frank and entitled “Infinite Memory Fabric Hardware Implementation with Router;”

**[0017]** U.S. patent application Ser. No. 15/371,393 (Attorney Docket Number 8620-15) filed concurrent herewith and entitled “Memory Fabric Software Implementation;” and

**[0018]** U.S. Patent Application No. \_\_\_\_\_ (Attorney Docket Number 8620-17) filed concurrent herewith and

entitled “Object Memory Interfaces Across Shared Links” of which the entire disclosure of each is incorporated herein by reference for all purposes.

BACKGROUND

**[0019]** Embodiments of the present invention relate generally to methods and systems for improving performance of processing nodes in a fabric and more particularly to changing the way in which processing, memory, storage, network, and cloud computing, are managed to significantly improve the efficiency and performance of commodity hardware.

**[0020]** As the size and complexity of data and the processes performed thereon continually increases, computer hardware is challenged to meet these demands. Current commodity hardware and software solutions from established server, network and storage providers are unable to meet the demands of Cloud Computing and Big Data environments. This is due, at least in part, to the way in which processing, memory, and storage are managed by those systems. Specifically, processing is separated from memory which is in turn separated from storage in current systems and each of processing, memory, and storage is managed separately by software. Each server and other computing device (referred to herein as a node) is in turn separated from other nodes by a physical computer network, managed separately by software and in turn the separate processing, memory, and storage associated with each node are managed by software on that node.

**[0021]** FIG. 1 is a block diagram illustrating an example of the separation data storage, memory, and processing within prior art commodity servers and network components. This example illustrates a system **100** in which commodity servers **105** and **110** are communicatively coupled with each other via a physical network **115** and network software **155** as known in the art. Also as known in the art, the servers can each execute any number of one or more applications **120a**, **120b**, **120c** of any variety. As known in the art, each application **120a**, **120b**, **120c** executes on a processor (not shown) and memory (not shown) of the server **105** and **110** using data stored in physical storage **150**. Each server **105** and **110** maintains a directory **125** mapping the location of the data used by the applications **120a**, **120b**, **120c**. Additionally, each server implements for each executing application **120a**, **120b**, **120c** a software stack which includes an application representation **130** of the data, a database representation **135**, a file system representation **140**, and a storage representation **145**.

**[0022]** While effective, there are three reasons that such implementations on current commodity hardware and software solutions from established server, network and storage providers are unable to meet the increasing demands of Cloud Computing and Big Data environments. One reason for the shortcomings of these implementations is their complexity. The software stack must be in place and every application must manage the separation of storage, memory, and processing as well as applying parallel server resources. Each application must trade-off algorithm parallelism, data organization and data movement which is extremely challenging to get correct, let alone considerations of performance and economics. This tends to lead to implementation of more batch oriented solutions in the applications, rather than the integrated real-time solutions preferred by most businesses. Additionally, separation of storage, memory, and processing, in such implementations also creates significant

inefficiency for each layer of the software stack to find, move, and access a block of data due to the required instruction execution and latencies of each layer of the software stack and between the layers. Furthermore, this inefficiency limits the economic scaling possible and limits the data-size for all but the most extremely parallel algorithms. The reason for the latter is that the efficiency with which servers (processors or threads) can interact limits the amount of parallelism due to Amdahl's law. Hence, there is a need for improved methods and systems for managing processing, memory, and storage to significantly improve the performance of processing nodes.

#### BRIEF SUMMARY

**[0023]** Embodiments of the invention provide systems and methods for managing processing, memory, storage, network, and cloud computing to significantly improve the efficiency and performance of processing nodes. Embodiments described herein can implement an object-based memory fabric in which memory objects in the memory fabric are distributed and tracked across a hierarchy of processing nodes. Each processing node can track memory objects and blocks within the memory objects that are present on paths from that node toward its leaf nodes in the hierarchy. Additionally, each processing node can utilize the same algorithms for memory object management such as memory object creation, block allocation, block coherency, etc. In this way, each higher level of the hierarchy creates an ever-larger cache which can significantly reduce the bandwidth in and out of the processing nodes at that level.

**[0024]** According to one embodiment, fault tolerance capability can be implemented based on this hierarchical distribution and tracking by enabling memory objects, on a per-object basis, to be stored in more than a single node. This distribution of memory objects across multiple nodes can be across the hierarchy and/or across multiple physical locations. Memory object fault tolerance copies can be handled by a block coherency mechanism as part of memory fabric operation. In this way, each memory object can be made to be present on multiple different nodes. The memory object can be contained as a whole within each of the multiple nodes or at a given level of the hierarchy or may be stored as different portions with each portion being contained within multiple different nodes.

**[0025]** According to one embodiment, a hardware-based processing node of a plurality of hardware-based processing nodes in an object memory fabric can comprise a memory module storing and managing a plurality of memory objects in a hierarchy of the object memory fabric. Each memory object can be created natively within the memory module, accessed using a single memory reference instruction without Input/Output (I/O) instructions, and managed by the memory module at a single memory layer. The object memory fabric can distribute and track the memory objects across the hierarchy of the object memory fabric and the plurality of hardware-based processing nodes on a per-object basis. Distributing the memory objects across the hierarchy of the object memory fabric and the plurality of hardware-based processing nodes can comprise storing, on a per-object basis, each memory object on two or more nodes of the plurality of hardware-based processing nodes of the object memory fabric. The two or more nodes of the plurality of hardware-based processing nodes can be remote

from each other in the hierarchy of the object memory fabric and/or in different physical locations.

**[0026]** Tracking the memory objects across the hierarchy of the object memory fabric can comprise tracking, by the hardware-based processing node, the memory objects and blocks within the memory objects that are present on the hardware-based processing node. Additionally or alternatively, tracking the memory objects across the hierarchy of the object memory fabric can comprise tracking the memory objects and blocks within the memory objects that are present on each level of the hierarchy of the object memory fabric and branches from each level of the hierarchy towards leaves of the hierarchy of the object memory fabric. The hardware-based processing node can utilize a same algorithm for object management as each other node of the plurality of hardware-based processing nodes. The algorithm for object management can comprise an algorithm for object creation, an algorithm for block allocation, and/or an algorithm for block coherency.

**[0027]** In one implementation, the hardware-based processing node can comprise a Dual In-line Memory Module (DIMM) card. In other cases, the hardware-based processing node can comprise a commodity server and wherein the memory module comprises a DIMM card installed within the commodity server. In other cases, the hardware-based processing node can comprise a mobile computing device. In yet other implementations, the hardware-based processing node can comprise a single chip.

**[0028]** According to another embodiment, an object memory fabric can comprise a plurality of hardware-based processing nodes. Each hardware-based processing node can comprise a memory module storing and managing a plurality of memory objects in a hierarchy of the object memory fabric. Each memory object can be created natively within the memory module, accessed using a single memory reference instruction without Input/Output (I/O) instructions, and managed at a single memory layer. A node router can be communicatively coupled with each of the one or more memory modules of the node and can be adapted to route memory objects or portions of memory objects between the one or more memory modules of the node. One or more inter-node routers can be communicatively coupled with each node router. Each of the plurality of nodes of the object memory fabric can be communicatively coupled with at least one of the inter-node routers and can be adapted to route memory objects or portions of memory objects between the plurality of nodes.

**[0029]** The object memory fabric can distribute and track the memory objects across the hierarchy of the object memory fabric and the plurality of hardware-based processing nodes on a per-object basis. Distributing the memory objects across the hierarchy of the object memory fabric and the plurality of hardware-based processing nodes can comprise storing, on a per-object basis, each memory object on two or more nodes of the plurality of hardware-based processing nodes of the object memory fabric. The two or more nodes of the plurality of hardware-based processing nodes can be remote from each other in the hierarchy of the object memory fabric and/or in different physical locations.

**[0030]** Tracking the memory objects across the hierarchy of the object memory fabric can comprise tracking, by each hardware-based processing node, the memory objects and blocks within the memory objects that are present on the hardware-based processing node. Tracking the memory

objects across the hierarchy of the object memory fabric can also comprise tracking the memory objects and blocks within the memory objects that are present on each level of the hierarchy of the object memory fabric and branches from each level of the hierarchy towards leaves of the hierarchy of the object memory fabric. Each hardware-based processing node can utilize a same algorithm for object management. The algorithm for object management can comprise an algorithm for object creation, an algorithm for block allocation, and /or an algorithm for block coherency.

**[0031]** According to yet another embodiment, a method for providing coherency and fault tolerance in an object memory fabric including a plurality of hardware-based processing nodes can comprise creating, by the hardware-based processing nodes of the object-based memory fabric, each memory object natively within a memory module of the hardware-based processing node, accessing, by the hardware-based processing nodes, each memory object using a single memory reference instruction without Input/Output (I/O) instructions, and managing, by the hardware-based processing nodes, each memory object within the memory module at a single memory layer. The memory objects can be distributed and tracked across a hierarchy of the object memory fabric and the plurality of hardware-based processing nodes on a per-object basis. Distributing the memory objects across the hierarchy of the object memory fabric and the plurality of hardware-based processing nodes can comprise storing, on a per-object basis, each memory object on two or more nodes of the plurality of hardware-based processing nodes of the object memory fabric. The two or more nodes can be remote from each other in the hierarchy of the object memory fabric and/or in different physical locations. Tracking the memory objects across the hierarchy of the object memory fabric can comprise tracking the memory objects and blocks within the memory objects that are present on the hardware-based processing node and/or that are present on each level of the hierarchy of the object memory fabric and branches from each level of the hierarchy towards leaves of the hierarchy of the object memory fabric.

#### BRIEF DESCRIPTION OF THE DRAWINGS

**[0032]** FIG. 1 is a block diagram illustrating an example of the separation data storage, memory, processing, network, and cloud computing within prior art commodity servers and network components.

**[0033]** FIG. 2 is a block diagram illustrating components of an exemplary distributed system in which various embodiments of the present invention may be implemented.

**[0034]** FIG. 3 is a block diagram illustrating an exemplary computer system in which embodiments of the present invention may be implemented.

**[0035]** FIG. 4 is a block diagram illustrating an exemplary object memory fabric architecture according to one embodiment of the present invention.

**[0036]** FIG. 5 is a block diagram illustrating an exemplary memory fabric object memory according to one embodiment of the present invention.

**[0037]** FIG. 6 is a block diagram illustrating an exemplary object memory dynamics and physical organization according to one embodiment of the present invention.

**[0038]** FIG. 7 is a block diagram illustrating aspects of object memory fabric hierarchy of object memory, which

localizes working sets and allows for virtually unlimited scalability, according to one embodiment of the present invention.

**[0039]** FIG. 8 is a block diagram illustrating aspects of an example relationship between object address space, virtual address, and physical address, according to one embodiment of the present invention.

**[0040]** FIG. 9 is a block diagram illustrating aspects of an example relationship between object sizes and object address space pointers, according to one embodiment of the present invention.

**[0041]** FIG. 10 is a block diagram illustrating aspects of an example object memory fabric distributed object memory and index structure, according to one embodiment of the present invention.

**[0042]** FIG. 11 illustrates aspects of an object memory hit case that executes completely within the object memory, according to one embodiment of the present invention.

**[0043]** FIG. 12 illustrates aspects of an object memory miss case and the distributed nature of the object memory and object index, according to one embodiment of the present invention.

**[0044]** FIG. 13 is a block diagram illustrating aspects of an example of leaf level object memory in view of the object memory fabric distributed object memory and index structure, according to one embodiment of the present invention.

**[0045]** FIG. 14 is a block diagram illustrating aspects of an example of object memory fabric router object index structure, according to one embodiment of the present invention.

**[0046]** FIGS. 15A and 15B are block diagrams illustrating aspects of example index tree structures, including node index tree structure and leaf index tree, according to one embodiment of the present invention.

**[0047]** FIG. 16 is a block diagram illustrating aspects of an example physical memory organization, according to one embodiment of the present invention.

**[0048]** FIG. 17A is a block diagram illustrating aspects of example object addressing, according to one embodiment of the present invention.

**[0049]** FIG. 17B is a block diagram illustrating aspects of example object memory fabric pointer and block addressing, according to one embodiment of the present invention.

**[0050]** FIG. 18 is a block diagram illustrating aspects of example object metadata, according to one embodiment of the present invention.

**[0051]** FIG. 19 is a block diagram illustrating aspects of an example micro-thread model, according to one embodiment of the present invention.

**[0052]** FIG. 20 is a block diagram illustrating aspects of an example relationship of code, frame, and object, according to one embodiment of the present invention.

**[0053]** FIG. 21 is a block diagram illustrating aspects of an example of micro-thread concurrency, according to one embodiment of the present invention.

**[0054]** FIG. 22A is a block diagram illustrating an example of streams present on a node with a hardware-based object memory fabric inter-node object router, in accordance with certain embodiments of the present disclosure.

**[0055]** FIG. 22B is a block diagram illustrating an example of software emulation of object memory and router on the node, in accordance with certain embodiments of the present disclosure.

[0056] FIG. 23 is a block diagram illustrating an example of streams within a memory fabric router, in accordance with certain embodiments of the present disclosure.

[0057] FIG. 24 is a block diagram illustrating a product family hardware implementation architecture, in accordance with certain embodiments of the present disclosure.

[0058] FIG. 25 is a block diagram illustrating an alternative product family hardware implementation architecture, in accordance with certain embodiments of the present disclosure.

[0059] FIG. 26 is a block diagram illustrating a memory fabric server view of a hardware implementation architecture, in accordance with certain embodiments of the present disclosure.

[0060] FIG. 27 is a block diagram illustrating a memory module view of a hardware implementation architecture, in accordance with certain embodiments of the present disclosure.

[0061] FIG. 28 is a block diagram illustrating a memory module view of a hardware implementation architecture, in accordance with an alternative embodiment of the present disclosure.

[0062] FIG. 29 is a block diagram illustrating a node router view of a hardware implementation architecture, in accordance with certain embodiments of the present disclosure.

[0063] FIG. 30 is a block diagram illustrating an inter-node router view of a hardware implementation architecture, in accordance with certain embodiments of the present disclosure.

[0064] FIG. 31 is a block diagram illustrating a memory fabric router view of a hardware implementation architecture, in accordance with certain embodiments of the present disclosure.

[0065] FIG. 32 is a block diagram illustrating object memory fabric functions that can replace software functions according to one embodiment of the present disclosure.

[0066] FIG. 33 is a block diagram illustrating an object memory fabric software stack according to one embodiment of the present disclosure.

[0067] FIG. 34 is a block diagram illustrating a summary of memory module caching according to one embodiment.

[0068] FIG. 35 is a diagram illustrating an exemplary partitioning of the DDR4 dram cache for several functions related to the memory module according to one embodiment.

[0069] FIG. 36 is a block diagram illustrating node and leaf caching according to one embodiment.

#### DETAILED DESCRIPTION

[0070] In the following description, for the purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of various embodiments of the present invention. It will be apparent, however, to one skilled in the art that embodiments of the present invention may be practiced without some of these specific details. In other instances, well-known structures and devices are shown in block diagram form.

[0071] The ensuing description provides exemplary embodiments only, and is not intended to limit the scope, applicability, or configuration of the disclosure. Rather, the ensuing description of the exemplary embodiments will provide those skilled in the art with an enabling description for implementing an exemplary embodiment. It should be

understood that various changes may be made in the function and arrangement of elements without departing from the spirit and scope of the invention as set forth in the appended claims.

[0072] Specific details are given in the following description to provide a thorough understanding of the embodiments. However, it will be understood by one of ordinary skill in the art that the embodiments may be practiced without these specific details. For example, circuits, systems, networks, processes, and other components may be shown as components in block diagram form in order not to obscure the embodiments in unnecessary detail. In other instances, well-known circuits, processes, algorithms, structures, and techniques may be shown without unnecessary detail in order to avoid obscuring the embodiments.

[0073] Also, it is noted that individual embodiments may be described as a process which is depicted as a flowchart, a flow diagram, a data flow diagram, a structure diagram, or a block diagram. Although a flowchart may describe the operations as a sequential process, many of the operations can be performed in parallel or concurrently. In addition, the order of the operations may be re-arranged. A process is terminated when its operations are completed, but could have additional steps not included in a figure. A process may correspond to a method, a function, a procedure, a subroutine, a subprogram, etc. When a process corresponds to a function, its termination can correspond to a return of the function to the calling function or the main function.

[0074] The term “machine-readable medium” includes, but is not limited to portable or fixed storage devices, optical storage devices, wireless channels and various other mediums capable of storing, containing or carrying instruction(s) and/or data. A code segment or machine-executable instructions may represent a procedure, a function, a subprogram, a program, a routine, a subroutine, a module, a software package, a class, or any combination of instructions, data structures, or program statements. A code segment may be coupled to another code segment or a hardware circuit by passing and/or receiving information, data, arguments, parameters, data, etc. may be passed, forwarded, or transmitted via any suitable means including memory sharing, message passing, token passing, network transmission, etc. Various other terms used herein are now defined for the sake of clarity.

[0075] Virtual memory is a memory management technique that gives the illusion to each software process that memory is as large as the virtual address space. The operating system in conjunction with differing degrees of hardware manages the physical memory as a cache of the virtual address space, which is placed in secondary storage and accessible through Input/Output instructions. Virtual memory is separate from, but can interact with, a file system.

[0076] A single level store is an extension of virtual memory in which there are no files, only persistent objects or segments which are mapped into a processes' address space using virtual memory techniques. The entire storage of the computing system is thought of as a segment and address within a segment. Thus at least three separate address spaces, i.e., physical memory address/node, virtual address/process, and secondary storage address/disk, are managed by software.

[0077] Object storage refers to the way units of storage called objects are organized. Every object consists of a

container that holds three things: actual data; expandable metadata; and a globally unique identifier referred to herein as the object address. The metadata of the object is used to define contextual information about the data and how it should be used and managed including relationship to other objects.

**[0078]** The object address space is managed by software over storage devices, nodes, and network to find an object without knowing its physical location. Object storage is separate from virtual memory and single level store, but can certainly inter-operate through software.

**[0079]** Block storage consists of evenly sized blocks of data with an address based on a physical location and without metadata.

**[0080]** A network address is a physical address of a node within an IP network that is associated with a physical location.

**[0081]** A node or processing node is a physical unit of computing delineated by a shared physical memory that be addressed by any processor within the node.

**[0082]** Object memory is an object store directly accessible as memory by processor memory reference instructions and without implicit or explicit software or Input/Output instructions required. Object capabilities are directly provided within the object memory to processing through memory reference instructions.

**[0083]** An object memory fabric connects object memory modules and nodes into a single object memory where any object is local to any object memory module by direct management, in hardware, of object data, meta-data and object address.

**[0084]** An object router routes objects or portions of objects in an object memory fabric based on an object address. This is distinct from a conventional router which forwards data packets to appropriate part of a network based on a network address.

**[0085]** Embodiments may be implemented by hardware, software, firmware, middleware, microcode, hardware description languages, or any combination thereof. When implemented in software, firmware, middleware or microcode, the program code or code segments to perform the necessary tasks may be stored in a machine readable medium. A processor(s) may perform the necessary tasks.

**[0086]** Embodiments of the invention provide systems and methods for managing processing, memory, storage, network, and cloud computing to significantly improve the efficiency and performance of processing nodes. Embodiments described herein can be implemented in a set of hardware components that, in essence, change the way in which processing, memory, and storage, network, and cloud computing are managed by breaking down the artificial distinctions between processing, memory, storage and networking in today's commodity solutions to significantly improve the efficiency and performance of commodity hardware. For example, the hardware elements can include a standard format memory module, such as a (DIMM) and a set of one or more object routers. The memory module can be added to commodity or "off-the-shelf" hardware such a server node and acts as a big data accelerator within that node. Object routers can be used to interconnect two or more servers or other nodes adapted with the memory modules and help to manage processing, memory, and storage across these different servers. Nodes can be physically close or far apart. Together, these hardware components can be used

with commodity servers or other types of computing nodes in any combination to implement the embodiments described herein.

**[0087]** According to one embodiment, such hardware components can implement an object-based memory which manages the objects within the memory and at the memory layer rather than in the application layer. That is, the objects and associated properties are implemented and managed natively in memory enabling the object memory system to provide increased functionality without any software and increasing performance by dynamically managing object characteristics including, but not limited to persistence, location and processing. Object properties can also propagate up to higher application levels.

**[0088]** Such hardware components can also eliminate the distinction between memory (temporary) and storage (persistent) by implementing and managing both within the objects. These components can eliminate the distinction between local and remote memory by transparently managing the location of objects (or portions of objects) so all objects appear simultaneously local to all nodes. These components can also eliminate the distinction between processing and memory through methods of the objects to place the processing within the memory itself.

**[0089]** According to one embodiment, such hardware components can eliminate typical size constraints on memory space of the commodity servers imposed by address sizes. Rather, physical addressing can be managed within the memory objects themselves and the objects can in turn be accessed and managed through the object name space.

**[0090]** Embodiment described herein can provide transparent and dynamic performance acceleration, especially with big data or other memory intensive applications by reducing or eliminating overhead typically associated with memory management, storage management, networking and data directories. Rather, management of the memory objects at the memory level can significantly shorten the pathways between storage and memory and between memory and processing, thereby eliminating the associated overhead between each. Various additional details of embodiments of the present invention will be described below with reference to the figures.

**[0091]** FIG. 2 is a block diagram illustrating components of an exemplary distributed system in which various embodiments of the present invention may be implemented. In the illustrated embodiment, distributed system 200 includes one or more client computing devices 202, 204, 206, and 208, which are configured to execute and operate a client application such as a web browser, proprietary client, or the like over one or more network(s) 210. Server 212 may be communicatively coupled with remote client computing devices 202, 204, 206, and 208 via network 210.

**[0092]** In various embodiments, server 212 may be adapted to run one or more services or software applications provided by one or more of the components of the system. In some embodiments, these services may be offered as web-based or cloud services or under a Software as a Service (SaaS) model to the users of client computing devices 202, 204, 206, and/or 208. Users operating client computing devices 202, 204, 206, and/or 208 may in turn utilize one or more client applications to interact with server 212 to utilize the services provided by these components. For the sake of clarity, it should be noted that server 212 and database 214, 216 can correspond to server 105 described above with

reference to FIG. 1. Network **210** can be part of or an extension to physical network **115**. It should also be understood that there can be any number of client computing devices **202**, **204**, **206**, **208** and servers **212**, each with one or more databases **214**, **216**.

**[0093]** In the configuration depicted in the figure, the software components **218**, **220** and **222** of system **200** are shown as being implemented on server **212**. In other embodiments, one or more of the components of system **200** and/or the services provided by these components may also be implemented by one or more of the client computing devices **202**, **204**, **206**, and/or **208**. Users operating the client computing devices may then utilize one or more client applications to use the services provided by these components. These components may be implemented in hardware, firmware, software, or combinations thereof. It should be appreciated that various different system configurations are possible, which may be different from distributed system **200**. The embodiment shown in the figure is thus one example of a distributed system for implementing an embodiment system and is not intended to be limiting.

**[0094]** Client computing devices **202**, **204**, **206**, and/or **208** may be portable handheld devices (e.g., an iPhone®, cellular telephone, an iPad®, computing tablet, a personal digital assistant (PDA)) or wearable devices (e.g., a Google Glass® head mounted display), running software such as Microsoft Windows Mobile®, and/or a variety of mobile operating systems such as iOS, Windows Phone, Android, BlackBerry 10, Palm OS, and the like, and being Internet, e-mail, short message service (SMS), Blackberry®, or other communication protocol enabled. The client computing devices can be general purpose personal computers including, by way of example, personal computers and/or laptop computers running various versions of Microsoft Windows®, Apple Macintosh®, and/or Linux operating systems. The client computing devices can be workstation computers running any of a variety of commercially-available UNIX® or UNIX-like operating systems, including without limitation the variety of GNU/Linux operating systems, such as for example, Google Chrome OS. Alternatively, or in addition, client computing devices **202**, **204**, **206**, and **208** may be any other electronic device, such as a thin-client computer, an Internet-enabled gaming system (e.g., a Microsoft Xbox gaming console with or without a Kinect® gesture input device), and/or a personal messaging device, capable of communicating over network(s) **210**.

**[0095]** Although exemplary distributed system **200** is shown with four client computing devices, any number of client computing devices may be supported. Other devices, such as devices with sensors, etc., may interact with server **212**.

**[0096]** Network(s) **210** in distributed system **200** may be any type of network familiar to those skilled in the art that can support data communications using any of a variety of commercially-available protocols, including without limitation TCP/IP (Transmission Control Protocol/Internet Protocol), SNA (Systems Network Architecture), IPX (Internet Packet Exchange), AppleTalk, and the like. Merely by way of example, network(s) **210** can be a Local Area Network (LAN), such as one based on Ethernet, Token-Ring and/or the like. Network(s) **210** can be a wide-area network and the Internet. It can include a virtual network, including without limitation a Virtual Private Network (VPN), an intranet, an extranet, a Public Switched Telephone Network (PSTN), an

infra-red network, a wireless network (e.g., a network operating under any of the Institute of Electrical and Electronics (IEEE) 802.11 suite of protocols, Bluetooth®, and/or any other wireless protocol); and/or any combination of these and/or other networks. Elements of such networks can have an arbitrary distance, i.e., can be remote or co-located. Software Defined Networks (SDNs) can be implemented with a combination of dumb routers and software running on servers.

**[0097]** Server **212** may be composed of one or more general purpose computers, specialized server computers (including, by way of example, Personal Computer (PC) servers, UNIX® servers, mid-range servers, mainframe computers, rack-mounted servers, etc.), server farms, server clusters, or any other appropriate arrangement and/or combination. In various embodiments, server **212** may be adapted to run one or more services or software applications described in the foregoing disclosure. For example, server **212** may correspond to a server for performing processing described above according to an embodiment of the present disclosure.

**[0098]** Server **212** may run an operating system including any of those discussed above, as well as any commercially available server operating system. Server **212** may also run any of a variety of additional server applications and/or mid-tier applications, including HyperText Transport Protocol (HTTP) servers, File Transfer Protocol (FTP) servers, Common Gateway Interface (CGI) servers, JAVA® servers, database servers, and the like. Exemplary database servers include without limitation those commercially available from Oracle, Microsoft, Sybase, International Business Machines (IBM), and the like.

**[0099]** In some implementations, server **212** may include one or more applications to analyze and consolidate data feeds and/or event updates received from users of client computing devices **202**, **204**, **206**, and **208**. As an example, data feeds and/or event updates may include, but are not limited to, Twitter® feeds, Facebook® updates or real-time updates received from one or more third party information sources and continuous data streams, which may include real-time events related to sensor data applications, financial tickers, network performance measuring tools (e.g., network monitoring and traffic management applications), click-stream analysis tools, automobile traffic monitoring, and the like. Server **212** may also include one or more applications to display the data feeds and/or real-time events via one or more display devices of client computing devices **202**, **204**, **206**, and **208**.

**[0100]** Distributed system **200** may also include one or more databases **214** and **216**. Databases **214** and **216** may reside in a variety of locations. By way of example, one or more of databases **214** and **216** may reside on a non-transitory storage medium local to (and/or resident in) server **212**. Alternatively, databases **214** and **216** may be remote from server **212** and in communication with server **212** via a network-based or dedicated connection. In one set of embodiments, databases **214** and **216** may reside in a Storage-Area Network (SAN). Similarly, any necessary files for performing the functions attributed to server **212** may be stored locally on server **212** and/or remotely, as appropriate. In one set of embodiments, databases **214** and **216** may include relational databases that are adapted to store, update, and retrieve data in response to commands, e.g., MySQL-formatted commands. Additionally or alternatively, server

**212** can provide and support big data processing on unstructured data including but not limited to Hadoop processing, NoSQL databases, graph databases etc. In yet other implementations, server **212** may perform non-database types of bog data applications including but not limited to machine learning.

**[0101]** FIG. 3 is a block diagram illustrating an exemplary computer system in which embodiments of the present invention may be implemented. The system **300** may be used to implement any of the computer systems described above. As shown in the figure, computer system **300** includes a processing unit **304** that communicates with a number of peripheral subsystems via a bus subsystem **302**. These peripheral subsystems may include a processing acceleration unit **306**, an I/O subsystem **308**, a storage subsystem **318** and a communications subsystem **324**. Storage subsystem **318** includes tangible computer-readable storage media **322** and a system memory **310**.

**[0102]** Bus subsystem **302** provides a mechanism for letting the various components and subsystems of computer system **300** communicate with each other as intended. Although bus subsystem **302** is shown schematically as a single bus, alternative embodiments of the bus subsystem may utilize multiple buses. Bus subsystem **302** may be any of several types of bus structures including a memory bus or memory controller, a peripheral bus, and a local bus using any of a variety of bus architectures. For example, such architectures may include an Industry Standard Architecture (ISA) bus, Micro Channel Architecture (MCA) bus, Enhanced ISA (EISA) bus, Video Electronics Standards Association (VESA) local bus, Peripheral Component Interconnect (PCI) bus, which can be implemented as a Mezzanine bus manufactured to the IEEE P1386.1 standard, or PCI enhanced (PCIe) bus.

**[0103]** Processing unit **304**, which can be implemented as one or more integrated circuits (e.g., a conventional microprocessor or microcontroller), controls the operation of computer system **300**. One or more processors may be included in processing unit **304**. These processors may include single core or multicore processors. In certain embodiments, processing unit **304** may be implemented as one or more independent processing units **332** and/or **334** with single or multicore processors included in each processing unit. In other embodiments, processing unit **304** may also be implemented as a quad-core processing unit formed by integrating two dual-core processors into a single chip.

**[0104]** In various embodiments, processing unit **304** can execute a variety of programs in response to program code and can maintain multiple concurrently executing programs or processes. At any given time, some or all of the program code to be executed can be resident in processor(s) **304** and/or in storage subsystem **318**. Through suitable programming, processor(s) **304** can provide various functionalities described above. Computer system **300** may additionally include a processing acceleration unit **306**, which can include a Digital Signal Processor (DSP), a special-purpose processor, and/or the like.

**[0105]** I/O subsystem **308** may include user interface input devices and user interface output devices. User interface input devices may include a keyboard, pointing devices such as a mouse or trackball, a touchpad or touch screen incorporated into a display, a scroll wheel, a click wheel, a dial, a button, a switch, a keypad, audio input devices with voice command recognition systems, microphones, and other

types of input devices. User interface input devices may include, for example, motion sensing and/or gesture recognition devices such as the Microsoft Kinect® motion sensor that enables users to control and interact with an input device, such as the Microsoft Xbox® 360 game controller, through a natural user interface using gestures and spoken commands. User interface input devices may also include eye gesture recognition devices such as the Google Glass® blink detector that detects eye activity (e.g., ‘blinking’ while taking pictures and/or making a menu selection) from users and transforms the eye gestures as input into an input device (e.g., Google Glass®). Additionally, user interface input devices may include voice recognition sensing devices that enable users to interact with voice recognition systems (e.g., Siri® navigator), through voice commands.

**[0106]** User interface input devices may also include, without limitation, three dimensional (3D) mice, joysticks or pointing sticks, gamepads and graphic tablets, and audio/visual devices such as speakers, digital cameras, digital camcorders, portable media players, webcams, image scanners, fingerprint scanners, barcode reader 3D scanners, 3D printers, laser rangefinders, and eye gaze tracking devices. Additionally, user interface input devices may include, for example, medical imaging input devices such as computed tomography, magnetic resonance imaging, position emission tomography, medical ultrasonography devices. User interface input devices may also include, for example, audio input devices such as MIDI keyboards, digital musical instruments and the like.

**[0107]** User interface output devices may include a display subsystem, indicator lights, or non-visual displays such as audio output devices, etc. The display subsystem may be a Cathode Ray Tube (CRT), a flat-panel device, such as that using a Liquid Crystal Display (LCD) or plasma display, a projection device, a touch screen, and the like. In general, use of the term “output device” is intended to include all possible types of devices and mechanisms for outputting information from computer system **300** to a user or other computer. For example, user interface output devices may include, without limitation, a variety of display devices that visually convey text, graphics and audio/video information such as monitors, printers, speakers, headphones, automotive navigation systems, plotters, voice output devices, and modems.

**[0108]** Computer system **300** may comprise a storage subsystem **318** that comprises software elements, shown as being currently located within a system memory **310**. System memory **310** may store program instructions that are loadable and executable on processing unit **304**, as well as data generated during the execution of these programs.

**[0109]** Depending on the configuration and type of computer system **300**, system memory **310** may be volatile (such as Random Access Memory (RAM)) and/or non-volatile (such as Read-Only Memory (ROM), flash memory, etc.) The RAM typically contains data and/or program modules that are immediately accessible to and/or presently being operated and executed by processing unit **304**. In some cases, system memory **310** can comprise one or more Double Data Rate fourth generation (DDR4) Dual Inline Memory Modules (DIMMs). In some implementations, system memory **310** may include multiple different types of memory, such as Static Random Access Memory (SRAM) or Dynamic Random Access Memory (DRAM). In some implementations, a Basic Input/Output System (BIOS), con-

taining the basic routines that help to transfer information between elements within computer system 300, such as during start-up, may typically be stored in the ROM. By way of example, and not limitation, system memory 310 also illustrates application programs 312, which may include client applications, Web browsers, mid-tier applications, Relational Database Management Systems (RDBMS), etc., program data 314, and an operating system 316. By way of example, operating system 316 may include various versions of Microsoft Windows®, Apple Macintosh®, and/or Linux operating systems, a variety of commercially-available UNIX® or UNIX-like operating systems (including without limitation the variety of GNU/Linux operating systems, the Google Chrome® OS, and the like) and/or mobile operating systems such as iOS, Windows® Phone, Android® OS, BlackBerry® 10 OS, and Palm® OS operating systems.

[0110] Storage subsystem 318 may also provide a tangible computer-readable storage medium for storing the basic programming and data constructs that provide the functionality of some embodiments. Software (programs, code modules, instructions) that when executed by a processor provide the functionality described above may be stored in storage subsystem 318. These software modules or instructions may be executed by processing unit 304. Storage subsystem 318 may also provide a repository for storing data used in accordance with the present invention.

[0111] Storage subsystem 300 may also include a computer-readable storage media reader 320 that can further be connected to computer-readable storage media 322. Together and, optionally, in combination with system memory 310, computer-readable storage media 322 may comprehensively represent remote, local, fixed, and/or removable storage devices plus storage media for temporarily and/or more permanently containing, storing, transmitting, and retrieving computer-readable information.

[0112] Computer-readable storage media 322 containing code, or portions of code, can also include any appropriate media known or used in the art, including storage media and communication media, such as but not limited to, volatile and non-volatile, removable and non-removable media implemented in any method or technology for storage and/or transmission of information. This can include tangible computer-readable storage media such as RAM, ROM, Electronically Erasable Programmable ROM (EEPROM), flash memory or other memory technology, CD-ROM, Digital Versatile Disk (DVD), or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or other tangible computer readable media. This can also include nontangible computer-readable media, such as data signals, data transmissions, or any other medium which can be used to transmit the desired information and which can be accessed by computing system 300.

[0113] By way of example, computer-readable storage media 322 may include a hard disk drive that reads from or writes to non-removable, nonvolatile magnetic media, a magnetic disk drive that reads from or writes to a removable, nonvolatile magnetic disk, and an optical disk drive that reads from or writes to a removable, nonvolatile optical disk such as a CD ROM, DVD, and Blu-Ray® disk, or other optical media. Computer-readable storage media 322 may include, but is not limited to, Zip® drives, flash memory cards, Universal Serial Bus (USB) flash drives, Secure

Digital (SD) cards, DVD disks, digital video tape, and the like. Computer-readable storage media 322 may also include, Solid-State Drives (SSD) based on non-volatile memory such as flash-memory based SSDs, enterprise flash drives, solid state ROM, and the like, SSDs based on volatile memory such as solid state RAM, dynamic RAM, static RAM, DRAM-based SSDs, Magnetoresistive RAM (MRAM) SSDs, and hybrid SSDs that use a combination of DRAM and flash memory based SSDs. The disk drives and their associated computer-readable media may provide non-volatile storage of computer-readable instructions, data structures, program modules, and other data for computer system 300.

[0114] Communications subsystem 324 provides an interface to other computer systems and networks. Communications subsystem 324 serves as an interface for receiving data from and transmitting data to other systems from computer system 300. For example, communications subsystem 324 may enable computer system 300 to connect to one or more devices via the Internet. In some embodiments communications subsystem 324 can include Radio Frequency (RF) transceiver components for accessing wireless voice and/or data networks (e.g., using cellular telephone technology, advanced data network technology, such as 3G, 4G or Enhanced Data rates for Global Evolution (EDGE), WiFi (IEEE 802.11 family standards, or other mobile communication technologies, or any combination thereof), Global Positioning System (GPS) receiver components, and/or other components. In some embodiments communications subsystem 324 can provide wired network connectivity (e.g., Ethernet) in addition to or instead of a wireless interface. In some cases, communications subsystem 324 can be implemented in whole or in part as one or more PCIe cards.

[0115] In some embodiments, communications subsystem 324 may also receive input communication in the form of structured and/or unstructured data feeds 326, event streams 328, event updates 330, and the like on behalf of one or more users who may use computer system 300.

[0116] By way of example, communications subsystem 324 may be configured to receive data feeds 326 in real-time from users of social networks and/or other communication services such as Twitter® feeds, Facebook® updates, web feeds such as Rich Site Summary (RSS) feeds, and/or real-time updates from one or more third party information sources.

[0117] Additionally, communications subsystem 324 may also be configured to receive data in the form of continuous data streams, which may include event streams 328 of real-time events and/or event updates 330, that may be continuous or unbounded in nature with no explicit end. Examples of applications that generate continuous data may include, for example, sensor data applications, financial tickers, network performance measuring tools (e.g. network monitoring and traffic management applications), click-stream analysis tools, automobile traffic monitoring, and the like.

[0118] Communications subsystem 324 may also be configured to output the structured and/or unstructured data feeds 326, event streams 328, event updates 330, and the like to one or more databases that may be in communication with one or more streaming data source computers coupled to computer system 300.



[0119] Computer system 300 can be one of various types, including a handheld portable device (e.g., an iPhone® cellular phone, an iPad® computing tablet, a PDA), a wearable device (e.g., a Google Glass® head mounted display), a PC, a workstation, a mainframe, a kiosk, a server rack, or any other data processing system.

[0120] Due to the ever-changing nature of computers and networks, the description of computer system 300 depicted in the figure is intended only as a specific example. Many other configurations having more or fewer components than the system depicted in the figure are possible. For example, customized hardware might also be used and/or particular elements might be implemented in hardware, firmware, software (including applets), or a combination. Further, connection to other computing devices, such as network input/output devices, may be employed. Based on the disclosure and teachings provided herein, a person of ordinary skill in the art will appreciate other ways and/or methods to implement the various embodiments.

[0121] As introduced above, embodiments of the invention provide systems and methods for managing processing, memory, storage, network, and cloud computing to significantly improve the efficiency and performance of processing nodes such as any of the servers or other computers or computing devices described above. Embodiments described herein can be implemented in a set of hardware components that, in essence, change the way in which processing, memory, storage, network, and cloud are managed by breaking down the artificial distinctions between processing, memory, storage and networking in today's commodity solutions to significantly improve the performance of commodity hardware. For example, the hardware elements can include a standard format memory module, such as a Dual Inline Memory Module (DIMM), which can be added to any of the computer systems described above. For example, the memory module can be added to commodity or "off-the-shelf" hardware such as a server node and acts as a big data accelerator within that node. The components can also include one or more object routers. Object routers can include, for example, a PCI express card added to the server node along with the memory module and one or more external object routers such as rack mounted routers, for example. Object routers can be used to interconnect two or more servers or other nodes adapted with the memory modules and help to manage processing, memory, and storage across these different servers. Object routers can forward objects or portions of objects based on object addresses and participate in operation of the object memory fabric. Together, these hardware components can be used with commodity servers or other types of computing nodes in any combination to implement an object memory fabric architecture.

[0122] FIG. 4 is a block diagram illustrating an exemplary object memory fabric architecture according to one embodiment of the present invention. As illustrated here, the architecture 400 comprises an object memory fabric 405 supporting any number of applications 410a-g. As will be described in greater detail below, this object memory fabric 405 can comprise any number of processing nodes such as one or more servers having installed one or more memory modules as described herein. These nodes can be interconnected by one or more internal and/or external object routers as described herein. While described as comprising one or more servers, it should be noted that the processing nodes of

the object memory fabric 405 can comprise any of a variety of different computers and/or computing devices adapted to operate within the object memory fabric 405 as described herein.

[0123] According to one embodiment, the object memory fabric 405 provides an object-based memory which manages memory objects within the memory of the nodes of the object memory fabric 405 and at the memory layer rather than in the application layer. That is, the objects and associated properties can be implemented and managed natively in the nodes of the object memory fabric 405 to provide increased functionality without any software and increasing efficiency and performance by dynamically managing object characteristics including, but not limited to persistence, location and processing. Object properties can also propagate to the applications 410a-g. The memory objects of the object memory fabric 405 can be used to eliminate typical size constraints on memory space of the commodity servers or other nodes imposed by address sizes. Rather, physical addressing can be managed within the memory objects themselves and the objects can in turn be accessed and managed through the object name space. The memory objects of the object memory fabric 405 can also be used to eliminate the distinction between memory (temporary) and storage (persistent) by implementing and managing both within the objects. The object memory fabric 405 can also eliminate the distinction between local and remote memory by transparently managing the location of objects (or portions of objects) so all objects appear simultaneously local to all nodes. The memory objects can also eliminate the distinction between processing and memory through methods of the objects to place the processing within the memory itself. In other words, embodiments of the present invention provide a single-level memory that puts the computes with the storage and the storage with the computes, directly and thereby eliminating numerous levels of software overhead communicating across these levels and the artificial overhead of moving data to be processed.

[0124] In these ways, embodiments of the object memory fabric 405 and components thereof as described herein can provide transparent and dynamic performance acceleration, especially with big data or other memory intensive applications by reducing or eliminating overhead typically associated with memory management, storage management, networking, data directories, and data buffers at both the system and application software layers. Rather, management of the memory objects at the memory level can significantly shorten the pathways between storage and memory and between memory and processing, thereby eliminating the associated overhead between each.

[0125] Embodiments provide coherent, hardware-based, infinite memory managed as memory objects with performance accelerated in-memory, spanning all nodes, and scalable across all nodes. This enables transparent dynamic performance acceleration based on the object and end application. Using an architecture according to embodiments of the present invention, applications and system software can be treated the same and as simple as a single, standard server but additionally allowing memory fabric objects to capture heuristics. Embodiments provide multiple dimensions of accelerated performance including locality acceleration. According to one embodiment, object memory fabric metadata associated with the memory objects can include triggers which enable the object memory fabric architecture to

localize and move data to fast dram memory ahead of use. Triggers can be a fundamental generalization that enables the memory system to execute arbitrary functions based on memory access. Various embodiments can also include an instruction set which can provide a unique instruction model for the object memory fabric based on the triggers defined in the metadata associated with each memory object and that supports core operations and optimizations and allows the memory intensive portion of applications to be more efficiently executed in a highly parallel manner within IMF.

[0126] Embodiments can also decrease software path-length by substituting a small number of memory references for a complex application, storage and network stack. This can be accomplished when memory and storage is directly addressable as memory under embodiments of the present invention. Embodiments can additionally provide accelerated performance of high level memory operations. For many cases, embodiments of the object memory fabric architecture can eliminate the need to move data to the processor and back to memory, which is extremely inefficient for today's modern processors with three or more levels of caches.

[0127] FIG. 5 is a block diagram illustrating an exemplary memory fabric object memory according to one embodiment of the present invention. More specifically, this example illustrates an application view of how memory fabric object memory can be organized. Memory fabric object address space 500 can be a 128 bit linear address space where the object ID corresponds to the start of the addressable object. Objects 510 can be variable size from 212 to 264 bytes. The address space 500 can efficiently be utilized sparsely within and across objects as object storage is allocated on a per block basis. The size of the object space 500 is meant to be large enough that garbage collection is not necessary and to enable disjoint systems to be easily combined.

[0128] Object metadata 505 associated with each object 510 can be transparent with respect to the object address space 500 and can utilize the object memory fabric to manage objects and blocks within objects and can be accessible at appropriate privilege by applications 515a-g through Application Program Interfaces (APIs) of the object memory fabric. This API provides functions for applications to set up and maintain the object memory fabric, for example by using modified Linux libc. With a small amount of additional effort applications such as a SQL database or graph database can utilize the API to create memory objects and provide and/or augment object metadata to allow the object memory fabric to better manage objects. Object metadata 505 can include object methods, which enable performance optimization through dynamic object-based processing, distribution, and parallelization. Metadata can enable each object to have a definable security policy and access encapsulation within an object.

[0129] According to embodiments of the present invention, applications 515a-g can now access a single object that captures it's working and/or persistent data (such as App0 515a) or multiple objects for finer granularity (such as App1 515b). Applications can also share objects. Object memory 500 according to these embodiments can physically achieves this powerfully simple application view with a combination of physical organization, which will be described in greater detail below with reference to FIG. 6, and object memory dynamics. Generally speaking, the object memory 500 can be organized as a distributed hierarchy that creates hierar-

chical neighborhoods for object storage and applications 515a-g. Object memory dynamics interact and leverage the hierarchal organization to dynamically create locals of objects and applications (object methods) that operate on objects. Since object methods can be associated with memory objects, as objects migrate and replicate on the memory fabric, object methods naturally gain increased parallelism as object size warrants. The hierarchy in conjunction with object dynamics can further create neighborhoods of neighborhoods based on the size and dynamics of the object methods.

[0130] FIG. 6 is a block diagram illustrating an exemplary object memory dynamics and physical organization according to one embodiment of the present invention. As illustrated in this example, an object memory fabric 600 as described above can include any number of processing nodes 605 and 610 communicatively coupled via one or more external object routers 615. Each node 605 and 610 can also include an internal object router 620 and one or more memory modules. Each memory module 625 can include a node object memory 635 supporting any number of applications 515a-g. Generally speaking, the memory module 625, node object router 620 and inter-node object router 615 can all share a common functionality with respect to the object memory 635 and index thereof. In other words, the underlying design objects can be reused in all three providing a common design adaptable to hardware of any of a variety of different form factors and types in addition to those implementations described here by way of example.

[0131] More specifically, a node can comprise a single node object router 620 and one or more memory modules 625 and 630. According to one embodiment, a node 605 can comprise a commodity or "off-the-shelf" server, the memory module 625 can comprise a standard format memory card such as a Dual-Inline Memory Module (DIMM) card, and the node object router 620 can similarly comprise a standard format card such as a Peripheral Component Interconnect express (PCIe) card. The node object router 620 can implement an object index covering the objects/blocks held within the object memory(s) 635 of the memory modules 625 and 630 within the same node 605. Each memory module 625 and 630 can hold the actual objects and blocks within objects, corresponding object meta-data, and object index covering objects currently stored local to that memory module. Each memory module 625 and 630 can independently manage both dram memory (fast and relatively expensive) and flash memory (not as fast, but much less expensive) in a manner that the processor (not shown) of the node 605 thinks that there is the flash amount of fast dram. The memory modules 625 and 630 and the node object router 620 can both manage free storage through a free storage index implemented in the same manner as for other indexes. Memory modules 625 and 630 can be directly accessed over the standard DDR memory bus by processor caches and processor memory reference instructions. In this way, the memory objects of the memory modules 625 and 630 can be accessed using only conventional memory reference instructions and without implicit or explicit Input/Output (I/O) instructions.

[0132] Objects within the object memory 635 of each node 625 can be created and maintained through an object memory fabric API (not shown). The node object router 620 can communicate with the API through a modified object memory fabric version of libc and an object memory fabric

driver (not shown). The node object router **620** can then update a local object index, send commands toward a root, i.e., towards the inter-node object router **615**, as required and communicate with the appropriate memory module **625** or **630** to complete the API command locally. The memory module **625** or **630** can communicate administrative requests back to the node object router **620** which can handle them appropriately.

**[0133]** According to one embodiment, the internal architecture of the node object router **620** can be very similar to the memory module **625** with the differences related to routing functionality such as managing a node memory object index and routing appropriate packets to and from the memory modules **625** and **630** and the inter-node object router **615**. That is, the node object router **620** can have additional routing functionality but does not need to actually store memory objects.

**[0134]** The inter-node object router **615** can be considered analogous to an IP router. However, the first difference is the addressing model used. IP routers utilize a fixed static address per each node and routes based on the destination IP address to a fixed physical node. However, the inter-node object router **615** of the object memory fabric **600** utilizes a memory fabric object address (OA) which specifies the object and specific block of the object. Objects and blocks can dynamically reside at any node. The inter-node object router **615** can route OA packages based on the dynamic location(s) of objects and blocks and track object/block location dynamically in real time. The second difference is that the object router can implement the object memory fabric distributed protocol which provides the dynamic nature of object/block location and object functions, for example including, but not limited to, triggers. The inter-node object router **615** can be implemented as a scaled up version of node object router **620** with increased object index storage capacity, processing rate and overall routing bandwidth. Also, instead of connecting to a single PCIe or other bus or channel to connect to memory modules, inter-node object router **615** can connect to multiple node object routers and/or multiple other inter-node object routers. According to one embodiment, a node object router **620** can communicate with the memory modules **625** and **630** with direct memory access over PCIe and the memory bus (not shown) of the node **605**. Node object routers of different nodes **605** and **610** can in turn connect with one or more inter-node object routers **615** over a high-speed network (not shown) such as 25/100GE fiber that uses several layers of Gigabit Ethernet protocol or object memory fabric protocol tunneled through standard IP, for example. Multiple inter-node object routers can connect with the same network.

**[0135]** In operation, the memory fabric object memory can physically achieve its powerfully simple application view described above with reference to FIGS. **4** and **5** with a combination of physical organization and object memory dynamics. According to one embodiment and as introduced above with reference to FIG. **5**, the memory fabric object memory can be organized as a distributed hierarchy that creates hierarchical neighborhoods for object storage and applications **515a-g**. The node object routers can keep track of which objects and portions of objects are local to a neighborhood. The actual object memory can be located on nodes **605** or **610** close to applications **515a-g** and memory fabric object methods.

**[0136]** Also as introduced above, object memory dynamics can interact and leverage the hierarchical organization to dynamically create locals of objects and applications (object methods) that operate on objects. Since object methods can be associated with objects as objects migrate and replicate across nodes, object methods naturally gain increased parallelism as object size warrants. This object hierarchy, in conjunction with object dynamics, can in turn create neighborhoods of neighborhoods based on the size and dynamics of the object methods.

**[0137]** For example, App0 **515a** spans multiple memory modules **625** and **630** within a single level object memory fabric neighborhood, in this case node **605**. Object movement can stay within that neighborhood and its node object router **620** without requiring any other communication links or routers. The self-organizing nature along the hierarchy defined neighborhoods provides efficiency from a performance and minimum bandwidth perspective. In another example, App1 (A1) **515b** can have the same characteristic but in a different neighborhood, i.e., in node **610**. App2 (A2) **515c** can be a parallel application across a two-level hierarchy neighborhood, i.e., nodes **605** and **610**. Interactions can be self-contained in the respective neighborhood.

**[0138]** As noted above, certain embodiments may include a data types and metadata architecture certain embodiments can also include a data types and metadata architecture that facilitate multiple advantages of the present invention. With respect to the architecture, the following description discloses various aspects of: object memory fabric address spaces; an object memory fabric coherent object address space; an object memory fabric distributed object memory and index; an object memory fabric index; object memory fabric objects; and an extended instruction execution model. Various embodiments may include any one or combination of such aspects.

**[0139]** FIG. **7** is a block diagram illustrating an aspect of object memory fabric hierarchy of object memory, which localizes working sets and allows for virtually unlimited scalability, according to one embodiment of the present invention. As disclosed herein, certain embodiments may include core organization and data types that enable the object memory fabric to dynamically operate to provide the object memory application view. The core organization and data types facilitate the fractal-like characteristics of the system which allow the system to behave identically in a scale-independent fashion. In the depicted example, an object memory fabric **700** as disclosed herein can include any number of processing nodes **705** and **710** communicatively coupled at higher levels via one or more external object routers, such as object router **715**, which may in turn be coupled to one or more higher level object routers.

**[0140]** Specifically, the system may be a fat-tree built from nodes, from leaf nodes to root node(s). According to certain embodiments, each node may just understand whether its scope encompasses an object and based on that whether to route a request/response toward the root or leaf. Putting these nodes together enables a system to dynamically scale to any capacity, without impacting the operation or perspective of any node. In some embodiments, the leaf node may be a DIMM built from standard memory chips, plus object memory fabric **700** implemented within an FPGA. In some embodiments, standard memory chips could have object memory fabric **700** imbedded. In various embodiments,

implementations may have remote nodes such as mobile phones, drones, cars, interne of things components, and/or the like.

[0141] To facilitate various advantageous properties of object memory fabric 700, certain embodiments may employ coherent object memory fabric address spaces. Table 1 below identifies non-limiting examples of various aspects of address spaces, in accordance with certain embodiments of the present disclosure. All nodes that are connected to a single object memory fabric 700, local or distributed, can be considered part of a single system environment according to certain embodiments. As indicated in Table 1, object memory fabric 700 can provide a coherent object address space. In some embodiments, a 128-bit object address space may be provided. However, other embodiments are possible. There are several reasons for a large object address space, including the following. The object address space is to directly uniquely address and manage all memory, storage across all nodes within an object memory fabric system, and provide a unique address for conventional storage outside of an object memory fabric system. The object address space can allow an address to be used once and never garbage collect, which is a major efficiency. The object address space can allow a distinction between allocating address space and allocating storage. In other words, the object address space can be used sparsely as an effective technique for simplicity, performance, and flexibility.

[0142] As further indicated in Table 1, the object memory fabric 700 can directly support per-process virtual address spaces and physical address spaces. With some embodiments, the per-process virtual address spaces and physical address spaces may be compatible with x86-64 architecture. In certain embodiments, the span of a single virtual address space may be within a single instance of Linux OS, and may be usually coincident with a single node. The object memory fabric 700 may enable the same virtual address space to span more than a single node. The physical address space may be the actual physical memory addressing (e.g., within an x86-64 node in some embodiments).

[0143] FIG. 8 is a block diagram illustrating an example relationship 800 between object address space 805, virtual addresses 810, and physical addresses 815, in accordance with certain embodiments of the present disclosure. With object address space 805, a single object can range in size. By way of example without limitation, a single object can range in size from 2 megabytes ( $2^{21}$ ) to 16 petabytes ( $2^{64}$ ). Other ranges are possible. Within the object memory fabric 700, object address space 805 may be allocated on an object granularity basis in some embodiments. In some embodiments, storage may be allocated on a 4 k byte block basis (e.g., blocks 806, 807). Thus, the object address space block 806, 807 in some embodiments may correspond to the 4 k byte page size within x86-64 architecture. When the object address space 805 is created, only the address space and object metadata may exist. When storage is allocated on a per block basis, there can be data stored in the corresponding block of the object. Block storage can be allocated in a sparse or non-sparse manner and pre and/or demand allocated. For example, in some embodiments, software can use an object as a hash function and only allocate physical storage for the valid hashes.

[0144] Referring to the example of FIG. 8, within a node 820, 825, which could be a conventional server in some embodiments, physical pages corresponding to physical addresses 815 may be allocated on a dynamic basis corresponding to the virtual addresses 810. Since object memory fabric 700 actually provides the physical memory within a node 820, 825 by way of the object memory fabric DIMM, when a virtual address segment 811, 812, 813, 814 is allocated, an object address space 805 object which corresponds to the particular segment 811, 812, 813, 814 can also be created. This enables the same or a different virtual address 810 across nodes 820, 825 to address and access the same object. The actual physical address 815 at which a block/page within an object resides within a node 820, 825 can vary over time within or across nodes 820, 825, transparently to application software.

TABLE 1

Address Spaces			
Parameter	Object memory fabric Object Address Space	Virtual Address	Physical Address
Description	Object memory fabric address	Process address handle to object memory fabric address	Cache of object memory fabric
Scope	Global	Per process, can be shared	Per node
Size	$2^{128}$	$2^{64}$ ( $2^{48}$ Haswell)	$2^{46}$ (Haswell)
Object Support	Yes, object memory fabric object index tree and per object index tree	Yes, page tables	Yes, object memory fabric metadata
Object Sizes		$2^{(1 2 21 30 39 48)}$	
Address Space Allocation	Sparse - with or without storage, object units	Sparse - with or without storage, object units	Sparse - page
Storage Allocation	Object or block (page)	Based on object memory fabric	Page
Security (Access)	Through virtual address, operating system, and file system	Operating system	Operating system/object memory fabric

[0145] Certain embodiments of the object memory fabric 700 may provide key advantages: embodiments of object memory fabric 700 may provide integrated addressing, objects with transparent invariant pointers (no swizzling required), and methods to access a large address space across nodes—a with certain embodiments being compatible with x84-64, Linux, and applications. Normally, systems have numerous different addresses (e.g., for memory address with separate address space, sectors, cylinders, physical disks, database systems, file systems, etc.), which requires significant software overhead for converting, buffering, and moving objects and blocks between different layers of addresses. Using integrated addressing to address objects, and blocks within objects, and using the object namespace eliminates layers of software by having single-level addressing invariant across all nodes/systems. With a sufficiently large address space, one address system is not invariant with particular database application and all these systems working together.

[0146] Thus, a node may include a memory module may store and manage one or more memory objects, where physical address of memory and storage is managed with

block diagram illustrating an example relationship 900 between object sizes 905 and object address space pointers 910, in accordance with certain embodiments of the present disclosure. Table 2 below identifies non-limiting examples of aspects of the object address space pointer 910, in accordance with certain embodiments of the present disclosure. As indicated by Table 2, some example embodiments can support three pointer formats. The object address space format may be an object memory fabric native 128 bit format and can provide a single pointer with full addressability for any object and offset within object. Object memory fabric 700 can support additional formats, for example, two additional formats in 64 bit format to enable direct compatibility with x86-64 virtual memory and virtual address. Once a relationship between an object memory fabric object and a virtual address segment is established by object memory fabric API (which can be handled transparently to the application in Linux libc, in some embodiments), standard x86 Linux programs can directly reference data within an object (x86 segment) efficiently and transparently utilizing the x86-64 addressing mechanisms.

TABLE 2

Object Address Space Pointer Formats				
Pointer Type	Object memory fabric Pointer	Object Address Space Generation	Transformation to Virtual Address	Virtual Address Format
Object memory fabric Address	128 bit Storage	Direct	None	None
Object Relative	Offset (64 bit)	Obj Start + ObjOffset	None	virtual address base + offset address mode
Object Virtual Address	Offset (64 bit)	ObjStart + ObjOffset	Add virtual address base to offset	48 bit virtual address with 64 bit data type

each of the one or more memory objects based at least in part on an object address space that is allocated on a per-object basis with a single-level object addressing scheme. The node may be configured to utilize the object addressing scheme to operatively couple to one or more additional nodes to operate as a set of nodes of an object memory fabric, where the set of nodes operates so that all memory objects of the set of nodes are accessible based at least in part on the object addressing scheme, the object addressing scheme defining invariant object addresses for the one or more memory objects that are invariant with respect to physical memory storage locations and storage location changes of the one or more memory objects within the memory module and across all modules interfacing the object memory fabric. Accordingly, the object addresses are invariant within a module and across all modules that interface to object memory fabric, regardless of whether the objects are in a single server or not. Even though the objects can be stored on any or all modules, the object addresses are still invariant no matter at which physical memory locations the objects are currently or will be stored. The following provides details of certain embodiments that may provide such advantages through the object address space and object address space pointers.

[0147] Certain embodiments of object memory fabric 700 may support multiple, various pointer formats. FIG. 9 is a

[0148] Table 3 below identifies non-limiting examples of aspects of the object address space pointers in relation to object sizes, in accordance with certain embodiments of the present disclosure. Embodiments of object address space can supports multiple segment sizes, for example, six segment sizes from 221 to 264 as illustrated in Table 3 below. The object sizes correspond to the x86-64 virtual memory segment and large page sizes. Objects can start on a modulo 0 object size boundary. Object address space pointers 910 may be broken into ObjStart and ObjOffset fields, the sizes of which are dependent on the object size as shown in the example below. The ObjStart field corresponds to the object address space start of the object and also the ObjectID. The ObjOffset is an unsigned value in a range from zero to (ObjectSize-1) with specifies the offset within an object. Object metadata can specify the object size and object memory fabric interpretation of the object address space pointer 910. Objects of arbitrary size and sparseness can be specified by only allocating storage for blocks of interest within an object.

[0149] Because of the nature of most applications and object nature of object memory fabric 700, most addressing can be relative to an object. In some embodiments, all the object memory fabric address pointer formats can be natively stored and loaded by the processor. Object Relative

and Object Virtual Address can work directly with x86-64 addressing modes in some embodiments. Object Virtual Address pointer can be or include a process virtual address that works within the x86-64 segment and corresponding object memory fabric object. Object memory fabric object address space can be calculated by using the Object Virtual Address as an object offset. Object Relative pointer can be or include an offset into an x86-64 virtual address segment, thus base plus index addressing mode works perfectly. Object memory fabric object address space can be calculated by using the Object Relative as an object offset. Table 3 below identifies non-limiting examples of details of generating a 128 bit object address space from an Object Virtual Address or Object Relative pointer as a function of object size, in accordance with certain embodiments of the present disclosure.

TABLE 3

Object Address Space Generation	
Object Size	Object Address Space Generation from Object Relative and Object Virtual Address Pointers
$2^{21}$	$IA[127:00] = (ObjBase[127:21], zero[20:0]) + (zero[127:21], ObjOffset[20:0])$
$2^{30}$	$IA[127:00] = (ObjBase[127:30], zero[29:0]) + (zero[127:30], ObjOffset[29:0])$
$2^{39}$	$IA[127:00] = (ObjBase[127:39], zero[38:0]) + (zero[127:39], ObjOffset[38:0])$
$2^{48}$	$IA[127:00] = (ObjBase[127:48], zero[47:0]) + (zero[127:48], ObjOffset[47:0])$
$2^{57}$	$IA[127:00] = (ObjBase[127:57], zero[56:0]) + (zero[127:57], ObjOffset[56:0])$
$2^{64}$	$IA[127:00] = (ObjBase[127:21], zero[20:0]) + (zero[127:21], ObjOffset[20:0])$

[0150] As disclosed herein, certain embodiments may include an object memory fabric distributed object memory and index. With the distributed index, individual nodes may index local objects and blocks of objects on a per-object basis. Certain embodiments of object memory fabric distributed object memory and index may be based at least in part on an intersection concept of cellular automata and fat trees. Prior distributed hardware and software systems with real-time dynamic indices used two approaches: a centralized index or a distributed single conceptual index. Embodiments of object memory fabric may use a new approach which overlays an independent local index function on top of a fat-tree hierarchical network.

[0151] FIG. 10 is a block diagram illustrating an example object memory fabric distributed object memory and index structure 1000, in accordance with certain embodiments of the present disclosure. At leaves of the structure 1000 are any number of processing nodes 1005 and 1010 object memories 1035. These object memories 1035 may each have an object index that describes the objects and portions of objects currently stored locally in the object memories 1035. A number of object memories 1035, which in some embodiments may be DDR4-DIMM interface compatible cards within a single node are logically connected with an object memory fabric node object index 1040. The object memory fabric node object indices 1040 may each have an object index that describes the objects and portions of objects currently stored locally and/or currently stored in the object memories 1035. In some embodiments, the object memory

fabric node object index 1040 can be instantiated as a PCIe card. With some embodiments, the object memory fabric object memory DDR4-DIMM and object memory fabric node object index PCIe card can communicate over PCIe and memory bus.

[0152] In some embodiments, the object memory fabric node object index 1040 works identically to the object index within the object memory 1035, except that the object memory fabric node object index 1040 tracks all objects and portions of objects that are within any of the connected object memories 1035 and maps the objects and portions of objects to particular object memory 1035. The next level up in the tree is an node object router object index 1020 that may be provided by an object memory fabric router that performs the same object index function for all the object memory fabric node object indices 1040 to which it is connected. The node object router object indices 1020 may each have an object index that describes the objects and portions of objects currently stored locally in lower levels (e.g., at 1040, 1035). Thus, according to some embodiments, router modules may have directory and router functions, whereas memory modules may have directory and router functions, as well as memory functions to store memory objects. However, other embodiments are possible, and, in alternative embodiments, the router modules may additionally have memory functions to store memory objects.

[0153] The pattern may illustrated by the structure 1000 may continue to another higher level inter-node object router object index 1015 that may be provided by an object memory fabric router that performs the same object index function for all the object memory fabric node object indices to which it is connected, and so on to the root of the tree. Thus, in certain embodiments, each object index and each level may perform the same function, independently, but, the aggregate of object indices and levels as a tree network may provide a real time dynamic distributed index, with great scalability properties, that efficiently tracks and localizes memory objects and blocks. An additional property may be that the combination of tree, distributed indices, and caching enable a significant reduction in bandwidth requirements. This may be illustrated by the hierarchically indicated neighborhoods that are delineated by object memory fabric router to leaves (down in this case). As the level of the defined hierarchy increases, so does the aggregate object memory caching capacity. So, as an application working set fits within the aggregate capacity of a given level, the bandwidth requirement at the level toward the root may go to zero.

[0154] As disclosed herein, each processing node is configured to utilize a set of algorithms to operatively couple to one or more additional processing nodes to operate as a set of processing nodes independently of a scale of the set of processing nodes. The set of nodes may operate so that all memory objects of the set of nodes are accessible by any node of the processing set of nodes. At the processing nodes, object memory modules may store and manage memory objects, each instantiated natively therein and managed at a memory layer, and object directories that index the memory objects and blocks thereof on a per-object basis. A memory module may process requests based at least in part on the one or more object directories, which requests may be received from an application layer. In some case, the requests may be received from one or more additional processing nodes. Responsive to the requests, a given memory module may process an object identifier corre-

sponding to a given request and may determine whether the memory module has requested object data. If the memory module has the requested object data, the memory module may generate a response to the request based at least in part on the requested object data. If the memory module does not have the requested object data, an object routing module may route the first request to another node in the tree. The routing of the request may be based at least in part on the object routing module making a determination about a location of object data responsive to the request. If the object routing module identifies the location based at least in part on the object routing module's directory function, the object routing module may route the request down toward the location (i.e., a lower level node possessing the requested object data). However, if the object routing module determines that the location is unknown, the object routing module may route the request toward a root node (i.e., to one or more higher level object routers—inter-node object routers) for further determinations at each level until the requested object is located, accessed, and returned to the original memory module.

**[0155]** In addition, as disclosed herein, triggers may be defined for objects and/or blocks within objects in object metadata. The object-based triggers may predict what operations will be needed and may provide acceleration by performing the operations ahead of time. When a node receives a request that specifies an object (e.g., with a 128-bit object address), the node uses an object directory to determine if the node has any part of the object. If so, the object directory points to a per-object tree (a separate one, where the size is based on the size of the object) which may be used to locate local the blocks of interest. There could be additional trigger metadata that indicates, for the particular blocks of interest, to interpret the particular addresses in a predefined manner as the blocks are transferred to/through the memory module. The triggers may specify one or more pre-defined hardware and/or software actions on a per-block basis with respect to one or more data blocks within an object (e.g., fetch a particular address, run a more complicated trigger program, perform pre-fetching, calculate these other three blocks and send signal to software, etc.). Triggers may correspond to a hardware way to dynamically move data and/or perform other actions ahead of when such actions are needed as objects flow through any memory module of the object memory fabric. Accordingly, such actions may be effected when a particular memory object having one or more trigger is located at a respective memory module and accessed as part of the respective memory module processing one or more other requests.

**[0156]** FIGS. 11 and 12 are block diagrams illustrating examples at a logical level of how the distributed nature of the object index operates and interoperates with the object memory fabric protocol layering, in accordance with certain embodiments of the present disclosure. Certain embodiments of object memory fabric protocol layering may be similar to, but have important differences from, a conventional layered communication protocol. A communications protocol may be essentially stateless, but embodiments of the object memory fabric protocol may maintain object state and directly enable distributed and parallel execution—all without any centralized coordination.

**[0157]** FIG. 11 illustrates an object memory hit case 1100 that executes completely within the object memory 1135, in accordance with certain embodiments of the present disclo-

sure. Object memory 1135 may receive a processor request 1105 or background trigger activity 1106. The object memory 1135 may manage the local DRAM memory as a cache 1130, based on processor physical address. The most frequent case may be that the requested physical address is present and it gets immediately returned to the processor, as indicated at 1110. The object memory 1135 may use triggers to transparently move data from slower flash memory into the fast DRAM memory, as indicated at 1115.

**[0158]** For the case of a miss 1115 or background trigger activity 1106, some embodiments may include one or a combination of the following. In some embodiments, an object memory fabric object address may be generated from the physical address, as indicated by block 1140. The object index may generate the location in local flash memory from the object address space, as indicated by block 1145. Object index lookup can be accelerated by two methods: (1) a hardware-based assist for index lookup; and (2) results of the object index lookup being locally cached. Object memory fabric cache coherency may be used to determine whether the local state is sufficient of the intended operation, as indicated by block 1150. Based on the index, a lookup may be performed to determine whether the object and/or block within object are local, as indicated by block 1155. In the case of a hit 1160, the data corresponding to request 1105 or trigger activity 1106 may be transferred, as indicated by 1165. And, in some embodiments, when the cache state is sufficient, a decision may be made to cache the block into DRAM.

**[0159]** FIG. 12 illustrates an object memory miss case 1200 and the distributed nature of the object memory and object index, in accordance with certain embodiments of the present disclosure. The object memory 1235 may go through steps described previously, but the routing/decision stage 125 may determine that the object and/or block is not local. As a result, the algorithm may involve the request traversing 1270 up the tree toward the root, until the object/block is found. Any number of levels and corresponding node elements may be traversed until the object/block is found. In some embodiments, at each step along the path, the same or similar process steps may be followed to independently determine the next step on the path. No central coordination is required. Additionally, as disclosed herein, object memory fabric API and triggers normally get executed in the leafs, but can be executed in a distributed manner at any index.

**[0160]** As a simplified example, in the case depicted the request traverses 1270 up from the object memory fabric node object index 1240 corresponding to object memory 1235 to the object router 1220. The object router 1220, with its an object router object index, may identify the request object/block as being down the branch toward object memory fabric node object index 1241. Hence, at the index of object router 1220, the request may then be routed 1275 toward the leaf(s) that can supply the object/block. In the example depicted, the object memory 1236 can supply the object/block. At the object memory 1236, memory access/caching 1241 may be performed (which may include previously described process steps for a hit case being performed), and the object/block may be returned 1280 back to the original requesting leaf 1235 for the ultimate return 1290. Again, in some embodiments, at each step along the path, the same or similar process steps may be followed to independently determine the next step on the path. For example, the original requesting leaf 1235 may perform

previously described process steps **1285** for a hit case, and then return **1290** the requested data.

**[0161]** As disclosed herein, the operation of a single object memory fabric index structure, the object memory fabric index structure may be based on several layers of the same tree implementation. Certain embodiments employing tree structure may have several uses within object memory fabric as described in Table 4 below. However, various other embodiments are possible.

TABLE 4

Tree Structure Uses			
Use	Object Memory	Node Object Index	Object Memory Fabric Router
Determine local location of objects and blocks comprising objects as function of object address space	Yes		
Determine which children hold objects, and blocks comprising objects, as a function of object address space		Yes	Yes
Generate object address space as function of local physical address (single level)	Yes		
Object virtual address to object address space		Yes	
Application defined	Yes		

**[0162]** FIG. 13 is a block diagram illustrating an example of leaf level object memory structure **1300** in view of the object memory fabric distributed object memory and index structure, in accordance with certain embodiments of the present disclosure. In some embodiments, the leaf level object memory structure **1300** may include a nested set of B-trees. The root tree may be the object index tree (OIT) **1305**, which may index objects locally present. The index for the object index tree **1305** may be the object memory fabric object address, since objects start at object size modulo zero. There may be one object index tree **1305** for each object that has at least a single block stored locally within the object memory.

**[0163]** The object index tree **1305** may provide one or more pointers (e.g., local pointers) to one or more per object index trees (POIT) **1310**. For example, every local object may have a per object index tree **1310**. A per object index tree **1310** may index object metadata and blocks belonging to the object that are locally present. The per object index tree **1310** leaves point to the corresponding metadata and blocks (e.g., based on offset within object) in DRAM **1315** and flash **1320**. A leaf for a specific block can point to both DRAM **1315** and flash **1320**, as in the case of leaf **1325**, for example. Organization of object metadata and data is disclosed further herein.

**[0164]** The tree structure utilized may be a modified B-tree that is copy-on-write (COW) friendly. COW is an optimization strategy that enables multiple tasks to share information efficiently without duplicating all storage where most of the data is not modified. COW stores modified blocks in a new location which works well for flash memory and caching. In certain embodiments, the tree structure utilized may be similar to that of the open source Linux file system btrfs, with major differences being utilization for a single object/memory space, hardware acceleration, and the ability of independent local indices to aggregate as described previously. By utilizing multiple layers of B-trees, there can be a higher degree of sharing and less rippling of changes. Applications, such as file systems and database storage managers, can utilize this underlying efficient mechanism for higher level operation.

**[0165]** FIG. 14 is a block diagram illustrating an example of object memory fabric router object index structure **1400**, in accordance with certain embodiments of the present disclosure. With some embodiments, the object memory fabric router object index and the node object index may use an almost identical structure of object index trees **1405** and per object index trees **1410** for each object. The object index trees **1405** may index objects locally present. Each object described in an object index tree **1405** may have a per object index tree **1410**. The per object index trees **1410** may index blocks and segments that are locally present.

**[0166]** The object memory fabric router object index and the node object index may index objects and blocks within objects that are present in the children **1415** within the tree structure **1400**, namely child router(s) or leaf object memory. An entry within a leaf in the per object index tree **1410** has the ability to represent multiple blocks within the object. Since blocks of an object may tend to cluster together naturally and due to background housekeeping, each object tends to be represented much more compactly in object indices that are closer to the tree root. The object index trees **1405** and per object index trees **1410** may enable deduplication at the object and block level, since multiple leaves can point to the same blocks, as in the case of leaves **1425** and **1430**, for example. Index Copy-On-Write (COW) support enables, for example, only modified blocks to be updated for an object.

**[0167]** FIGS. 15A and 15B are block diagrams illustrating non-limiting examples of index tree structures, including node index tree structure **1500** and leaf index tree **1550**, in accordance with certain embodiments of the present disclosure. Further non-limiting examples of various aspects of index tree fields are identified in Table 5 below. Other embodiments are possible. An individual index tree may include node blocks and leaf blocks. Each node or leaf block may include of a variable number of entries based on the type and size. Type specifies type of node, node block, leaf, and/or leaf block.

TABLE 5

Index Tree Fields		
Name	Description	Size
NSize	Encoded node size field. Single value for OIT node. Multiple values for POIT node based on object size corresponding to POIT index. Implies the size of NValue field.	3
ObjSize	Encoded Object Size	3



TABLE 5-continued

Index Tree Fields		
Name	Description	Size
ObjectID	Maximum size object ID	107
Object Offset	4k block Based on Object size corresponding to POIT index (9-52)	52
L.Pointer (LP)	References local 4k block in flash or dram. Includes 32 bits of pointer and a single bit specifying dram address space.	32
L.Parent (LPt)	Local Parent references the local 4k block of the parent node in flash or dram. Includes 32 bits of pointer and a single bit specifying dram address space.	33
L.Size	Encoded leaf LValue size.	3
Otype	Type of OIT Leaf	2
Ptype	Type of POIT Leaf	2
Etype	Type of OIT or POIT Entry Node	3
Rtype	Type of reserved Leaf	3
num	May be utilized to increase the size of data that the leaf specifies to increase the efficiency of index tree and storage device. Values may include: 1 block 4 blocks (flash page) 512 blocks (minimum size object, 2 Mbyte)	0
Children	Specifies a remote device number	32
Block State	Encoding of 4k block cache coherency state	8
	Block referenced count (unsigned)	7
	Modified - Indicates that the block has been modified with respect to persistent store. Only valid for blocks while they are present in volatile memory.	1
DS State [15:0]	DownStream State [15:0] - Enumerates the state of for the block within object specified by Object Offset for each of 16 devices.	128

**[0168]** Size specifies independently the size of the LPointer and IndexVal (or object offset). Within a balanced tree, a single block may point to all node blocks or all leaf blocks. In order to deliver highest performance, the tree may become un-balanced, such as for example where the number of levels for all paths through the tree are equivalent. Node blocks and leaf blocks may provide fields to support un-balanced trees. A background activity may re-balance the trees that are part of other background operations. For example, an interior node (non-leaf) in OIT may include L Pointer and NValue fields. NValue may include object size and object ID. Object ID requires 107 (128-21) bits to specify the smallest possible object. Each LPointer may point to the next level of interior node or a leaf node. LPointer may require enough bits to represent all the blocks within its local storage (approximately 32 bits representing 16 terabytes). For a node in the POIT, the NValue may consist of the object offset based on object size. The object size may be encoded within the NSize field. The size field may enable a node to hold the maximum number of LPointer and NValue fields based on usage. An index tree root node may be stored at multiple locations on multiple flash devices to achieve reliable cold boot of the OIT. Tree root block updates may be alternated among mirrors to provide wear leveling.

**[0169]** By default, each POIT Leaf entry may point to the location of a single block (e.g., 4 k bytes). POIT Leaf OM entry and POIT Leaf Router entry may contain a field to enable support beyond single block to enable more compressed index trees, higher resulting index tree performance and higher persistent storage performance by being able to match the page size for persistent storage.

**[0170]** Nodes and leafs may be differentiated by the Type field at the start of each 4 k block. The NNize field may

encode the size of NValue field within a node, and LSize field may encode the size of the LValue field within a leaf. The size of the LPointer field may be determined by the physical addressing of local storage is fixed for a single devices (e.g., RDIMM, node router, or router). The LPointer may be only valid within a single device and not across devices. The LPointer may specify whether the corresponding block is stored in persistent memory (e.g., flash) or faster memory (e.g., DRAM). Blocks that are stored in DRAM may also have storage allocated within persistent memory, so that two entries are present to indicate the two storage locations for a block, node or leaf. Within a single block type, all NValue and/or LValue fields may be a single size.

**[0171]** The OIT Node may include several node level fields (Type, NSize, and LParent) and entries including OIT Node Entry or OIT Leaf Entry. Since an index tree can be un-balanced at times a node can include both node and leaf entries. The POIT Node may include one or more node level fields (e.g., Type, NSize, and/or LParent) and entries including OIT Leaf Entry. OIT Leaf types may be differentiated by the otype field. OIT Leaf (Object Index Table Leaf) may point to the head of a POIT (Per Object Index Table) that specifies object blocks and object metadata. OIT Leaf R may point to a remote head of an POIT. This may be utilized to reference an object that is residing on a remote device across a network. This leaf may enable the remote device to manage the object.

**[0172]** POIT Leaf types may be differentiated by the ptype field. POIT Leaf OM may point to a block of object memory or metadata. The Object offset field may be one bit greater than the number of bits to specify the offset for a specific object size to specify metadata. For example, for 221 object size 10 bits may be required (9 plus 1 bits). The implementation can choose to represent the offset in two's comple-

ment form (signed form, first block metadata is -1), or in one's complement where the additional bit indicates metadata (first block of metadata is represented by 1, with metadata bit set).

**[0173]** POIT Leaf Remote may point to a block of object memory or metadata that is remote from the local DIMM. This may be used to reference a block that is residing on a remote device across a network through the stream package interface. For example, this device could be a mobile device. This leaf may enable object memory fabric hardware to manage coherence on a block basis for the remote device.

**[0174]** POIT Leaf Router may be utilized within node object routers and inter-node object routers to specify the state of the corresponding object memory fabric Block Object Address for each of up to 16 downstream nodes. If within a node object router, up to 16 DIMMs may be specified in some embodiments (or more in other embodiments). If within an inter-node object router up to 16 downstream routers or node object routers (e.g., server nodes) may be specified in some embodiments (or more in other embodiments). The Block Object Address can be present in one or more downstream devices based on valid state combinations.

**[0175]** Index lookups, index COW updates, and index caching may be directly supported in object memory fabric hardware in Object Memory, node object index, and object memory fabric Router. In addition to the node formats for object memory fabric indices, application-defined indices may be supported. These may be initialized through the object memory fabric API. An advantage of application-defined indices may be that object memory fabric hardware-based index lookup, COW update, index caching, and parallelism may be supported.

**[0176]** Various embodiments may provide for background operations and garbage collection. As each DIMM and Router within object memory fabric may maintain its own directory and storage locally, background operations and garbage collection may be accomplished locally and independently. Each DIMM or Router may have a memory hierarchy for storing index trees and data blocks, that may include on-chip cache, fast memory (e.g., DDR4 or HMC DRAM) and slower nonvolatile memory (e.g., flash) that it can manage, as well as index trees.

**[0177]** Each level within the hierarchy may perform the following operations: (1) Tree balancing to optimize lookup time; (2) Reference count and aging to determine when blocks are moved between different storage; (3) Free list updating for each local level of hierarchy as well as keeping a parameters of fill level of the major levels of the local hierarchy; (4) Delivering periodic fill levels to the next level of hierarchy to enable load balancing of storage between DIMMs on a local server and between levels of object memory fabric hierarchy; (5) If a Router, then load balancing between child nodes.

**[0178]** Block reference count may be utilized object memory fabric to indicate the relative frequency of access. Higher value may indicate more frequent use over time, lower less frequent use. When block reference count is associated with a block in persistent memory, blocks which have lowest values may be candidates to move to another DIMM or node that has more available space. Each time a block is accelerated into volatile memory, the reference count may be incremented. Low frequency background scanning may decrement the value if it is not in volatile

memory and increments the value if it is in volatile memory. It may be expected that the scanning algorithm may evolve over time to increment or decrement based on reference value to provide appropriate hysteresis. Blocks that are frequently accelerated into or present in volatile memory may have higher reference count values.

**[0179]** When a block reference count is associated with a block in volatile memory, blocks which have lowest values may be candidates to move back to persistent memory or memory within another DIMM or node. When a block moves into volatile memory, reference count may be initialized based on the instruction or use case that initiated the movement. For example, a demand miss may set the value to a midpoint, and a speculative fetch may set it to a quarter point. Single use may set it to below the quarter point. Moderate frequency background scanning may decrement the referenced value. Thus, demand fetches may be initially weighted higher than speculative fetches. If a speculative fetch is not utilized, it may quickly fall to the lower referenced values that may be replaced first. Single use may be weighted low to be candidate for replacement sooner than other blocks. Thus, single use and speculative blocks may not replace other frequently accessed blocks.

**[0180]** FIG. 16 is a block diagrams illustrating an aspect of example physical memory organization 1600, in accordance with certain embodiments of the present disclosure. Object memory fabric may provide multiple methods to access objects and blocks. For example, a direct method may be based on execution units within object memory fabric or devices that can directly generate full 128-bit memory fabric addresses may have full direct access.

**[0181]** An associated method may consider conventional servers having limited virtual address and physical address spaces. Object memory fabric may provide an API to dynamically associate objects (e.g., segments) and blocks (e.g., pages) with the larger object memory fabric 128-bit memory fabric address. The associations provided by AssocObj and AssocBlk operations may be utilized by object memory fabric driver (e.g., Linux driver) and object memory fabric system library (Syslib) interfacing with the standard processor memory management to enable object memory fabric to behave transparently to both the operating system and applications. Object memory fabric may provide: (a) an API to associate a processor segment and its range of virtual addresses with an object memory fabric object thus ensuring seamless pointer and virtual addressing compatibility; (b) an API to associate a page of virtual address space and the corresponding object memory fabric block with a page/block of local physical memory within an object memory fabric DIMM (which may ensure processor memory management and physical addressing compatibility); and/or (c) local physical memory divided into standard conventional server DIMM slots, with 512 Gbytes (239 bytes) per DIMM slot. On a per slot basis, object memory fabric may keep an additional directory indexed by physical address of the object memory fabric address of each block that has been associated with the corresponding physical address as illustrated in the following diagram.

**[0182]** FIG. 16 is a block diagram illustrating an example physical memory organization 1600, in accordance with certain embodiments of the present disclosure. A physical memory directory 1605 for physical memory 1630 may include: object memory fabric object block address 1610; object size 1615; reference count 1620; a modified field

1625 which may indicate whether the block has been modified with respect to persistent memory; and/or write enable 1630 which may indicate whether local block cache state is sufficient for writing. For example, if the cache state were copy, writes may be blocked, and object memory fabric would may with sufficient state for writing. The physical address range may be assigned to each by system BIOS on boot based object memory fabric DIMM SPD (Serial Presence Detect) configuration.

[0183] FIG. 17A is a block diagram illustrating an example object addressing 1700, in accordance with certain embodiments of the present disclosure. FIG. 17B is a block diagram illustrating example aspects of object memory fabric pointer and block addressing 1750, in accordance with certain embodiments of the present disclosure. Object memory fabric objects 1705 may include object data 1710 and metadata 1715, both divided into 4 k blocks in some embodiments as one unit of storage allocation, referenced by the object memory fabric address space 1720. The object starting address may be the ObjectID 1755. Data 1710 may be accessed as a positive offset from ObjectID 1755. The largest offset may be based on ObjectSize 1760.

[0184] Object metadata 1715 may be accessed as a negative offset from ObjectStart 1725 (ObjectID). Metadata 1715 can be also referenced by an object memory fabric address in the top 1/16th of object address space 1720. The start of a specific objects metadata may be 2128-2124+ObjStart/16. This arrangement may enable the POIT to compactly represent metadata 1715 and the metadata 1715 to have an object address space so it can be managed coherently just like data. Although the full object address space may be allocated for object data 1710 and metadata 1715, storage may be sparsely allocated on a block basis. At a minimum, an object 1705 has a single block of storage allocated for the first block of metadata 1715, in some embodiments. Object access privilege may be determined through object memory fabric Filesystem ACL or the like. Since object memory fabric manages objects in units of 4 k blocks, addressing within the object memory fabric object memory are block addresses, called Block Object Address 1765 (BOA), which corresponds to object address space [127:12]. BOA [11:0] may be utilized by the object memory for ObjectSize (BOA[7:0]) and object metadata indication (BOA[2:0])

[0185] FIG. 18 is a block diagram illustrating example aspects 1800 of object metadata 1805, in accordance with certain embodiments of the present disclosure. Table 6 below indicates metadata of the first block 1810 of metadata 1805 per certain embodiments. In some embodiments, the first block 1810 of metadata 1805 may hold metadata for an object as depicted.

TABLE 6

Metadata First Block		
Name	Description	Size
Object address space	Object ID. Number of significant bits determined by object size	16
Object size	Object Size	
CRC	Reserved for optional object crc	16
Parity pointer	Pointer to pages used for optional object block parity	16
Compression Flags	OID of compression object	16
Encryption Flags	OID of encryption object	16

TABLE 6-continued

Metadata First Block		
Name	Description	Size
System Defined Application	Reserved for software defined OS functions	256
Defined Others	Reserved for software defined owning application functions	256
Remote Object Table	Specifies Objects accessible from this object. Specifies 64 OIDs (128 bit). The zero entry is used to specify object or metadata within this	1024
Triggers	Triggers or Trigger B-Tree root	2048
		4096

[0186] System-defined metadata may include any Linux-related data to coordinate use of certain objects seamlessly across servers. Application-defined metadata may include application related data from a file system or database storage manager to enable searches and/or relationships between objects that are managed by the application.

[0187] For an object with a small number of triggers, base triggers may be stored within the first block; otherwise, a trigger B-tree root may reference metadata expansion area for the corresponding object. Trigger B-tree leaf may specify base triggers. A base trigger may be a single trigger action. When greater than a single action is required, a trigger program may be invoked. When trigger programs are invoked, they may reside in the expansion area. The remote object table may specify objects that are accessible from this object by the extended instruction set.

[0188] Certain embodiments may provide for an extended instruction execution model. One goal of the extended execution model may be to provide a lightweight dynamic mechanism to provide memory and execution parallelism. The dynamic mechanism enables a dataflow method of execution that enables a high degree of parallelism combined with tolerance of variation in access delay of portion of objects. Work may be accomplished based on the actual dependencies, not a single access delay holding up the computation.

[0189] Various embodiments may include one or a combination of the following. Loads and memory references may be split transactions, with separate request and response so that the thread and memory path are not utilized during the entire transaction. Each thread and execution unit may be able to issue multiple loads into object memory fabric (local and remote) prior to receiving a response. Object memory fabric may be a pipeline to handle multiple requests and responses from multiple sources so that memory resources can be fully utilized. The execution unit may be able to accept responses in a different order from that the requests were issued. Execution units can switch to different threads to be fully utilized. Object memory fabric can implement policies to dynamically determine when to move objects or portions of objects versus moving a thread versus creating a thread.

[0190] FIG. 19 is a block diagram illustrating aspects of an example micro-thread model 1900, in accordance with certain embodiments of the present disclosure. A thread may be the basic unit of execution. A thread may be defined at least in part by an instruction pointer (IP) and a frame pointer (FP). The instruction pointer may specify the current instruction that is being executed. The frame pointer may specify the location of the current execution state of the thread.

[0191] A thread can include multiple micro-threads. In the example depicted, the thread **1905** include micro-threads **1906** and **1907**. However, a thread can include greater numbers of micro-threads. The micro-threads of a particular thread may share the same frame pointer but have different instruction pointers. In the example depicted, frame pointers **1905-1** and **1905-2** specify the same location, but instruction pointers **1910** and **1911** specify different instructions.

[0192] One purpose of micro-threads may be to enable data-flow like operation within a thread by enabling multiple asynchronous pending memory operations. Micro-threads may be created by a version of the fork instruction and may be rejoined by the join instruction. The extended instruction set may treat the frame pointer as a top of stack or register set by performing operations on offsets from the frame pointer. Load and store instructions may move data between the frame and the object.

[0193] FIG. 20 is a block diagram illustrating aspects of an example relationship **2000** of code, frame, and object, in accordance with certain embodiments of the present disclosure. Specifically, FIG. 20 illustrates how object data **2005** is referenced through the frame **2010**. The default may be for load and store instructions to reference the object **2005** within local scope. Access to object **2005** beyond local scope can be given in a secure manner by access control and security policies. Once this access is given, objects **2005** within local and non-local scope can be accessed with equal efficiency. Object memory fabric encourages strong security by encouraging efficient object encapsulation. By sharing the frame, micro-threads provide a very lightweight mechanism to achieve dynamic and data-flow memory and execution parallelism, for example, on the order of **10-20** micro-threads or more. The multiple threads enable virtually unlimited memory based parallelism.

[0194] FIG. 21 is a block diagram illustrating aspects of an example of micro-thread concurrency **2100**, in accordance with certain embodiments of the present disclosure. Specifically, FIG. 21 illustrates the parallel data-flow concurrency for a simple example of summing several randomly located values. A serial version **2105** and a parallel version **2110** are juxtaposed, in accordance with certain embodiments of the present disclosure. The parallel version **2110** can be almost  $n$  times faster since loads are overlapped in parallel.

[0195] Referring again to FIG. 20, the approach can be extended to interactive and recursive approaches in a dynamic manner. The advantages of prefetching ahead can now be achieved in cases with minimal locality without using prefetch. When an object is created, a single default thread **2015** (single micro-thread **2020** is created) may be waiting to start with a start message to the default thread **2015**. The default thread **2015** then can create micro-threads with the thread or use a version of the fork instruction to create a new thread.

[0196] In some embodiments, both the instruction pointer and the frame pointer may be restricted to the expansion metadata region **1815** starting at block two and extending to  $\text{SegSize}/16$ . As the number of objects, object size, and object capacity increase, the thread and micro-thread parallelism may increase. Since threads and micro-threads may be tied to objects, as objects move and distribute so may the threads and micro-threads. Embodiments of object memory fabric may have the dynamic choice of moving objects or portions of objects to threads or distributing threads to the object(s).

This may be facilitated by the encapsulated object methods implemented by the extended execution model.

[0197] As further noted above, embodiments of the present invention may also include an object memory fabric instruction set which can provide a unique instruction model based on triggers that support core operations and optimizations and allow the memory intensive portion of applications to be more efficiently executed in a highly parallel manner within the object memory fabric.

[0198] The object memory fabric instruction set can be data-enabling due to several characteristics. First, the sequence of instructions can be triggered flexibly by data access by a conventional processor, object memory fabric activity, another sequence or an explicit object memory fabric API call. Second, sequences can be of arbitrary length, but short sequences can be more efficient. Third, the object memory fabric instruction set can have a highly multi-threaded memory scale. Fourth, the object memory fabric instruction set can provide efficient co-threading with conventional processors.

[0199] Embodiments of the present invention include two categories of instructions. The first category of instructions is trigger instructions. Trigger instructions include a single instruction and action based on a reference to a specific Object Address (OA). A trigger instruction can invoke extended instructions. The second category of instructions is extended instructions. Extended instructions define arbitrary parallel functionality ranging from API calls to complete high level software functions. After a discussion of the instruction set model, these two categories of instructions will be discussed in turn. As noted, trigger instructions enable efficient single purpose memory related functions with no context outside of the trigger.

[0200] Using the metadata and triggers defined above an execution model based on memory data flow can be implemented. This model can represent a dynamic dataflow method of execution in which processes are performed based on actual dependencies of the memory objects. This provides a high degree of memory and execution parallelism which in turn provides tolerance of variations in access delays between memory objects. In this model, sequences of instructions are executed and managed based on data access. These sequences can be of arbitrary length but short sequences are more efficient and provide greater parallelism.

[0201] The extended instruction set enables efficient, highly threaded, in-memory execution. The instruction set gains it's efficiency in several manners. First, the instruction set can include direct object address manipulation and generation without the overhead of complex address translation and software layers to manage differing address spaces. Second, the instruction set can include direct object authentication with no runtime overhead that can be set based on secure third party authentication software. Third, the instruction set can include object related memory computing. For example, as objects move, the computing can move with them. Fourth, the instruction set can include parallelism that is dynamic and transparent based on scale and activity. Fifth, the instruction set can include an object memory fabric operation that can be implemented with the integrated memory instruction set so that memory behavior can be tailored to application requirements. Sixth, the instruction set can handle functionality for memory-intensive computing directory in the memory. This includes adding operations as memory is touched. Possible opera-

tions may include, but are not limited to, searching, image/signal processing, encryption, and compression. Inefficient interactions with conventional processors are significantly reduced.

[0202] The extended instruction capability can be targeted at memory intensive computing which is dominated with memory references for interesting size problems that are larger than caches or main memory, and simple operations based on these references. Some examples can include but are not limited to:

[0203] Defining API macros from conventional processors.

[0204] Defining the streams of interaction between hierarchical components of the object memory fabric. Each component can use a core set of instruction sequences to implement object memory fabric functionality.

[0205] Short sequences for macros to accelerate key application kernels such as BFS (Breath First Search), etc. BFS is a core strategy for searching a graph and is heavily used by graph databases and graph applications. For example, BFS is used across a wide variety of problem spaces to find a shortest or optimal path. It is a representative algorithm that illustrates the challenges for analyzing large scale graphs namely, no locality because graphs are larger than caches and main memory and virtually all the work is through memory references. In the case of BFS, the extended instruction capability described herein coupled with threads handles almost the entire BFS by recursive instantiation of threads to search adjacency lists based on graph size and available nodes. Highly parallel direct in-memory processing and high-level memory operations reduce software path-length. When combined with object memory fabric capability described above to bring all data in-memory and localize it ahead of use, the performance and efficiency per node is significantly increased.

[0206] Complete layer functionality, such as:

[0207] Storage engine for hierarchical file system built on top of a flat object memory. A storage engine is, for example, what stores, handles, and retrieves the appropriate object(s) and information from within an object. For MySQL, the object may be a table. For a file system, the object may be a file or directory. For a graph database, the object may be a graph and information may consist of vertices and edges. Operators supported may be, for example, based on type of object (file, graph, SQL, etc.).

[0208] Storage engine for structured database such as MySQL

[0209] Storage engine for unstructured data such as graph database

[0210] Storage engine for NoSQL key-value store

[0211] Complete application: Filesystem, structured database such as MySQL, unstructured data such as graph database or NoSQL key-value store

[0212] User programmable.

[0213] According to one embodiment, a base trigger may invoke a single trigger action based on reference to a specific OA. There can be a single base trigger per OA. When greater than a single action is required, a trigger program can be invoked with the TrigFunction base trigger. Base triggers may consist of the instructions included in Table 7 below.

TABLE 7

Example Base Trigger Instruction Set	
Base Trigger	Description
Trigger	Fetch the block specified in the pointer at the specified object offset based on specified trigger conditions and actions
TrigFunction	Execute the trigger program starting at specified meta-data offset when the specified data object offset and specified trigger conditions.

[0214] As noted, the Trigger instruction set can include fetching the block specified in the pointer at the specified object offset based on the specified trigger conditions and actions. The Trigger instruction binary format can be expressed as:

[0215] Trigger PtrType TrigType TrigAction RefPolicy ObjOffset

[0216] An example set of operands for the Trigger instruction set are included in Tables 8-12 below.

TABLE 8

PtrType—Pointer Type		
Encoding	Symbol	Description
	None	No pointer
	OA	Object Address
	ObjReg	Object Relative
	ObjVA	Object Virtual Address
	Reserved	Reserved

TABLE 9

TrigType—Trigger Type		
Encoding	Symbol	Description
	None	
	demand	Trigger by demand miss for block
	prefetch	Trigger by preched block
	access	Triggered by actual processor access to cache block
	emptyfill	Trigger by empty or fill instructions. Enables trigger on specific processor action
	any	Any trigger type
	reserved	Reserved

TABLE 10

TrigAction—Trigger Action		
Encoding	Symbol	Description
	None	
	Cache	Trigger by demand miss for block
	Clean	Trigger by preched block
	reserved	Triggered by actual processor access to cache block

TABLE 11

RefPolicy—Reference Count and Policy		
Encoding	Symbol	Description
	InitLowA	Initial reference count of prefetch page to low value, policy A

TABLE 11-continued

RefPolicy—Reference Count and Policy		
Encoding	Symbol	Description
	InitMidA	Initial reference count of prefetch page to mid value, policy A
	InitHighA	Initial reference count of prefetch page to high value, policy A
	InitLowB	Initial reference count of prefetch page to low value, policy B
	InitMidB	Initial reference count of prefetch page to mid value, policy B
	InitHighB	Initial reference count of prefetch page to high value, policy B

TABLE 12

ObjOffset—Object Offset Description	
Object offset based on Object size. Trigger can be evaluated based on TriggerType and trigger action taken if TriggerType is satisfied is define by TriggerAction and RefPolicy.	

[0217] As noted, the TrigFunction (or TriggerFunc) instruction set can include executing the trigger program starting at specified meta-data offset when the specified data object offset and specified trigger conditions. TriggerFunc can enable more complex sequences than a single Trigger instruction to be executed. The TrigFunc Instruction binary format can be expressed as:

[0218] TrigFunc PtrType TrigType MetaDataOffset ObjOffset

[0219] An example set of operands for the Trigger instruction set are included in Tables 13-16 below.

TABLE 13

PtrType—Pointer Type		
Encoding	Symbol	Description
	None	No pointer
	OA	Object Address
	ObjReg	Object Relative
	ObjVA	Object Virtual Address
	Reserved	Reserved

TABLE 14

TrigType—Trigger Type		
Encoding	Symbol	Description
	None	
	demand	Trigger by demand miss for block
	prefetch	Trigger by preched block
	access	Triggered by actual processor access to cache block
	emptyfill	Trigger by empty or fill instructions. Enables trigger on specific processor action

TABLE 14-continued

TrigType—Trigger Type		
Encoding	Symbol	Description
	any	Any trigger type
	reserved	Reserved

TABLE 15

MetaDataOffset—Meta-Data Offset Description	
Meta-Data offset based on Object size. TriggerFunction can be evaluated based on TriggerType. The trigger program starting at MetaDataOffset is executed if TriggerType is satisfied.	

TABLE 16

ObjOffset—Object Offset Description	
Object offset based on Object size. TriggerFunction can be evaluated based on TriggerType at ObjOffset. The trigger program starting at MetaDataOffset is executed if TriggerType is satisfied.	

[0220] According to one embodiment, extended instructions can be interpreted in 64 bit word chunks in 3 formats, including short (2 instructions per word), long (single instruction per word), and reserved.

TABLE 17

Extended Instruction Format			
Format	bits[63:62]	bits[61:31]	bits[30:0]
Short	0x00	s_instruction[1] (31 bits)	s_instruction[0] (31 bits)
Long	0x01	l_instruction (62 bits)	
Reserved	0x1*		

[0221] Generally speaking, triggers in combination with the extended instruction set can be used to define arbitrary, parallel functionality such as: direct object address manipulation and generation without the overhead of complex address translation and software layers to manage differing address space; direct object authentication with no runtime overhead that can be set based on secure 3rd party authentication software; object related memory computing in which, as objects move between nodes, the computing can move with them; and parallelism that is dynamically and transparent based on scale and activity. These instructions are divided into three conceptual classes: memory reference including load, store, and special memory fabric instructions; control flow including fork, join, and branches; and execute including arithmetic and comparison instructions.

[0222] A list of the different types of memory reference instructions are shown in Table 18 below.

TABLE 18

Memory Reference Instructions					
Instruction	[30:23] Encoding/Options	[22:17] FPA	[16:11] FPB	[10:5] FPC	[4:0] Predicate
Pull	encode[7:0]	oid	offset	prior, plstate	src_pred
Push	encode[7:0]	oid	offset	prior, plstate	src_pred
Ack	encode[7:0]	oid	offset		src_pred
Load	encode[4:0], osize[2:0]	src oid	src offset	dst fp	src_pred
Store	encode[4:0], osize[2:0]	dst oid	dst offset	src fp	src_pred
ReadPA	encode[7:0]		src pa	dst fp	src_pred
WritePA	encode[7:0]		dst pa	src fp	src_pred
Empty	encode[7:0]	src oid	src offset	dst fp	src_pred
Fill	encode[7:0]	dst oid	dst offset	src fp	src_pred
Pointer	encode[5:0], opt[1:0]	dst oid	dst offset		src_pred
PrePtrChn	encode[4:0], opt[2:0]	src oid	src offset st	src offset end	src_pred
ScanEF	encode[4:0], opt[2:0]	src oid	src offset	dst fp	src_pred
Create					src_pred
CopyObj					src_pred
CopyBlk					src_pred
Allocate					src_pred
Deallocate					src_pred
Destroy					src_pred
Persist					src_pred
AssocObj					src_pred
DeAssocObj					src_pred
AssocBlk	encode[5:0], opt[1:0]	src oid	src pa	dst ls	src_pred
DeAssocBlk	encode[7:0]				src_pred
OpenObj					src_pred
OpenBlk					src_pred
Btree					src_pred

[0223] The pull instruction may be utilized within the object memory fabric as a request to copy or move the specified block to (e.g. local) storage. The 4 k byte block operand in the object specified by src\_oid at the object offset specified by src\_offset may be requested with the state specified by pull\_state with the priority specified by priority. The data may be subsequently moved by a push instruction. The Pull instruction binary format can be expressed as:

Pull Instruction (binary format)					
[30:23]	[22:17]	[16:11]	[10:9]	[8:5]	[4:0]
src_old	src_offset	priority	pull_state	Predicate	

[0224] An example set of operands for the Pull instruction set are included in Tables 19-23 below.

TABLE 19

predicate—Predicate Description
Specifies a single bit predicate register. If the predicate value is true, the instruction executes, if false the instruction does not execute.

TABLE 20

src_oid—Source Object Identifier Description
Index into the remote object table to specify the specific object identifier for this memory operation. Index value of 0 always corresponds to local object.

TABLE 21

src_off—Source Object Offset Description
Specifies the unsigned offset from the thread frame pointer to read the source operand corresponding to the object offset.

TABLE 22

priority—How object memory fabric treats the requests		
Encoding	Symbol	Description
0x0	required-high	Highest priority handling of requests. Highest priority requests are always handled in the order received.
0x1	required-low	Can be optionally reordered with respect to required-high by object memory fabric only to prioritize required-high requests for short time periods. Must be completed. Typically most requests are of required-low priority.

TABLE 22-continued

priority—How object memory fabric treats the requests		
Encoding	Symbol	Description
0x2	optional-high	Requests can be considered optional by object memory fabric and can be delayed or deleted as required to manage object memory fabric load. Optional-high requests are always considered ahead of optional-low requests.
0x3	optional-low	Request can be considered optional by object memory fabric and can be delayed or deleted as required to manage object memory fabric load. Optional-low requests are treated at the lowest priority. Typically most optional requests are o the optional-low priority.

TABLE 23

pull_state—Requested object memory fabric state for block		
States can be listed in order of weakest to strongest. State can be returned in a stronger state. Modified with respect to persistent memory can be indicated by _m suffix.		
Encoding	Symbol	Description
0x0	invalid	
0x1	snapcopy	Snapshot copy. This copy can be updated when a block is persisted. Utilized for object fault tolerance. Can be configured on an object basis redundancy and geographic dispersion.
0x2	shadcopy	Shadow copy. Can be updated on a lazy basis (eventually consistent), usually after a period of time or some number of writes and/or transactions. Can also be used for fault tolerant block copies.
0x3	copy	Read-only copy. Will be updated for owner modifications as they occur. Insures sequential consistency.
0x4	own_snapcopy	Exclusive owner with snapshot copy. Enables local write privilege without any updates required. Snapshot copies may exist, but are only updated when corresponding block is persisted and through and push instruction with push_state = pstate_sncopy.
0x8	own-snapcopy_m	
0x5	own_shadcopy	Non-exclusive owner with shadow copies. Enables write privilege shadow copies or snapshot copies to exist which are updated from writes on a lazy basis-eventually consistent.
0x9	own-shadcopy_m	
0x6	own_copy	Non-exclusive owner with copies. Enables write privilege and copies, shadow copies or snapshot copies to exist which are updated from writes. Multiple writes to the same block can occur with a single update.
0xa	own_copy_m	
0x7	own	Exclusive owner. Enables local write privilege. No copies, shadow copies or snapshot copies exist.
0xb	own_m	
0xc	error	Error has been encountered on corresponding block.
0xd-0xf	reserved	Reserved

[0225] Push instruction may be utilized to copy or move the specified block from local storage to a remote location. The 4 k byte block operand in the object specified by src\_oid at the object offset specified by src\_offset may be requested with the state specified by pull\_state with the priority specified by priority. The data may be previously requested by a pull instruction. The Push instruction binary format can be expressed as:

Push Instruction (binary format)					
[30:23]	[22:17]	[16:11]	[10:9]	[8:5]	[4:0]
src_oid	src_offset	priority	push_state	Predicate	

[0226] An example set of operands for the Push instruction set are included in Tables 24-28 below.

TABLE 24

predicate—Predicate
Description
Specifies a single bit predicate register. If the predicate value is true, the instruction executes, if false the instruction does not execute.

TABLE 25

src_oid—Source Object Identifier
Description
Index into the remote object table to specify the specific object identifier for this memory operation. Index value of 0 always corresponds to local object.



TABLE 26

src_off—Source Object Offset	Description
	Specifies the unsigned offset from the thread frame pointer to read the source operand corresponding to the object offset.

TABLE 27

priority—How object memory fabric treats the requests		
Encoding	Symbol	Description
0x0	required-high	Highest priority handling of requests. Highest priority requests are always handled in the order received.
0x1	required-low	Can be optionally reordered with respect to required-high by object memory fabric only to prioritize required-high requests for short time periods. Must be completed. Typically most requests are of required-low priority.
0x2	optional-high	Requests can be considered optional by object memory fabric and can be delayed or deleted as required to manage object memory fabric load. Optional-high requests are always considered ahead of optional-low requests.
0x3	optional-low	Request can be considered optional by object memory fabric and can be delayed or deleted as required to manage object memory fabric load. Optional-low requests are treated at the lowest priority. Typically most optional requests are o the optional-low priority.

TABLE 28

push_state—Requested object memory fabric state for block Modified with respect to persistent memory can be indicated by <code>_m</code> suffix.		
Encoding	Symbol	Description
0x0	invalid	
0x1	snapcopy	Snapshot copy. This copy can be updated when a block is persisted. Utilized for object fault tolerance. Can be configured on an object basis redundancy and geographic dispersion.
0x2	shadcopy	Shadow copy. Will be updated on a lazy basis-eventually consistent, usually after a period of time or some number of writes and/or transaction. Can also be used for fault tolerant block copies.
0x3	copy	Read-only copy. Can be updated for owner modifications as they occur. Insures sequential consistency.
0x4	own_snapcopy	Exclusive owner with snapshot copy. Enables local write privilege without any updates required. Snapshot copies may exist, but are only updated when corresponding block is persisted and through and push instruction with <code>push_state = pstate_sncopy</code> .
0x8	own_snapcopy_m	
0x5	own_shadcopy	Non-exclusive owner with shadow copies. Enables write privilege shadow copies or snapshot copies to exist which are updated from writes on a lazy basis-eventually consistent.
0x9	own_shadcopy_m	
0x6	own_copy	Non-exclusive owner with copies. Enables write privilege and copies, shadow copies or snapshot copies to exist which are updated from writes. Multiple writes to the same block can occur with a single update.
0xa	own_copy_m	
0x7	own	Exclusive owner. Enables local write privilege. No copies, shadow copies or snapshot copies exist.
0xb	own_m	
0xc	error	Error has been encountered on corresponding block.
0xd-0xf	reserved	

[0227] PushAck or Ack instruction may be utilized to acknowledge that the block associated with a Push has been accepted at one or more locations. The 4 k byte block operand in the object specified by src\_oid at the object offset specified by src\_offset may be acknowledged. The Ack instruction binary format can be expressed as:

Ack Instruction (binary format)					
[30:23]	[22:17]	[16:11]	[10:9]	[8:5]	[4:0]
src_oid	src_offset				Predicate

[0228] An example set of operands for the Push instruction set are included in Tables 29-31 below.

TABLE 29

predicate—Predicate Description
Specifies a single bit predicate register. If the predicate value is true, the instruction executes, if false the instruction does not execute.

TABLE 30

src_oid—Source Object Identifier Description
Index into the remote object table to specify the specific object identifier for this memory operation. Index value of 0 always corresponds to local object.

TABLE 31

src_off—Source Object Offset Description
Specifies the unsigned offset from the thread frame pointer to read the source operand corresponding to the object offset.

[0229] The Load instruction set includes the osize operand in the object specified by src\_oid at the object offset specified by src\_offset. src\_offset can be written to the word offset from the frame pointer specified by dst\_fp. The load instruction ignores the empty state.

Load Instruction (binary format)					
[30:26]	[25:23]	[22:17]	[16:11]	[10:5]	[4:0]
osize	src_oid	src_offset	dst_fp		Predicate

[0230] An example set of operands for the Load instruction set are included in Tables 32-36 below.

TABLE 32

osize—Object operand size		
Encoding	Symbol	Description
0x0	8 bit unsigned	8 bit source is zero extended to 64 bit dst_fp

TABLE 32-continued

osize—Object operand size		
Encoding	Symbol	Description
0x1	16 bit unsigned	16 bit source is zero extended to 64 bit dst_fp
0x2	32 bit unsigned	32 bit source is zero extended to 64 bit dst_fp
0x3	64 bit	64 bit source is loaded into 64 bit dst_fp
0x4	8 bit signed	8 bit source is sign extended to 64 bit dst_fp
0x5	16 bit signed	16 bit source is sign extended to 64 bit dst_fp
0x6	32 bit signed	32 bit source is sign extended to 64 bit dst_fp
0x7	reserved	

TABLE 33

predicate—Predicate Description
Specifies a single bit predicate register. If the predicate value is true, the instruction executes, if false the instruction does not execute.

TABLE 34

src_oid—Source Object Identifier Description
Index into the remote object table to specify the specific object identifier for this memory operation. Index value of 0 always corresponds to local object.

TABLE 35

src_off—Source Object Offset Description
Specifies the unsigned offset from the thread frame pointer to read the source operand corresponding to the object offset.

TABLE 36

dst_fp—Destination offset from frame pointer Description
Specifies the unsigned offset from the thread frame pointer to write the source operand.

[0231] The Store instruction set includes the word specified by src\_fp can be truncated to the size specified by osize and stored into the object specified by dst\_oid at offset of dst\_offset. For example, only the ssize bytes are stored. The store instruction ignores the empty state. The Store instruction binary format can be expressed as:

Store Instruction (binary format)					
[30:25]	[24:23]	[22:17]	[16:11]	[10:5]	[4:0]
ssize	dst_oid	dst_offset	src_fp		Predicate

[0232] An example set of operands for the Store instruction set are included in Tables 37-41 below.

TABLE 37

ssize—Store Object operand size		
Encoding	Symbol	Description
0x0	8 bit	Least significant 8 bits are stored
0x1	16 bit	Least significant 16 bits are stored
0x2	32 bit	Least significant 32 bits are stored
0x3	64 bit	Least significant 64 bits are stored

TABLE 38

predicate—Predicate Description
Specifies a single bit predicate register. If the predicate value is true, the instruction executes, if false the instruction does not execute.

TABLE 39

dst_oid—Source Object Identifier Description
Index into the remote object table to specify the specific object identifier for this memory operation. Index value of 0 always corresponds to local object.

TABLE 40

dst_off—Source Object Offset Description
Specifies the unsigned offset from the thread frame pointer to read the source operand corresponding to the object offset.

TABLE 41

src_fp—Destination offset from frame pointer Description
Specifies the unsigned offset from the thread frame pointer to read the source operand.

**[0233]** The ReadPA instruction reads 64 bytes by physical address of the local memory module. The operand in the object specified by src\_pa can be written to the word offset from the frame pointer specified by dst\_fp. The ReadPA instruction binary format can be expressed as:

ReadPA Instruction (binary format)					
[30:26]	[25:23]	[22:17]	[16:11]	[10:5]	[4:0]
			src_pa	dst_fp	Predicate

**[0234]** An example set of operands for the ReadPA instruction set are included in Tables 42-44 below.

TABLE 42

predicate—Predicate Description
Specifies a single bit predicate register. If the predicate value is true, the instruction executes, if false the instruction does not execute.

TABLE 43

src_pa—Source Physical Address Description
Specifies a physical address local to the current node/server.

TABLE 44

dst_fp—Destination offset from frame pointer Description
Specifies the unsigned offset from the thread frame pointer to write the source operand

**[0235]** The WritePA instruction writes 64 bytes by physical address of the local memory module. The 64 bytes specified by src\_fp is stored into the physical address specified by dst\_pa. The ReadPA instruction binary format can be expressed as:

WritePA Instruction (binary format)					
[30:25]	[24:23]	[22:17]	[16:11]	[10:5]	[4:0]
			dst_pa	src_fp	Predicate

**[0236]** An example set of operands for the WritePA instruction set are included in Tables 45-47 below.

TABLE 45

predicate—Predicate Description
Specifies a single bit predicate register. If the predicate value is true, the instruction executes, if false the instruction does not execute.

TABLE 46

dst_pa—Destination physical address Description
Specifies a physical address local to the current node/server

TABLE 47

src_fp—Source frame pointer Description
Specifies the unsigned offset from the thread frame pointer to read the source operand.

**[0237]** Each word within an object memory fabric object can include an state to indicate empty or full states. An empty state conceptually means that the value of the corre-

sponding word has been emptied. A full state conceptually means the value of the corresponding word has been filled. This state can be used by certain instructions to indivisibly insure that only a single thread can read or write the word. Empty instructions can operate similar to a load, as shown below in Table 48.

TABLE 48

State	Result
Empty	Memory doesn't respond until word transitions to full state
Full	Completes as load and indivisibly transitions state to empty

**[0238]** The *osize* operand in the object specified by *src\_oid* at the object offset specified by *src\_offset* can be written to the word offset from the frame pointer specified by *dst\_fp*. The Empty instruction binary format can be expressed as:

Empty Instruction (binary format)					
[30:26]	[25:23]	[22:17]	[16:11]	[10:5]	[4:0]
<i>src_oid</i>	<i>src_offset</i>	<i>dst_fp</i>	Predicate		

**[0239]** An example set of operands for the Empty instruction set are included in Tables 49-52 below.

TABLE 49

predicate—Predicate Description
Specifies a single bit predicate register. If the predicate value is true, the instruction executes, if false the instruction does not execute.

TABLE 50

<i>src_oid</i> —Source Object Identifier Description
Index into the remote object table to specify the specific object identifier for this memory operation. Index value of 0 always corresponds to local object.

TABLE 51

<i>src_off</i> —Source Object Offset Description
Specifies the unsigned offset from the thread frame pointer to read the source operand corresponding to the object offset.

TABLE 52

<i>dst_fp</i> —Destination offset from frame pointer Description
Specifies the unsigned offset from the thread frame pointer to write the source operand.

**[0240]** Each word within a memory fabric object can include an state to indicate empty or full states. Empty state

conceptually means that the value of the corresponding word has been emptied. Full state conceptually means the value of the corresponding word has been filled. This state can be used by certain instructions to indivisibly insure that only a single thread can read or write the word. The Fill instruction binary format can be expressed as:

Fill Instruction (binary format)

[30:25]	[24:23]	[22:17]	[16:11]	[10:5]	[4:0]
<i>dst_oid</i>	<i>dst_offset</i>	<i>src_fp</i>	Predicate		

**[0241]** Fill instruction operates similar to a store, as shown below in Table 53.

TABLE 53

State	Result
Empty	The fill instruction completes as a store and transitions state to full.
Full	The fill instruction

**[0242]** The word specified by *src\_fp* can be stored into the object specified by *dst\_oid* at offset of *dst\_offset*. Only the *ssize* bytes are stored. Store ignores empty state. An example set of operands for the Fill instruction set are included in Tables 54-57 below.

TABLE 54

predicate—Predicate Description
Specifies a single bit predicate register. If the predicate value is true, the instruction executes, if false the instruction does not execute.

TABLE 55

<i>dst_oid</i> —Source Object Identifier Description
Index into the remote object table to specify the specific object identifier for this memory operation. Index value of 0 always corresponds to local object.

TABLE 56

<i>dst_off</i> —Source Object Offset Description
Specifies the unsigned offset from the thread frame pointer to read the source operand corresponding to the object offset.

TABLE 57

<i>src_fp</i> —Destination offset from frame pointer Description
Specifies the unsigned offset from the thread frame pointer to read the source operand.

**[0243]** The Pointer instruction set can specify to the object memory fabric that a pointer of *ptr\_type* can be located in

the object specified by scrod at object offset specified by src\_offset. This information can be utilized by the object memory fabric to pre-stage data movement. The Pointer instruction binary format can be expressed as:

Pointer Instruction (binary format)					
[30:26]	[24:23]	[22:17]	[16:11]	[10:5]	[4:0]
ptr_type	src_oid	src_offset			Predicate

[0244] An example set of operands for the Pointer instruction set are included in Tables 58-61 below.

TABLE 58

ptr_type—Pointer Type		
Encoding	Symbol	Description
0x0	none	No pointer at this object offset
0x1	MF Address	Full 128 Memory Fabric Address pointer at this object offset
0x2	Object Relative	64 bit object relative pointer at this object offset. The range of the object relative pointer can be determined by object size
0x3	Object-VA	64 bit object virtual address pointer at this object offset. The range of the object relative pointer can be determined by object size.

TABLE 59

predicate—Predicate Description	
	Specifies a single bit predicate register. If the predicate value is true, the instruction executes, if false the instruction does not execute.

TABLE 60

src_oid—Source Object Identifier Description	
	Index into the remote object table to specify the specific object identifier for this memory operation. Index value of 0 always corresponds to local object.

TABLE 61

src_off- Source Object Offset Description	
	Specifies the unsigned offset from the thread frame pointer to read the source operand corresponding to the object offset.

[0245] The Prefetch Pointer Chain instruction set can be based on the policy specified by policy in the object specified by src\_oid, in the range specified by src\_offset\_st to src\_offset\_end. The osize operand in the object specified by src\_oid at the object offset specified by src\_offset can be written to the word offset from the frame pointer specified by dst\_fp. Load ignores empty state. The PrePtrChn instruction binary format can be expressed as:

PrePtrChn Instruction (binary format)					
[30:26]	[25:23]	[22:17]	[16:11]	[10:5]	[4:0]
policy	src_oid	src_off- set_st	src_off- set_end		src_pred

[0246] An example set of operands for the Prefetch Pointer Chain instruction set are included in Tables 62-66 below.

TABLE 62

Policy- Prefetch PointerChain Policy		
Encoding	Symbol	Description
0x0	none_ahead	Just prefetch blocks corresponding to pointers in chain
0x1	breath_1ahead	Breath first prefetch. Fetch each pointer in chain then fetch one ahead of each pointer
0x2	breath_2ahead	Breath first prefetch. Fetch each pointer in chain then recursively fetch two ahead of each pointer
0x3	breath_3ahead	Breath first prefetch. Fetch each pointer in chain then recursively fetch three ahead of each pointer
0x4	reserved	reserved
0x5	depth_1ahead	Depth first prefetch 1 deep.
0x6	depth_2ahead	Depth first prefetch 2 deep.
0x7	depth_3ahead	Depth first prefetch 3 deep.

TABLE 63

predicate- Predicate Description	
	Specifies a single bit predicate register. If the predicate value is true, the instruction executes, if false the instruction does not execute.

TABLE 64

src_oid- Source Object Identifier Description	
	Index into the remote object table to specify the specific object identifier for this memory operation. Index value of 0 always corresponds to local object.

TABLE 65

src_off_st- Source Object Offset Description	
	Specifies the unsigned offset from the thread frame pointer to read the source operand corresponding to starting object offset..

TABLE 66

src_off_end- Destination offset from frame pointer Description	
	Specifies the unsigned offset from the thread frame pointer to read the source operand corresponding to ending object offset.

[0247] The Scan and Set Empty or Full instruction set can be initialed in an object specified by src\_oid, at offset

specified by `src_offset` with specified policy. Scan can be used to do a breath first or depth first search and empty or fill the next available location. The ScanEF instruction binary format can be expressed as:

ScanEF Instruction (binary format)					
[30:26]	[25:23]	[22:17]	[16:11]	[10:5]	[4:0]
policy	src_oid	src_offset	dst_fp	Predicate	

**[0248]** An example set of operands for the Scan and Set Empty or Full instruction set are included in Tables 67-71 below.

TABLE 67

osize- Object operand size		
Encoding	Symbol	Description
0x0	scan_empty	Scan object until empty state and set to full. Terminates on full with null value. The object offset when the condition was met can be placed into <code>dst_fp</code> . If the scan terminated without condition being met, a value of <code>-0x1</code> can be placed into <code>dst_fp</code> .
0x1	scan_full	Scan object to full state and set to empty. Terminates on empty with null value. The object offset when the condition was met can be placed into <code>dst_fp</code> . If the scan terminated without condition being met, a value of <code>-0x1</code> can be placed into <code>dst_fp</code> .
0x2	ptr_full	Follow pointer chain until full and set to empty. Terminates on null pointer. The object offset when the condition was met can be placed into <code>dst_fp</code> . If the scan terminated without condition being met, a value of <code>-0x1</code> can be placed into <code>dst_fp</code> .
0x3	ptr_empty	Follow pointer chain until empty and set to full. Terminates on null pointer. The object offset when the condition was met can be placed into <code>dst_fp</code> . If the scan terminated without condition being met, a value of <code>-0x1</code> can be placed into <code>dst_fp</code> .

TABLE 68

predicate- Predicate Description
Specifies a single bit predicate register. If the predicate value is true, the instruction executes, if false the instruction does not execute.

TABLE 69

src_oid- Source Object Identifier Description
Index into the remote object table to specify the specific object identifier for this memory operation. Index value of 0 always corresponds to local object.

TABLE 70

src_off- Source Object Offset Description
Specifies the unsigned offset from the thread frame pointer to read the source operand corresponding to the object offset.

TABLE 71

dst_fp- Destination offset from frame pointer Description
Specifies the object offset when the condition was met. If the scan terminated without condition being met, a value of <code>-0x1</code> can be placed into <code>dst_fp</code> .

**[0249]** The Create instruction set includes an object memory fabric object of the specified `ObjSize` with an object ID of OA and initialization parameters of `DataInit` and `Type`. No data block storage can be allocated and storage for the first meta-data block can be allocated. The Create instruction binary format can be expressed as:

**[0250]** Create Type Redundancy `ObjSize` `OID`

**[0251]** An example set of operands for the Create instruction set are included in Tables 72-75 below.

TABLE 72

Type		
Encoding	Symbol	Description
	volatile	temp object that does not need to be persisted
	persistant	object must be persisted
	reserved	reserved

TABLE 73

Redundancy		
Encoding	Symbol	Description
	nonredundant	Object memory fabric does not provide object redundancy
	redundant	Object memory fabric guarantees that object can be persisted in at least 2 separate nodes
	remote_redundant	Object memory fabric guarantees that object can be persisted in at least 2 separate nodes which are remote with respect to each other
	reserved	reserved

TABLE 74

ObjSize- Object Size Description
Specifies the object size.

TABLE 75

OID- Object Id Description
Object memory fabric object ID which also the starting address for the object.

**[0252]** The CopyObj instruction set includes copies source object specified by SOID to destination object specified by DOID. If DOID is larger object than SOID, all DOID blocks beyond SOID size are copied as unallocated. If SOID is larger object than DOID, then the copy ends at DOID size. The CopyObj instruction binary format can be expressed as:

**[0253]** CopyObj Ctype SOID DOID

**[0254]** An example set of operands for the CopyObj instruction set are included in Tables 76-78 below.

76. Ctype- Copy type

Encoding	Symbol	Description
	copy	One time copy from SOID to DOID. Allocated blocks are one time copied and non-allocated block SOID blocks become unallocated DOID blocks, object memory fabric has the option of treating the copy initially as cow and executing the copy in the background.
	cow	All allocated blocks are treated as copy on write. Newly allocated blocks after cow are considered modified.
	reserved	reserved

TABLE 77

SOID- Source Object ID Description
Object memory fabric object ID which is the source for the copy.

TABLE 78

DOID- Destination Object ID Description
Object memory fabric object ID which is the destination for the copy.

**[0255]** The CopyBlk instruction set includes copies cnum source blocks starting at SourceObjectAddress (SOA) to destination starting at DestinationObjectAddress (DOA). If cnum blocks extends beyond the object size associated with SOA, then the undefined blocks are copied as unallocated. The CopyBlk instruction binary format can be expressed as:

**[0256]** CopyBlk ctype cnum SOA DOA

**[0257]** An example set of operands for the CopBlk instruction set are included in Tables 79-82 below.

TABLE 79

Ctype- Copy type		
Encoding	Symbol	Description
	copy	One time copy of cnum blocks starting at SOA to destination blocks starting at DOA. Allocated blocks are one time copied and non-allocated SOA

TABLE 79-continued

Ctype- Copy type		
Encoding	Symbol	Description
		blocks become unallocated SOA blocks, object memory fabric has the option of treating the copy initially as cow and executing the copy in the background.
	cow	All allocated blocks are treated as copy on write. Newly allocated blocks after cow are considered modified.
	reserved	reserved

TABLE 80

cnum- Number of blocks to copy Description
Specifies the number of blocks to copy.

TABLE 81

SOA- Source object memory fabric Block Object Address Description
Object memory fabric block object address which is the source for the copy.

TABLE 82

DOA- Destination object memory fabric Block Object Address Description
Object memory fabric block object address which is the destination for the copy.

**[0258]** The Allocate instruction set includes storage to the object specified by OID. The Allocate instruction binary format can be expressed as:

**[0259]** Allocate init ASize OID

**[0260]** An example set of operands for the Allocate instruction set are included in Tables 83-85 below.

TABLE 83

init- Initialization		
Encoding	Symbol	Description
	zero	Zero all data
	random	Random data.
	reserved	reserved

TABLE 84

ASize- Allocation Size		
Encoding	Symbol	Description
	block	single block
	object	full object
	size21	2 <sup>9</sup> blocks
	size30	2 <sup>18</sup> blocks
	size39	2 <sup>27</sup> blocks

TABLE 85

OID- Object ID Description
Object memory fabric object ID for which storage is allocated.

[0261] The Deallocate instruction set includes storage for cnum blocks starting at OA. If deallocation reaches the end of the object, the operation ends. The Deallocate instruction binary format can be expressed as:

[0262] Deallocate cnum OA

[0263] An example set of operands for the Deallocate instruction set are included in Tables 86 and 87 below.

TABLE 86

cnum- Number of blocks to copy Description
Specifies the number of blocks to deallocate.

TABLE 87

OA- Object Address Description
Object memory fabric block object address which is starting block number for deallocation.

[0264] The Destroy instruction set includes completely deleting all data and meta-data corresponding to object specified by OID. The Destroy instruction binary format can be expressed as:

[0265] Destroy OID

[0266] An example set of operands for the Destroy instruction set are included in Table 88 below.

TABLE 88

OID - Object ID Description
Object ID of the object to be deleted.

[0267] The Persist instruction set includes persisting any modified blocks for the specified OID. The Persist instruction binary format can be expressed as:

[0268] Persist OID

[0269] An example set of operands for the Persist instruction set are included in Table 89 below.

TABLE 89

OID—Object ID Description
Object ID of the object to be persisted.

[0270] The AssocObj instruction set includes associating the object OID with the VaSegment and ProcessID. Associating an OID and VaSegment enables ObjectRelative and ObjectVA addresses to be properly accessed by the object

memory fabric. The AssocObj instruction binary format can be expressed as:

[0271] AssocObj OID ProcessID VaSegment

[0272] An example set of operands for the AssocObj instruction set are included in Tables 90-92 below.

TABLE 90

OID—Object ID Description
Object ID of the object to be associated.

TABLE 91

ProcessID—Process ID Description
Process ID associated with the VaSegment.

TABLE 92

OID—Object ID Description
Object ID of the object to be associated.

[0273] The DeAssocObj instruction set includes de-associating the object OID with the VaSegment and ProcessID. An error can be returned if the ProcessID and VaSegment do not match those previously associated with the OID. The DeAssocObj instruction binary format can be expressed as:

[0274] DeAssocObj OID ProcessID VaSegment

[0275] An example set of operands for the DeAssocObj instruction set are included in Tables 93-95 below.

TABLE 93

OID—Object ID Description
Object ID of the object to be de-associated.

TABLE 94

ProcessID—Process ID Description
Process ID associated with the VaSegment.

TABLE 95

OID—Object ID Description
Object ID of the object to be de-associated.

[0276] The AssocBlk instruction set includes associating the block OA with the local physical address PA. This enables an Object Memory to associate an object memory



fabric block with a PA block for local processor access. The AssocBlk instruction binary format can be expressed as:

[0277] AssocBlk place OA PA LS[15:00]

[0278] An example set of operands for the AssocBlk instruction set are included in Tables 96-99 below.

TABLE 96

place—Physical Placement		
Encoding	Symbol	Description
0x0	match	Associate PA must match physical DIMM with allocated block. If currently not allocated on any physical DIMM will associate and allocate on DIMM specified. Returns status within ack_detail package file of SUCCESS or NOT_ALLOC If not allocated the LS field provides a bitmap of current physical
0x1	force	Force associate and implicit allocate on DIMM specified.
0x2	dynamic	Memory fabric associates a free PA with the OA and returns PA.
0x3	reserved	reserved

TABLE 97

OA—object memory fabric Block Object Address	Description
	Object ID of the object to be associated.

TABLE 98

PA—Physical block Address	Description
	Local physical block address of the block to be associated.

TABLE 99

LS[15:00]—Local State[15:00]	Description
	Valid for ackdetail::NOT_ASSOC which indicates that the OA is allocated on a different physical DIMM. Local state specifies a single bit indicating which DIMM(s) have currently allocated the corresponding OA. Value is return in operand3, with bit0 corresponding to DIMM0.

[0279] The DeAssocBlk instruction set includes de-associating the block OA with the local physical address PA. This OA will then no longer be accessible from a local PA. The DeAssocBlk instruction binary format can be expressed as:

[0280] DeAssocBlk OA PA

[0281] An example set of operands for the DeAssocBlk instruction set are included in Tables 100 and 101 below.

TABLE 100

OA—object memory fabric Block Object Address	Description
	Block object address of block to be de-associated.

TABLE 101

PA—Physical block Address	Description
	Local physical block address of the block to be de-associated. Corresponds to Operand2 within the package header.

[0282] The OpenObj instruction set includes caching the object specified by OID in the manner specified by TypeFetch and CacheMode on an advisory basis. The OpenObj instruction binary format can be expressed as:

[0283] OpenObj TypeFetch CacheMode OID

[0284] An example set of operands for the OpenObj instruction set are included in Tables 102-104 below.

TABLE 102

OID—Object ID	Description
	Object ID of the object to be associated.

TABLE 103

TypeFetch—Type of Prefetch		
Encoding	Symbol	Description
	MetaData	Cache MetaData only
	First 8 Blocks	Cache MetaData and first 8 data blocks
	First 32 Blocks	Cache MetaData and first 32 data blocks
	Reserved	Reserved

TABLE 104

CacheMode- Advisory Block State		
Encoding	Symbol	Description
	copy	Copy block state if possible. All updates can be propagated immediately
	shadcopy	Shadow copy block state if possible. Updates can be propagated in a lazy manner
	snapcopy	Snapshot copy. Copy only updated on persist.
	own	Own block state is possible. No other copies in memory fabric
	owncopy	Own block state with 0 or more copies if possible.
	own_shadcopy	Own block state with 0 or more shadow copies (no copy block state)
	own_snapcopy	Own block state with 0 or more snapshot copes. (no copy or shadow copy block state)

[0285] The OpenBlk instruction set includes caching the block(s) specified by OID in the manner specified by TypeFetch and CacheMode. The prefetch terminates when it's beyond the end of the object. The OpenBlk instruction binary format can be expressed as:

[0286] OpenBlk TypeFetch CacheMode OID

[0287] An example set of operands for the OpenBlk instruction set are included in Tables 105-107 below.

TABLE 105

OID - Object ID Description
Object ID of the object to be associated.

TABLE 106

Encoding	Symbol	TypeFetch- Type of Prefetch Description
	1 Block	Cache MetaData only
	First 8 Blocks	Cache MetaData and 8 data blocks starting at OID
	First 32 Blocks	Cache MetaData and 32 data blocks starting at OID
	Reserved	Reserved

TABLE 107

Encoding	Symbol	CacheMode- Advisory Block State Description
	copy	Copy block state if possible. All updates can be propagated immediately
	shadcopy	Shadow copy block state if possible. Updates can be propagated in a lazy manner
	snapcopy	Snapshot copy. Copy only updated on persist.
	own	Own block state is possible. No other copies in memory fabric
	owncopy	Own block state with 0 or more copies if possible.
	own_shadcopy	Own block state with 0 or more shadow copies (no copy block state)
	own_snapcopy	Own block state with 0 or more snapshot copes. (no copy or shadow copy block state)

[0288] An example set of operands for the Control Flow (short instruction format) instruction set are included in Table 108 below.

TABLE 108

Instruction	Encoding/ Options	[30:23]	[22:17]	[16:11]	[10:5]	[4:0]	Predicate
Fork	encode[6:0], IP fpobj[0]	IP	FP	count	count	src_pred	
Join	encode[6:0], IP fpobj[0]	IP	FP	count	count	src_pred	
Branch BranchLink					disp[5:0]	src_pred src_pred	

[0289] The fork instruction set provides an instruction mechanism to create a new thread or micro-thread. Fork specifies the New Instruction Pointer (NIP) and new Frame Pointer for the newly created thread. At the conclusion of the fork instruction, the thread (or micro-thread) which executed the instruction and the new thread (e.g. micro-thread) are running with fork\_count (count) incremented by one. If the new FP has no relationship to the old FP, it may be considered a new thread, or otherwise a new micro-thread. The Fork instruction binary format can be expressed as:

Fork Instruction (binary format)

[30:24]	[23]	[22:17]	[16:11]	[10:5]	[4:0]
where		NIP	NFP	count	Predicate

[0290] An example set of operands for the Fork instruction set are included in Tables 109-113 below.

TABLE 109

Encoding	Symbol	Description
0x0	frame	Fork count can be stored directly on the frame. Faster, but only accessible to micro-threads within the same thread on a single node
0x1	object	Fork count can be stored within the object which enables distributed operation.

TABLE 110

Encoding	Symbol	predicate- Predicate Description
		Specifies a single bit predicate register. If the predicate value is true, the instruction executes, if false the instruction does not execute.

TABLE 111

Encoding	Symbol	NIP- New micro-thread Instruction Pointer Description
		Specifies the unsigned offset from the thread frame pointer to read the IP of the newly spawned micro-thread. The IP can be a valid object meta-data expansion space address.

TABLE 112

Encoding	Symbol	New micro-thread Frame Pointer Description
		Specifies the unsigned offset from the thread frame pointer to read the FP of the newly spawned micro-thread. The FP can be a valid object meta-data expansion space address.

TABLE 113

Encoding	Symbol	count- Fork count variable Description
		The fork_count variable keeps track of the number of forks that have not been paired with joins. If the where options indicates frame, the count specifies the unsigned offset from the thread frame pointer where fork_count can be located. If the where option indicates object, the count specifies the unsigned offset from the thread frame pointer to read the pointer to fork_count.

[0291] Join is the instruction mechanism to create a new thread or micro-thread. The join instruction set enables a micro-thread to be retired. The join instruction decrements fork count (count) and fork\_count is greater than zero there is no further action. If fork\_count is zero, then this indicates the micro-thread executing the join is the last spawned

micro-thread for this fork\_count and execution continues at the next sequential instruction with the FP specified by FP. The Join instruction binary format can be expressed as:

[30:24]	[23]	[22:17]	[16:11]	[10:5]	[4:0]
where		FP	count	Predicate	

[0292] An example set of operands for the Join instruction set are included in Tables 114-117 below.

TABLE 114

where-Where fork join count can be stored		
Encoding	Symbol	Description
0x0	frame	Fork count can be stored directly on the frame. Faster, but only accessible to micro-threads within the same thread on a single node
0x1	object	Fork count can be stored within the object which enables distributed operation.

TABLE 115

predicate- Predicate Description	
Specifies a single bit predicate register. If the predicate value is true, the instruction executes, if false the instruction does not execute.	

TABLE 116

NFP- Post join Frame Pointer Description	
Specifies the unsigned offset from the thread frame pointer to read the FP of the post join micro-thread. The FP can be a valid object meta-data expansion space address.	

TABLE 117

count—Fork count variable Description	
The fork_count variable keeps track of the number of forks that have not been paired with joins. If the where options indicates frame, the count specifies the unsigned offset from the thread frame pointer where fork_count can be located. If the where option indicates object, the count specifies the unsigned offset from the thread frame pointer to read the pointer to fork_count.	

[0293] The branch instruction set allows for branch and other conventional instructions to be added. The Branch instruction binary format can be expressed as:

Branch Instruction (binary format)					
[30:24]	[23]	[22:17]	[16:11]	[10:5]	[4:0]
					Predicate

[0294] An example set of operands for the Execute (short instruction format) instruction set are included in Table 118 below

TABLE 118

Short Instruction Format-Execute					
	[30:23]	[22:17]	[16:11]	[10:5]	[4:0]
Instruction Add	Encoding/Options encode[5:0], esize[1:0]	FPA srcA	FPB srcB	FPC dst	Predicate src_pred
Compare	encode[5:0], esize[1:0]	srcA	srcB	dpred	src_pred

[0295] Object Memory Fabric Streams and APIs

[0296] Object memory fabric streams facilitate a mechanism that object memory fabric utilizes to implement a distributed coherent object memory with distributed object methods. According to certain embodiments, object memory fabric streams may define a general mechanism that enables hardware and software modules in any combination to communicate in a single direction. Ring streams may support a pipelined ring organization, where a ring of two modules may be just two one-way streams.

[0297] A stream format API may be defined at least in part as two one-way streams. Thus, as part of providing the infinite memory fabric architecture in some embodiments, communication between two or more modules may be executed with the stream format API, which at least partially defines the communication according to the object memory fabric stream protocol so that the communication is based on different unidirectional streams.

[0298] Each stream may be logically composed of instruction packages. Each instruction package may contain an extended instruction and associated data. In some embodiments, each stream may interleave sequences of requests and responses. Streams may include short and long packages. The short package may be referenced herein as simply an “instruction package,” which may be descriptive of the instruction packages containing bookkeeping information and commands. The short package may include either the Pull or Ack instructions and object information. The long package may be referenced herein as an “object data package,” which may be descriptive of the object data packages carrying object data, as distinguished from the short package (“instruction packages”) which do not carry object data. The object data package may include one or more push instructions, object information, and a single block specified by the object address space block address. All other instructions and data may be communicated within the block.

[0299] In some embodiments, for example, the short package may be 64 bytes (1 chunk), and the long package may be 4160 bytes (65 chunks). However, other embodiments are possible. In some embodiments, there may be a separator (e.g., a 1 byte separator). Object memory fabric streams may be connectionless in a manner similar to UDP and may be efficiently embedded over UDP or a UDP-type protocol having certain characteristics common with, or similar to, UDP. In various embodiments, attributes may include any one or combination of:

[0300] Transaction-oriented request-response to enable efficient movement of object memory fabric-named (e.g., 128-bit object memory fabric object address) data blocks.

**[0301]** Packages may be routed based on the location of block, the request object memory fabric object address (object address space), and object memory fabric instruction—not be based on a static IP-like node address.

**[0302]** Coherency and object memory fabric protocol may be implemented directly.

**[0303]** Reliability may be provided within the object memory fabric end-to-end protocol.

**[0304]** Connectionless.

**[0305]** The only state in the system may be the individual block coherency state at each end node, which may be summarized at object memory fabric routing nodes for efficiency.

**[0306]** Table 119 below identifies non-limiting examples of various aspects of a short package definition, in accordance with certain embodiments of the present disclosure.

TABLE 119

Short Package Definition		
Name	Description	Size (bytes)
Instruction	For the short extended instruction format, only s_instruction[0] may be utilized. Pull and Ack may be short extended instructions.	8
ObjID, ObjOff, ObjSize	ObjSize (bit[7:0]) may define the ObjID and ObjOff fields as defined in object memory fabric Coherent Object Address (Object Address Space) Space disclosure above. Bit [11] set specifies meta-data.	16
NodeID	Hierarchical node number. Nodes can be hardware and/or software based. May utilize to route a response back to the original requestor.	8
Acknowledge	Accumulated acknowledge fields. These may be utilized to signal acknowledgement across objects as defined below.	1
Operand2	Utilized for PA address for PA instructions. Utilized for optional streaming block count for other instructions	8
Operand3		8
Checksum	Checksum of the package. This assures correctness all package chunks and correct number of chunks per package.	8
Acknowledge Detail	This may include status or error codes specific to each instruction, shown in the Table CIII below.	1
Local use	Source of the incoming package	1
Local destination	Destination of the outgoing package	1
Local mod_ref	Utilized to locally pass modified and referenced information	1
Reserved	Reserved.	2
Total Size	Short package size.	64

**[0307]** Table 120 below identifies non-limiting examples of various aspects of a long package definition, in accordance with certain embodiments of the present disclosure.

TABLE 120

Long Package Definition		
Name	Description	Size (bytes)
Short package	Push may be long package instruction.	64
Block Data		4096
Total Size	Short package size.	128

**[0308]** Table 121 below identifies non-limiting examples of various aspects of object size encoding, in accordance with certain embodiments of the present disclosure.

TABLE 121

Object Size Encoding	
Encoding	ObjSize
0x0	2 <sup>21</sup>
0x1	2 <sup>30</sup>
0x2	2 <sup>39</sup>
0x3	2 <sup>48</sup>
0x4	2 <sup>57</sup>
0x5	2 <sup>64</sup>
0x6-0xff	reserved

**[0309]** Software and/or hardware based objects may interface to 2 one-way streams, one in each direction. Depending on the object, there can be additional lower level protocol layering including encryption, checksum, and reliable link

protocol. The object memory fabric stream protocol provides for matching request response package pairs (and timeouts) to enforce reliability for packages that traverse over an arbitrary number of streams.

**[0310]** In certain cases, each request-response package pair is approximately 50% short package and 50% long package on the average, the average efficiency relative to a block transfer is 204%, with the equation:

$$\text{efficiency} = 1 / (50\% * 4096 / (40 + 4136))$$

$$= 1 / (50\% * \text{blocksize} / (\text{smallpackage size} + \text{largepackage size}))$$

**[0311]** For links with stochastic error rates, a reliable link protocol may be utilized to detect the errors locally.

[0312] Node ID

[0313] Object address spaces (object memory fabric object addresses) can be dynamically present in any object memory within object memory fabric, as well as dynamically migrate. There still can be (or, for example, needs to be) a mechanism that enables object memory's and routers (collectively nodes) to communicate with each other for several purposes including book-keeping the original requestor, setup and maintenance. The NodeID field within packages can be utilized for these purposes. DIMMs and routers can be addressed based on their hierarchical organization. Non-leaf nodes can be addressed when the lesser significant fields are zero. The DIMM/software/mobile field can enable up to 256 DIMMs or more and the remainder proxied software threads and/or mobile devices. This addressing scheme can support up to 2<sup>40</sup> servers or server equivalents, up to 2<sup>48</sup> DIMMs and up to 2<sup>64</sup> mobile devices or software threads. Examples of these fields are shown below in Tables 122-124.

TABLE 122

Package NodeID Field					
[63:56]	[55:48]	[47:40]	[39:32]	[31:24]	[23:00]
Level5	Level4	Level3	Level2	Level1	DIMM/software/ Inter-Node mobile Object Router

TABLE 123

Leveln Field (n = 1 to 5)	
Encoding	Description
0x00-0xfd	Node address with hierarchy
0xfe	Add this router ID to NodeID when it first leaves this level toward root.
0xff	Indicates that the NodeID Field specifies an interior node that is one level above the field in which this value is specified. All fields lower than this Leveln Field should be specified as 0xff.

TABLE 124

DIMM/SW/Mobile NodeID Field	
Encoding	Description
0x000000-0x0000fe	Up to 256 DIMMs per logical server
0xfe	Add this router ID to NodeID when it first leaves this level toward root.
0x0000ff	Indicates that the NodeID Field specifies an interior node that is one level above the field in which this value is specified. All fields lower than this Leveln Field should be specified as 0xff.
0x000100-0xfffff	Up 2 <sup>24</sup> -2 <sup>8</sup> (16,776,960) SW threads or Mobile Devices per logical server

[0314] Table 125 and 126 below identifies non-limiting examples of various aspects of acknowledge fields and detail, in accordance with certain embodiments of the present disclosure.

TABLE 125

Acknowledge Fields			
Name	Description	Position	Size (bits)
Ack	Cleared when package first inserted into ring from another ring. Set by an object when it is able to respond to the request.	0	1
BusyAck	Cleared when package first inserted into ring from another ring. Set by an object when it is unable to evaluate or perform appropriate action on the Object Block Address. BusyAck may cause the package to be re-transmitted around the local ring.	1	1
SnapCopyAck	Cleared when package first inserted into ring from another ring. Set by an object to indicate it still has a snapshot copy of the Object Block Address. This information may be used to enable the proper state to be set when an object transfers the package between rings (hierarchy levels).	2	1
ShadCopyAck	Cleared when package first inserted into ring from another ring. Set by an object to indicate it still has a shadow copy of the Object Block Address. This information may be used to enable the proper state to be set when an object transfers the package between rings (hierarchy levels).	3	1
CopyAck	Cleared when package first inserted into ring from another ring. Set by an object to indicate it still has a copy of the Object Block Address. This information is used to enable the proper state to be set when an object transfers the package between rings (hierarchy levels).	4	1
ToRoot	Function may be to enable the uplink ring object that provides streams toward the root to not require a directory. Signals the uplink object that a package has traversed once around the ring and can now be sent toward the root.	4	1
Reserved	Reserved.		3
Total Size	Acknowledge field size.		8

[0315] Table 126 below identifies non-limiting examples of various aspects of the Acknowledge detail field, in accordance with certain embodiments of the present disclosure. The Acknowledge detail field may provide detailed status information of the corresponding request based on the package instruction field.

TABLE 126

Acknowledge Detail	
Instructions	Acknowledge Field Definition
Pull, Push, Ack	previous block state. Utilized for diagnostic and
Load, Store	0x0- Success 0x1- Fail
Empty, Fill	0x0- Success 0x1- Fail
Pointer, PrePtrChn, ScanEF	0x0- Success 0x1- Fail
Create	0x0- Success 0x1- Already created (fail) 0x2- Fail
Destroy	0x0- Success 0x1- Not valid (nothing to destroy) 0x2- Fail

TABLE 126-continued

Acknowledge Detail	
Instructions	Acknowledge Field Definition
Allocate	0x0- Success
	0x1- Already allocated (fail)
	0x2- Fail
Deallocate	0x0- Success
	0x1- not allocated (fail)
	0x2- Fail
CopyObj	0x0- Success
	0x1- Object doesn't exist (fail)
	0x2- Fail
CopyBlk	0x0- Success
	0x1- Block doesn't exist (fail)
	0x2- Fail
Persist	0x0- Success
	0x1- Object doesn't exist (fail)
	0x2- Fail
AssocObj	0x0- Success
	0x1- Object doesn't exist (fail)
	0x2- Fail
DeAssocObj	0x0- Success
	0x1- Object doesn't exist (fail)
	0x2- Object not associated (fail)
	0x3- Fail
AssocBlk	0x0- Success
	0x1- Object or block doesn't exist (fail)
	0x2- Fail
DeAssocBlk	0x0- Success
	0x1- Object or block doesn't exist (fail)
	0x2- Block not associated (fail)
	0x3- Fail
OpenObj	0x0- Success
	0x1- Object doesn't exist (fail)
	0x2- Object already open
	0x3- Fail
Btree	0x0- Success
	0x1- Object doesn't exist (fail)
	0x2- Fail

[0316] In some embodiments, the topology used within object memory fabric may be a unidirectional point-to-point ring. However, in various embodiments, the stream format would support other topologies. A logical ring may include any combination of hardware, firmware, and/or software stream object interfaces. A two-object ring may include two one-way streams between the objects. An object that connects to multiple rings may have the capability to move, translate, and/or generate packages between rings to create the object memory fabric hierarchy.

[0317] FIG. 22A is a block diagram illustrating an example of streams present on a node 2200 with a hardware-based object memory fabric inter-node object router 2205, in accordance with certain embodiments of the present disclosure. In some embodiments, the node 2200 may correspond to a server node. The inter-node object router 2205 may include ring objects 2210 which are connected with physical streams 2215 in a ring orientation. In various embodiments, the ring objects may be connected in a ring 2220, which may be a virtual (Time Division Multiplexed) TDM ring in some embodiments. The ring objects 2210 and streams 2215 can be any combination of physical objects and streams or TDM ring objects and streams when hardware is shared. As depicted, one ring object 2210 may connect within the inter-node object router ring 2220 and to a stream 2225 that goes toward the object memory fabric router. In some embodiments, more than one ring object 2210 may connect within the inter-node object router ring and corresponding streams.

[0318] As depicted, the node 2200 may include a PCIe 2230, node memory controllers and DD4 memory buses 2235, and object memory fabric object memories 2240. Each object memory fabric object memory 2240 may have at least one pair of streams that connect to a inter-node object router ring object 2210 over the DD4 memory bus 2235 and PCIe 2230, running at hardware performance. As depicted, there can be software objects 2245 running on any processor core 2250 that can be functioning as any combination of routing agent and/or object memory. The software objects 2245 may have streams that connect ring objects 2210 within the inter-node object router 2205. Thus, such software objects 2245 streams may stream over the PCIe 2230.

[0319] FIG. 22B is a block diagram illustrating an example of software emulation of object memory and router on the node 2200-1, in accordance with certain embodiments of the present disclosure. The software object 2245 may, for example, emulate object memory fabric object memory 2240. The software object 2245 may include the same data structures to track objects and blocks and respond to requests from the inter-node object router 2205 identically to the actual object memory fabric object memory 2240. The software object 2245-1 may, for example, correspond to a routing agent by emulating the inter-node object router 2205 functionality. In so doing, the software object 2245-1 may communicate streams over standard wired and/or wireless networks, for example, to mobile, wired, and/or Internet of Things (IoT) devices 2255.

[0320] In some embodiments, the entire inter-node object router function could be implemented in one or more software objects 2245 running on one or more processing cores 2250, with the only difference being performance. And, as noted, one or more processing cores 2250 can also directly access object memory fabric object memory per conventional memory reference.

[0321] FIG. 23 is a block diagram illustrating an example of streams within an object memory fabric node object router 2300, in accordance with certain embodiments of the present disclosure. The object memory fabric router 2300 may include ring objects 2305 which are connected with streams 2310. As depicted, ring objects 2305 may be connected by streams 2310 in a ring topology. The ring objects 2305 and streams 2310 can be any combination of physical or TDM. One or more ring objects 2305 may connect to a physical stream 2315 that goes toward a leaf node. As depicted, one ring object 2305 may connect to a physical stream 2320 that goes toward a root node. In some embodiments, more than one ring object 2305 may connect to a respective physical stream 2320 that goes toward a root node.

#### [0322] API Background

[0323] Although API which stands for Applications Programming Interface, sounds like it should be about how software interfaces to object memory fabric, the main interface to object memory fabric may correspond to memory in some embodiments. In some embodiments, the object memory fabric API may correspond to how object memory fabric is set up and maintained transparently for applications, e.g., by modified Linux libc. Applications such as a SQL database or graph database can utilize the API to create object memory fabric objects and provide/augment meta-data to enable object memory fabric to better manage objects.

**[0324]** In various embodiments, overall capabilities of the API may include:

- [0325]** 1. Creating objects and maintaining objects within object memory fabric;
- [0326]** 2. Associating object memory fabric objects with local virtual address and physical address;
- [0327]** 3. Providing and augmenting meta-data to enable object memory fabric to better manage objects; and/or
- [0328]** 4. Specifying extended instruction functions and methods.

**[0329]** API functions may utilize the last capability to implement all capabilities. By being able to create functions and methods, entire native processor sequences can be offloaded to object memory fabric, gaining efficiencies such as those disclosed above with respect to the extended instruction environment and extended instructions.

**[0330]** The API interface may be through the PCIe-based Server Object Index, also referred to as object memory fabric inter-node object router. The API Programming model may directly integrate with the application. Multi-threading (through in memory command queue) may be provided so that each application is logically issuing commands. Each command may provide return status and optional data. The API commands may be available as part of trigger programs.

**[0331]** As noted regarding “Memory Fabric Distributed Object Memory and Index” (e.g. with respect to FIGS. 10-12 described herein), three components were introduced to describe the data structures and operation of the Object memory and index. The three components are shown below in Table 127. This section will discuss the physical instantiations in more depth.

TABLE 127

Logical Abstraction	Physical Device	Form Factor
Object Memory	Memory module/DIMM	DDR4 DIMM
Server Object Index	Node router	PCIe Card (half height & length)
Inter-node Router Object Index	Inter-node router	0.5 U Rack mount

**[0332]** Since all three form factors share a common functionality with respect to Object Memory and Index, the underlying design objects may be reused in all three (a common design).

**[0333]** FIG. 24 is a block diagram illustrating a product family hardware implementation architecture, in accordance with certain embodiments of the present disclosure.

**[0334]** Within a server, memory modules or DIMMs may plug into standard DDR4 memory sockets. Each memory module/DIMM may independently manage both dram memory (fast and relatively expensive) and flash memory (not as fast, but much less expensive) in a manner that the processor thinks that there is the flash amount of fast dram (see, for example, “Object Memory Caching” section herein). There may be eight memory sockets per processor socket or sixteen for a two-socket server. The node router or “uRouter” may communicate with the memory modules/DIMM(s) with direct memory access over PCIe and memory bus. The memory fabric may reserve a portion of each memory module/DIMM physical memory map to enable communication to and from the PCIe based node router/uRouter. Thus the combination of PCIe, memory bus and memory fabric private portion of memory module/DIMM

memory may form a virtual high bandwidth link. This may all be transparent to application execution.

**[0335]** The node router/uRouter may connect with an inter-node router or “IMF-Router” over 25/100GE fiber that uses several layers of Gigabit Ethernet protocol. Inter-node routers may connect with same 25/100GE fiber. An inter-node router may provide sixteen downlinks and two uplinks toward root. One embodiment may utilize dedicated links. Another embodiment may interoperate with standard links and routers.

**[0336]** FIG. 25 is a block diagram illustrating an alternative product family hardware implementation architecture, in accordance with certain embodiments of the present disclosure. This embodiment may provide an additional memory trigger instruction set and extended object method execution resources. This may enable a reduction in the number of servers that are required because more of the database storage manager and engine can execute within the object memory without need of server processor resources. A server-less memory fabric node may consist of sixteen object memories with a node router/uRouter. Ten nodes may be packaged into a single 1U rack mount enclosure, providing sixteen times reduction in space and up to five-times the performance improvement.

**[0337]** Server Node

**[0338]** The server may consist of a single node router/uRouter and one or more memory modules/DIMMs. The node router may implement the object index covering all objects/blocks held within the object memory(s) (memory modules) within the same server. The memory module may hold the actual objects and blocks within objects, corresponding object meta-data and object index covering objects currently stored locally. Each memory module independently manages both dram memory (which may be, for example, fast and relatively expensive) and flash memory (which may be, for example, not as fast, but much less expensive) in a manner that the processor thinks that there is the flash amount of fast dram. Both memory module and node router may manage free storage through a free storage index, which may be implemented in the same manner as for other indexes.

**[0339]** FIG. 26 is a block diagram illustrating a memory fabric server view of a hardware implementation architecture, in accordance with certain embodiments of the present disclosure.

**[0340]** Objects may be created and maintained through the memory fabric API as described herein. The API may communicate to the node router/uRouter through the memory fabric version of libc and memory fabric driver. The node router may then update the local object index, send commands toward the root as required and communicate with the appropriate memory module/DIMM to complete the API command (e.g. locally). Memory module may communicate an administrative request back to the node router, which may handle them appropriately both with respect to the memory fabric and the local Linux. The node router and memory module may participate in moving objects and blocks (e.g. in the manner described in the “Object Memory Miss” with respect to FIG. 12).

**[0341]** Memory Module/RDIMM

**[0342]** The RDIMM may consist of dram (e.g. 32 Gbyte), flash memory (e.g. 4 Terabytes) and FPGA and DDR4 compatible buffers (first generation product capacities per memory module). The FPGA may include all the resources,

structure, and internal data structures to manage the dram and flash as Object Memory integrated within the memory fabric whole.

**[0343]** FIG. 27 is a block diagram illustrating a memory module view of a hardware implementation architecture, in accordance with certain embodiments of the present disclosure.

**[0344]** A single scalable and parametrizable architecture may be used to implement the memory fabric on a memory module/DIMM as well as node router/uRouter and inter-node router/IMF-Router.

**[0345]** The internal architecture may be organized around a high performance, scalable ring interconnect that may implement a local version of memory fabric coherency protocol. Each subsystem may connect the ring through a coherent cache. The type of meta-data, data and objects stored may depend on the functionality of the subsystem. The routing engines in all three subsystems may be synthesized from a common design, may be highly multi-threaded, and may have no long term threads or state. An example set of routing engines may be as follows:

**[0346]** 1. Dram Routing Engine (StreamEngine): Controls memory module/DDR4 access, monitors triggers for processor access data and includes DDR4 cache. StreamEngine may monitor DDR4 operations for triggers and validate DDR4 cache access through an internal table that maps the 0.5 Tbyte physical memory module address space. This table has several possible implementations including:

**[0347]** a. Fully associative: Table that may convert each page physical number (excludes low 12 bits of address) to a page offset in DDR4. This has the advantage that any arbitrary set of pages can be cached.

**[0348]** b. Partially associative: Same as associative technique except that RAS address bits for the associative set and give the StreamEngine time to do the translation. This enables associativity level of 16-32 way, which is very close to the performance of fully associative. This technique requires a table of approximately 128 k×4 bits (512 k bits).

**[0349]** 2. Memory Fabric Background & API Engine (ExecuteEngine): May provides core memory fabric algorithms such as coherency, triggers, memory fabric APIs to accelerate graph and other big data as well as higher level memory fabric instruction sequences. May provide higher level API and memory fabric trigger execution. Also may handle background maintenance.

**[0350]** 3. OIT/POIT Engine: Manages OIT/POIT and provides this service to the other engines. The engine can process a level within an index in 2 cycles providing high performance index search and management. Manages flash storage for objects, meta-data blocks, data blocks and indices.

**[0351]** FIG. 28 is a block diagram illustrating a memory module view of a hardware implementation architecture, in accordance with an alternative embodiment of the present disclosure.

**[0352]** According to this embodiment, the capability of the multi-threaded memory fabric background & API engine may be functionally increased to execute a wide range of memory fabric trigger instructions. Additional instances of the updated multi-threaded memory fabric background & API engine may be added for more memory fabric trigger program performance. The combination of functional addi-

tions and more instances may be intended to enable memory fabric to execute big-data and data-manager software with fewer servers as shown, for example, in FIG. 28.

**[0353]** Node Router

**[0354]** The internal architecture of the node router/uRouter may be the same as the memory module/DIMM, with the differences related to the functionality of the node router, manage memory fabric server object index, and route appropriate packets to/from PCIe (memory modules) and inter-node router/IMF-Router. It may have additional routing function and may not actually store objects. As noted, an example set of routing engines may be as follows:

**[0355]** FIG. 29 is a block diagram illustrating a node router view of a hardware implementation architecture, in accordance with certain embodiments of the present disclosure.

**[0356]** 1. Routing Engine: Controls routing of packets to/from PCIe (memory modules) and inter-node router. Typically packets enter through one path are processed internally and exit on one of the paths.

**[0357]** 2. OIT/POIT Engine (ObjMemEngine): Manages OIT/POIT and provides this service to the other engines. The engine can process a level within an index in 2 cycles providing high performance index search and management. Manages flash and HMC (Hybrid Memory Cube) storage for indices. Caches most frequently used indices in HMC.

**[0358]** 3. Memory fabric background & API engine: Provides higher level API and memory fabric trigger execution. Also handles background maintenance.

**[0359]** Inter-node Router

**[0360]** FIG. 30 is a block diagram illustrating an inter-node router view of a hardware implementation architecture, in accordance with certain embodiments of the present disclosure.

**[0361]** The inter-node router may be analogous to an IP router. A difference may be the addressing model and static vs. dynamic. IP routers may utilize a fixed static address per each node and routes based on the destination IP address to a fixed physical node (can be virtualized for medium and long timeframes). The inter-node router may utilize a memory fabric object address (OA) which may specify the object and specific block of the object. Objects and blocks may dynamically reside at any node. The inter-node router may route OA packages based on the dynamic location(s) of objects and blocks and may track object/block location dynamically in real time.

**[0362]** The inter-node router may be a scaled up version of node router. Instead of connecting to a single PCIe bus to connect to leaf memory modules, it may connect multiple (e.g. 12-16, but expected to be 16) downlink node routers or inter-node routers and two uplink inter-node routers. There may also be a scale up of the object index storage capacity, processing rate and overall routing bandwidth.

**[0363]** FIG. 31 is a block diagram illustrating a memory fabric router view of a hardware implementation architecture, in accordance with certain embodiments of the present disclosure. The memory fabric architecture may utilize a memory fabric router for each downlink or uplink it connects to. The memory fabric router may be virtually identical to the node router (e.g. with the exception of supporting the internal memory fabric ring—which may be the same as the on chip version—and deleted PCIe). The memory fabric ring may utilize Interlaken protocol between memory fabric



routers. Interlaken protocol at the packet level may be compatible with utilizing 10G and 100G ethernet for downlinks and uplinks. Each memory fabric router may have as much object index storage capacity, processing rate and routing bandwidth as the node router, thus allowing the inter-node router to scale up to support the number of downlinks and uplinks.

[0364] Each downlink memory fabric router's object index may reflect all objects or blocks that are downlink from it. So even an inter-node router may use a distributed internal object index and routing.

[0365] The inter-node routers at any level with respect to leafs may be identical. The larger aggregate hierarchical object memory (caches) at each level from leaf may tend to lower the data movement between levels since more data can be stored at each level. Data that is in high use may be stored in multiple locations.

[0366] Implementation with standard software

[0367] The object-based memory fabric described above can provide native functions that can replace portions of virtual memory, in-memory file systems and database storage managers and store their respective data in a very efficient format. FIG. 32 is a block diagram illustrating object memory fabric functions that can replace software functions according to one embodiment of the present disclosure. As described in detail above, these object-based memory fabric functions can include functions 3205 for in-memory handling of blocks within objects through the object address space and functions 3210 for handling of objects through the object address and the local virtual address space of the node. Building on these functions 3205 and 3210, the object-based memory fabric can also provide in-memory file handling functions 3215, in-memory database functions 3220, and other in-memory functions 3225. Each of these in-memory functions 3215, 3220, and 3225 can, as described above, operate on the memory objects within the object-based memory fabric through the object address space and the virtual address space of the individual nodes of the object-based memory fabric. The object-based memory fabric and the functions provided thereby can be transparent to end user applications with minor changes to storage managers. While minor, these changes can create a huge increase in efficiency by storing data in an in-memory object format in the object infinite address space. The efficiency increase is two-fold: 1) the underlying in-memory object format and; 2) eliminating the conversions from storage and various database and/or application formats.

[0368] As introduced above, embodiments of the invention provide interfaces to the object-based memory fabric that can be implemented below the application level in the software stack. In this way, differences between the object-based memory and a standard address space are transparent to the applications which can utilize the object-based memory without modification, with the functional and performance benefits of object-based memory. Instead, modified storage managers can interface system software, such as a standard operating system, e.g., Linux, to the object-based memory. These modified storage managers can provide for management of standard processor hardware, such as buffers and caches, can control portions of the object-based memory space visible to the narrower physical address space available to the processor, and can be accessible by the applications through the standard, system software. In this way, the applications can access and utilize the object-based memory

fabric through the system software, e.g., through the standard operating system memory allocation process, without modification.

[0369] FIG. 33 is a block diagram illustrating an object memory fabric software stack according to one embodiment of the present disclosure. As illustrated in this example, the stack 3300 begins with and is built on top of the object-based memory fabric 3305 as described in detail above. A memory fabric operating system driver 3310 can provide access to the object-based memory space of the object-based memory fabric 3305 through memory allocation functions of the operating system of the node. In some cases, the operating system can comprise Linux or Security-Enhanced Linux (SELinux). The memory fabric operating system driver 3310 can also provide hooks to one or more virtual machines of the operating system.

[0370] In one implementation, the stack 3300 can also comprise an object-based memory specific version of a library file 3315 of the operating system. For example, this library file 3315 can comprise an object-based memory fabric specific version of a standard c library, libc. This library file 3315 can handle memory allocation and file system APIs in a manner appropriate to the object-based memory and that takes advantage of object-based memory fabric leverage. Additionally, the use of this library file 3315 and the functions therein can be transparent to application programs and users, i.e., they do not need to be treated different from the corresponding standard library functions.

[0371] The stack 3300 can further include a set of storage managers 3325, 3330, 3335, 3340, and 3345. Generally speaking, the storage managers 3325, 3330, 3335, 3340, and 3345 can comprise a set of modified storage managers that are adapted to utilize the format and addressing of the object-based memory space. The storage managers 3325, 3330, 3335, 3340, and 3345 can provide an interface between the object-based memory space and an operating system executed by the processor and an alternate object memory based storage transparent to a file system, database, or other software using the interface layer. The storage managers 3325, 3330, 3335, 3340, and 3345 can include, but are not limited to, a graph database storage manager 3325, an SQL or other relational database storage manager 3330, a filesystem storage manager 3335, and/or one or more other storage managers 3340 of different types.

[0372] According to one embodiment, a direct access interface 3320 allows a direct in-memory storage manager 3334 to directly access the object memory fabric 3305 with interfacing through the object memory fabric library file 3315. Since the memory fabric 3305 manages objects in a complete and coherent manner the direct storage manager 3345 can directly access the memory fabric 3305. Both the direct access interface 3320 and the direct memory manager 3345 are enabled by the capability of the memory fabric 3305 to coherently manage objects. This gives a path for a modified application to directly interface to the memory fabric class library 3315 or directly to the memory fabric 3305.

[0373] The object-based memory fabric additions to the software stack 3300 sit below the application level to provide compatibility between a set of unmodified applications 3350, 3355, 3360, and 3365 and the object-based memory fabric 3305. Such applications can include, but are not limited to, one or more standard graph database applications 3350, one or more standard SQL or other relational

database applications **3355**, one or more standard filesystem access applications **3360**, and/or one or more other standard, unmodified applications **3365**. The object-based memory fabric additions to the software stack **3300**, including the memory fabric operating system driver **3310**, object-based memory specific library file **3315**, and storage managers **3325**, **3330**, **3335**, **3340**, and **3345** can therefore provide an interface between the applications **3350**, **3355**, **3360**, and **3365** and the object-based memory fabric **3305**. This interface layer can control portions of the object-based memory space visible to a virtual address space and physical address space of the processor, i.e., a page fault and page handler that controls what portion of the object address space is currently visible in each node's physical address space and coordinating the relationship between memory objects and application segments and files. According to one embodiment, object access privilege for each application **3350**, **3355**, **3360**, and **3365** can be determined through an object-based memory fabric Access Control List (ACL) or equivalent.

[0374] Stated another way, each hardware-based processing node of an object memory fabric **3305**, such as described in detail above, can comprise a memory module storing and managing one or more memory objects within an object-based memory space. Also as described above, each memory object can be created natively within the memory module, accessed using a single memory reference instruction without Input/Output (I/O) instructions, and managed by the memory module at a single memory layer. The memory module can provide an interface layer **3310**, **3315**, **3320**, **3325**, **3330**, **3335**, **3340**, and **3345** below an application layer **3350**, **3355**, **3360**, and **3365** of a software stack **3300**. The interface layer can comprise one or more storage managers **3325**, **3330**, **3335**, **3340**, and **3345** managing hardware of a processor and controlling portions of the object-based memory space visible to a virtual address space and physical address space of the processor of each hardware-based processing node of the object-based memory fabric **3305**. The one or more storage managers **3325**, **3330**, **3335**, **3340**, and **3345** can further provide an interface between the object-based memory space and an operating system executed by the processor of each hardware-based processing node and an alternate object memory based storage transparent to a file system, database, or other software of the application layer **3350**, **3355**, **3360**, and **3365** of a software stack **3300** using the interface layer **3310**, **3315**, **3320**, **3325**, **3330**, **3335**, **3340**, and **3345**. In some cases, the operating system can comprise Linux or Security-Enhanced Linux (SELinux). Memory objects created and managed by the memory fabric can be created and managed equivalently from any node with the memory fabric. Thus a multi-node memory fabric does not require a centralized storage manager or memory fabric class library.

[0375] The interface layer **3310**, **3315**, **3320**, **3325**, **3330**, **3335**, **3340**, and **3345** can provide access to the object-based memory space to one or more applications executing in the application layer of the software stack access through memory allocation functions of the operating system. In one implementation, the interface layer can comprise an object-based memory specific version of a library file **3315** of the operating system. The one or more storage managers **3325**, **3330**, **3335**, **3340**, and **3345** can utilize a format and addressing of the object-based memory space. The one or

more storage managers can comprise, for example, a database manager **3330**, a graph database manager **3325**, and/or a filesystem manager **3335**.

[0376] Operations and Coherency Using Fault-Tolerant Objects

[0377] As introduced above, embodiments described herein can implement an object-based memory fabric in which memory objects in the memory fabric are distributed and tracked across a hierarchy of processing nodes. Each processing node can track memory objects and blocks within the memory objects that are present on paths from that node toward its leaf nodes in the hierarchy. Additionally, each processing node can utilize the same algorithms for memory object management such as memory object creation, block allocation, block coherency, etc. In this way, each higher level of the hierarchy creates an ever-larger cache which can significantly reduce the bandwidth in and out of the processing nodes at that level.

[0378] Fault tolerance capability can be implemented based on this hierarchical distribution and tracking by enabling memory objects, on a per-object basis, to be stored in more than a single node. This distribution of memory objects across multiple nodes can be across the hierarchy and/or across multiple physical locations. Memory object fault tolerance copies can be handled by a block coherency mechanism as part of memory fabric operation. In this way, each memory object can be made to be present on multiple different nodes. The memory object can be contained as a whole, within each of the multiple nodes, or at a given level of the hierarchy or may be stored as different portions with each portion being contained within multiple different nodes.

[0379] For illustrative purposes, reference will now be made to FIG. 7. The object memory fabric **700** can distribute and track the memory objects across the hierarchy of the object memory fabric **700** and the plurality of hardware-based processing nodes **705** and **710** on a per-object basis. Distributing the memory objects across the hierarchy of the object memory fabric **700** and the plurality of hardware-based processing nodes **705** and **710** can comprise storing, on a per-object basis, each memory object on two or more nodes of the plurality of hardware-based processing nodes **705** and **710** of the object memory fabric **700**. The two or more nodes of the plurality of hardware-based processing nodes **705** and **710** can be remote from each other in the hierarchy of the object memory fabric and/or in different physical locations.

[0380] More specifically, and as described above, the hierarchy can be a fat-tree structure. The Object Index Tree (OIT) and Per Object Index Tree (POIT) at each hardware-based processing node, node router, or inter-node router can track objects and blocks within objects that are present on all paths from that node toward its leaves. Thus, the algorithms for object creation, block allocation and block coherency can be the same at each node in the hierarchy. The power of such a hierarchy provides simplicity and efficiency of these algorithms. Each level of the hierarchy can create an ever-larger cache, which reduces the bandwidth in/out of that level. In the normal operating state where the working set is held within the corresponding level of the hierarchy, the bandwidth in/out of that level approaches zero.

[0381] Algorithm operations are now described within a node in the hierarchy with any interaction toward the root and/or toward the leaf. As used herein, "toward the root"

(root-path) refers to the direction from the node to the tree root while “toward the leaf” (leaf-path) refers to the direction from the node to the tree leaf. Hardware-based processing node memory modules, e.g., DIMMs, are considered the leaf within the hierarchy.

**[0382]** Each node in the hierarchy can track some number of paths toward the leaf, called tree-span factor. In the one implementation, the spanning factor at any level of the hierarchy can be 16. Thus, a memory module can keep track of objects stored locally. Other nodes, such as node routers and inter-node routers, can track, for example, up to 16 paths toward the leaf. In such an example, node routers and inter-node router OITs can keep the object state for the 16 paths and the POITs track block state of the 16 paths. Steady-state, most objects or blocks can be present in a single (or small number) of leaf path(s).

**[0383]** Although the unit of tracking can be a single block, the POIT can be organized so that it can provide a single entry for a grouping of blocks to improve POIT storage efficiency. This grouping can correspond to the storage chunk the persistent level of hierarchy (typically 4 blocks) or a defined object size.

**[0384]** Objects are created using the CreateObject instruction issued from a processor within a server or from processing on a hierarchy leaf (memory module/DIMM). If the object is already valid within the memory fabric, an indication of “already valid” is returned. The leaf which receives the CreateObject instruction can check locally to determine if the object is valid. If it is not, then the leaf can pass the CreateObject instruction toward the root and can record a local state of “pending\_create.” Eventually the CreateObject instruction can reach the hierarch root. The object is created and the CreateObject package is passed as successful back toward the requesting leaf, with each node along the way transitioning from pending\_create to valid object. If there are simultaneous CreateObject instructions for the same object, then the first to reach the root can be successful and the other CreateObject Instructions return that the object is already valid. Thus, software does not have to keep external locks to control creation.

**[0385]** The DestroyObject instruction can delete blocks and meta-data associated with the corresponding object in a hierarchical manner similar to CreateObject.

**[0386]** Once an object is created, any individual block within the object can be allocated, creating storage within the memory fabric for that block. The AllocateBlk instruction can work much like the CreateObject instruction. The AllocateBlk instruction can allocate a block of storage at the specified IOA within the memory fabric. The block can be stored in at least one memory module within the memory fabric or within at 2 or more memory modules if object fault tolerance is enabled for the object that the block is part of.

**[0387]** The AllocateBlk instruction can be issued from a processor within the server or from processing on a hierarchy leaf (memory module/DIMM). If the block is already

allocated within the memory fabric, an already allocated response can be returned. The leaf which receives the AllocateBlk instruction can check locally to determine if the block is allocated. If it is not, then the leaf can pass the AllocateBlk instruction toward the root and record a local state of pending\_allocate. Eventually the AllocateBlk reaches the hierarchy root. The block can be allocated and the AllocateBlk package can be passed as successful back toward the requesting leaf, with each node along the way transitioning from pending\_allocate to valid block state (usually own). If there are simultaneous AllocateBlk instructions for the same block, then the first to reach the root can be successful and the other AllocateBlk Instructions can return that the block is already allocated. Thus, software does not have to keep external locks to control allocation.

**[0388]** When an OA reference from a leaf is made to an OA object that is shown by IOT as not valid (invalid object or no local object), a root-path search can be made to establish the object is valid and implicitly create that object in the leaf-path back toward the requesting leaf. If the object is not valid, a not-valid status can be returned. When an IOA reference from a leaf is made to an IOA block that is shown by IOT or PIOT as not allocated (invalid block), a rootpath search can be made to establish the block is allocated and route the request to the leafpath of the allocated block. The leaf can then return block data to the requesting leaf with the path between the responding leaf to requesting leaf through their closest common root. If the block is not allocated within the memory fabric, a not-allocated status can be returned.

**[0389]** According to one embodiment, the memory fabric can optionally support allocation and deallocation of multiple blocks for more efficient allocation. Additionally or alternatively, the memory fabric can utilize a protocol as will be described in greater detail below and that can provide sequential consistency and forward progress guarantees across the memory fabric for applications including databases and filesystems. For example, the protocol can be an AllCache, ownership based, supporting update and invalidate modes. Use of such a protocol can provide lockless synchronization and can support integral object/block fault tolerance. The protocol can match the hierarchal distributed nature of object indices so that the coherency algorithm can be identical at each router. Thus, proving correctness in a single router provides correctness by construction for any arbitrary size memory fabric.

**[0390]** According to one embodiment, the memory fabric can support coherency on a 4 k block basis for object data and metadata. Memory fabric block state describes the state a level of the hierarchy including all leaves. Memory fabric package has the ability to optionally chain together requests so that many blocks (e.g., up to 227 blocks) can move on a single, initial request. Under certain conditions, the chain may be broken and can be retried. Each chained package can be pushed as an individual package for purposes of coherency.

TABLE I28

Memory Fabric Block & Object States		
Encoding	Symbol	Description
	local_object	object created on DL node or router
	no_local_object	No object allocated on DL node or router. Usually another DL node or router from this node is in state local_object.

TABLE I28-continued

Memory Fabric Block & Object States		
Encoding	Symbol	Description
0x0c	invalid_object	Object OIT entry allocated, but invalid object corresponding to DL node
0x0d	invalid_block	Object allocated, but block not allocated on corresponding DL node
0x00	invalid	Block allocated locally, but no data present (valid)
0x01	snapcopy	Snapshot copy. This copy is updated only when a block is persisted. Utilized for object fault tolerance. Can be configured on an object basis redundancy and geographic dispersion.
0x02	shadcopy	Shadow copy. Will be updated on a lazy basis-eventually consistent, usually after a period of time or some number of writes and/or transaction. Can also be used for fault tolerant block copies.
0x03	copy	Read-only copy. Will be updated for owner modifications as they occur. Insures sequential consistency.
0x04	own_snapcopy	Exclusive owner with snapshot copy. Enables local write privilege without any updates required. Snapshot copies may exist, but are only updated when corresponding block is persisted and through and push instruction with push_state = pstate_snapcopy.
0x08	own_snapcopy_m	
0x05	own_shadcopy	Non-exclusive owner with shadow copies. Enables write privilege shadow copies or snapshot copies to exist which are updated from writes on a lazy basis-eventually consistent.
0x09	own_shadcopy_m	
0x06	own_copy	Non-exclusive owner with copies. Enables write privilege and copies, shadow copies or snapshot copies to exist which are updated from writes. Multiple writes to the same block can occur with a single update.
0x0a	own_copy_m	
0x07	own_own_m	Exclusive owner. Enables local write privilege. No copies, shadow copies or snapshot copies exist.
0x0b		
0x0e	error	Error has been encountered on corresponding block.
0x0f	reserved	

[0391] According to one embodiment, when blocks are being requested and/or moved within the memory fabric, instances of the Object Index within memory modules, node routers and inter-node routers can track the request and movement with pending block states. The pending states enable multiple simultaneous requests to get services with a single or minimal number of responses. Pending states can be divided into 2 categories, leaf requests and remote root requests. Leaf requests can be received by a node from the

leaf direction. Remote requests can be requests that are received from the root direction for requests that progressed the maximum required depth toward the IMF root. Although infrequent, a request can be busied for retry as a simple hardware mechanism to handle the most complex cases. Through the hardware-based hierarchical memory fabric with integrated Index Trees including pending states at each node as described herein, software can be spared the burden of detecting performance robbing boundary cases.

TABLE I29

Memory Fabric Pending Block States		
Encoding	Symbol	Description
	pending_create	Object is being created on corresponding DL router or node
	pending_allocate	Storage for block is being allocated on corresponding DL router or node
	pending_destroy	Object is being destroyed on corresponding DL router or node
	pending_remote_destory	Object is being destroyed on corresponding DL router or node. Destroy initiated remotely from root.
	pending_deallocate	Storage for block is being deallocated on corresponding DL router or node
	pending_remote_deallocate	Storage for block is being deallocated on corresponding DL router or node. Deallocate initiated remotely from root.
	invalid_pown	invalid, pending own
	invalid_powncopy	invalid, pending own_copy
	invalid_pcopy	invalid, pending copy
	invalid_psnapcopy	invalid, pending snapcopy
	invalid_pshadcopy	invalid, pending shadcopy
	snapcopy_pown	snapcopy, pending own

TABLE I29-continued

Memory Fabric Pending Block States		
Encoding	Symbol	Description
	snapcopy_pown_copy	snapcopy, pending own_copy
	snapcopy_pcopy	snapcopy, pending copy
	snapcopy_pshadcop	snapcopy, pending shadcop
	snapcopy_remote_pull	snapcopy, pending remote (from root) pull_snapcopy
	snapcopy_remote_inval	snapcopy, pending remote (from root) invalidate
	snapcopy_remote	snapcopy, pending remote (from root) update
	shadcop_pown	shadcop, pending own
	shadcop_pown_copy	shadcop, pending own_copy
	shadcop_pcopy	shadcop, pending copy
	shadcop_remote_pull	shadcop, pending remote (from root) pull_shadcop
	shadcop_remote_inval	shadcop, pending remote (from root) invalidate
	shadcop_remote	shadcop, pending remote (from root) update
	copy_pown	copy, pending own
	copy_pown_copy	copy, pending own_copy
	copy_remote_pull	copy, pending remote (from root) pull_copy
	copy_remote_inval	copy, pending remote (from root) invalidate
	copy_remote	copy, pending remote (from root) update
	own_snapcopy_m_pupdate	own_snapcopy_m, pending update
	own_snapcopy_m_inval	own_snapcopy_m, pending invalidate
	own_snapcopy_ppush	own_snapcopy, pending push
	own_snapcopy_m_ppush	own_snapcopy_m, pending push
	own_shadcop_m_pupdate	own_shadcop_m, pending update
	own_shadcop_m_inval	own_shadcop_m, pending invalidate
	own_shadcop_ppush	own_shadcop, pending push
	own_shadcop_m_ppush	own_shadcop_m, pending push
	own_copy_m_pupdate	own_copy_m, pending update
	own_copy_m_inval	own_copy_m, pending invalidate
	own_copy_ppush	own_copy, pending push
	own_copy_m_ppush	own_copy_m, pending push
	own_ppush	own, pending push
	own_m_ppush	own_m, pending push

TABLE I30

Within Router packet status bits		
Bit Name	Description	
0 Busy	Indicates one or more routing nodes is unable to handle the request	
1 Copy	Indicates one or more routing nodes toward root contain a copy of corresponding address	
2 Shadow Copy	Indicates one or more routing nodes toward root contain a shadow copy of corresponding address	
3 Snapshot Copy	Indicates one or more routing nodes toward root contain a snapshot copy of corresponding address	

[0392] Routing through the hierarchical memory fabric can be based on package instruction, directory match on object address, match on appropriate level of node ID, acknowledge fields, and/or cache state of block. The routing decision outputs can include, but are not limited to, where to route package, update to block state, and/or set/clear acknowledge fields.

[0393] In one implementation, the memory fabric memory module as described herein can comprise a DIMM 32 Gbyte of fast DDR4 dram and 2-4 Tbytes of slower/less expensive flash memory. From the processor viewpoint, it manages a 2-4 Tbyte object memory with a 32 Gbyte dram cache. The memory fabric can be managed as a three-level memory hierarchy by taking advantage of two ideas. First, DDR DIMMs can indicate to the processor up to a maximum capacity of 0.5 Tbyte through direct physical addressing (PA). The memory module can indicate to the processor that it has 0.5 Tbyte of dram through 0.5 Tbyte physical address range and can fake that amount through caching from the

larger flash memory. Second, the memory module can utilize object triggers to predict and move data into dram, which can be being managed as a cache ahead of use. A reference count algorithm can be utilized to determine which objects/blocks can be removed from dram cache. Background cleaning of modified pages can also be implemented. In the low probability event that a processor requested physical address within the 0.5 Tbyte range is not immediately available, the memory module can signal a special recovery exception, which can then be fielded by the memory fabric exception handler as well as makes the request physical address available. The application and access can then be restarted.

[0394] The performance advantages of this technique are two-fold. First, the memory module behaves as if it were over 125 times larger or it appears that flash memory has the performance close dram, which is 1000x faster. Second, the operating system overhead of switching page table entries (PTEs) and PTE shoot-down is virtually eliminated.

[0395] FIG. 34 is a block diagram illustrating a summary of memory module caching according to one embodiment. As illustrated in this example, the memory module can dynamically manage the 3 levels of cache within the memory module/DIMM. The DRAM cache can be managed based on a set associate approach where the RAS addresses form the set index. This can provide approximately 256-way associativity for the DRAM cache which by all cache studies closely approximates fully associative.

[0396] FIG. 35 is a diagram illustrating an exemplary partitioning of the DDR4 dram cache for several functions related to the memory module according to one embodiment. In this example, PA Memory is the memory that is

physically addressed from the processor. According to one embodiment, almost 90% of the dram is allocated for this function. This partition can cache the 512 Gbyte physical memory address space of the memory module.

**[0397]** The Index Tree Cache (OIT & POIT Cache) can be a partition that caches portions of the OIT and POIT to minimize flash access. The allocated space as indicated in this example can cache the index tree for approximately 10% of a 4 Terabyte object memory.

**[0398]** The PA Directory can consist of the PA IOA Directory and PA DS Directory, both indexed by PA. The PA OA Directory can hold the OA (Object Address) for each block that has been associated with a processor physical address. The PA DS Directory can hold the DRAM slot number corresponding to each processor physical address. The PA DS Directory can also hold valid, modified and reference could information with respect to PA accesses.

**[0399]** The DramSlot Directory can be a directory of the blocks from the PA Directory that are currently in DRAM and can be indexed by DRAM block address. If the DRAM block is in use, a corresponding entry can contain the IOA (-128 bits) and associated PA (if associated, 27+1 valid bit). If the corresponding block is free it can contain a pointer in a free list chain.

DIMMs utilize a reference count algorithm to maintain and move blocks between levels of cache. An exemplary algorithm is explained above with reference to background operations and garbage collection.

**[0400]** FIG. 36 is a block diagram illustrating node and leaf caching according to one embodiment. According to one embodiment, each memory module, node router, and inter-node router can keep a separate instance of the OIT and POITs based on their place within the memory fabric hierarchy. Each instance independently caches OIT and POIT nodes and leafs based on activity and likely activity. Additionally, the direct IOA to POIT translation can be cached at the fastest level. In this case, the OIT and POIT access is eliminated. The levels of cache from fastest to slowest are: 1. IOA (IMF Object Address) to POIT Leaf translation; 2. Object Memory on chip cache for OIT/POIT nodes and leafs; 3. Local dram cache for OIT/POIT nodes and leafs; and 4. Local flash. The local flash can also provide persistence.

**[0401]** By caching individual nodes and leaves, the latency for OIT and POIT access can be reduced and overall throughput memory access can be increased. By caching the IOA to POIT Leaf translation OIT and POIT look up can be reduced to a single on-chip reference. According to one embodiment, each memory module can utilize a reference count algorithm to maintain and move blocks between levels of cache. An exemplary algorithm is explained above with reference to background operations and garbage collection.

**[0402]** Fault tolerance capability can be implemented enabling objects on a per object basis to be stored in more than a single node and/or in multiple physical locations. Object fault tolerance copies can be handled by the standard block coherency mechanism as part of basic memory fabric operation. Thus, blocks can be tracked and copies only updated on the block being updated. Therefore, object fault tolerance can have the inherent high performance of the memory fabric as described herein since minimum data movement is done. For each block of a fault tolerant object there can be a copy, snapcopy or shadcopy in addition to the own\_copy, own\_snapcopy or own\_shadcopy. The snapcopy

and corresponding own\_snapcopy enable the second copy to be updated when the object is persisted. The shadcopy and corresponding own\_shadcopy enables a more realtime, but lazy update. The copy and corresponding own\_copy enable a realtime update. According to one embodiment, fault tolerance can be extended to a mode of 3-5 copies. This mode can utilize the 16 to 1 spanning of a router as described above to put copies in parallel on several leaf-paths. When nodes are physically distributed, object fault tolerance can provide copies on leafs on both sides of physical distribution.

**[0403]** Memory Fabric Protocol and Interfaces

**[0404]** Embodiments of the invention provide systems and methods for managing processing, memory, storage, network, and cloud computing to significantly improve the efficiency and performance of processing nodes. Embodiments described herein can implement an object-based memory fabric in which memory objects in the memory fabric are distributed and tracked across a hierarchy of processing nodes. Each processing node can track memory objects and blocks within the memory objects that are present on paths from that node toward its leaf nodes in the hierarchy. Additionally, each processing node can utilize the same algorithms for memory object management such as memory object creation, block allocation, block coherency, etc. In this way, each higher level of the hierarchy creates an ever-larger cache which can significantly reduce the bandwidth in and out of the processing nodes at that level.

**[0405]** According to one embodiment, a highly threaded and latency tolerant memory fabric protocol can be used between nodes and routers within the memory fabric. This protocol can be implemented across dedicated links, e.g., 25/100GE (Gigabit Ethernet), and or can be tunneled over a standard link, e.g., an Internet Protocol (IP) link, enabling standard routers, such as IP routers, to be utilized between memory fabric routers. With that change, physically distributed memory fabric systems can be built which can cache data locally with neighborhoods of equal distant nodes and/or aggressively cache and duplicate objects that are utilized in multiple physical locations. Using a memory fabric protocol over such links, movement of memory objects across the memory fabric can be initiated at any node in the memory fabric hierarchy on a per object basis. Once a memory object is duplicated across nodes, changes to that memory object can be propagated to the other node(s) by moving only the data that changes rather than replicating or copying the memory object.

**[0406]** For illustrative purposes, reference will now be made to FIG. 7. As illustrated in this example, an object memory fabric 700 can comprise a plurality of hardware-based processing nodes 705 and 710. Each hardware-based processing node 705 and 710 can comprise one or more memory modules 725 and 730 storing and managing a plurality of memory objects in a hierarchy of the object memory fabric 700. Each memory object can be created natively within the memory module 725 or 730, accessed using a single memory reference instruction without Input/Output (I/O) instructions, and managed by the memory module 725 or 730 at a single memory layer. A node router 720 can be communicatively coupled with each of the one or more memory modules 725 and 730 of the node 705 and can be adapted to route memory objects or portions of memory objects between the one or more memory modules 725 and 730 of the node 705 using a memory fabric protocol.

One or more inter-node routers **715** can be communicatively coupled with each node router **720**. Each of the plurality of nodes **705** and **710** of the object memory fabric **700** can be communicatively coupled with at least one of the inter-node routers **715** and can be adapted to route memory objects or portions of memory objects between the plurality of nodes **705** and **710** using the memory fabric protocol.

**[0407]** Two or more nodes **705** and **710** of the plurality of hardware-based processing nodes can be remote from each other in the hierarchy of the object memory fabric **700**. In such cases, distributing and tracking the memory objects across the object memory fabric **700** can comprise creating neighborhoods of equal-distance nodes in the hierarchy of the object memory fabric. Additionally or alternatively, two or more nodes **705** and **710** of the plurality of hardware-based processing nodes can each be in different physical locations. In such cases, distributing and tracking the memory objects across the object memory fabric can comprise caching and duplicating objects in a plurality of different physical locations.

**[0408]** Distributing the memory objects across the object memory fabric can be initiated by one of the hardware-based processing nodes **705** on a per-object basis. Tracking the memory objects across the object memory fabric **705** can comprise tracking, by the hardware-based processing node **705**, the memory objects and blocks within the memory objects that are present on the hardware-based processing node **705**. Tracking the memory objects across the object memory fabric **700** can also comprise propagating changes to the memory objects from one of the hardware-based processing node **705** to one or more other nodes **710** of the plurality of hardware-based processing nodes. Propagating changes to the memory objects from the hardware-based processing node **705** to one or more other nodes **710** of the plurality of hardware-based processing nodes can comprise moving only data within the memory objects that has changed and without replicating or copying the memory object.

**[0409]** As introduced above, the memory fabric **700** can utilize a protocol that can provide sequential consistency and forward progress guarantees across the memory fabric **700** for applications including databases and filesystems. Use of such a protocol can provide lockless synchronization and can support integral object/block fault tolerance. The protocol can match the hierarchical distributed nature of object indices so that the coherency algorithm can be identical at each router. Thus, proving correctness in a single router provides correctness by construction for any arbitrary size memory fabric. The underlying memory fabric protocol between inter-node routers **715** and/or node routers **720** can be highly threaded and latency tolerant.

**[0410]** The object memory fabric **700** can utilize the memory fabric protocol between the hardware-based processing nodes **705** and one or more other nodes **710** of the plurality of hardware-based processing nodes to distribute and track the memory objects across the object memory fabric **700**. The memory fabric protocol can be utilized across a dedicated link between the hardware-based processing node **705** and one or more other nodes **710** of the plurality of hardware-based processing nodes. For example, the dedicated link can comprise an Ethernet link. In other implementations, the memory fabric protocol can be tunneled across a shared link between the hardware-based processing node **705** and one or more other nodes **710** of the

plurality of hardware-based processing nodes. For example, the shared link can comprise an Internet Protocol (IP) link. In such cases, the memory fabric protocol can provide a dedicated communication link between the hardware-based processing node **705** and one or more other nodes **710** of the plurality of hardware-based processing nodes and the shared link supports communications other than the memory fabric protocol.

**[0411]** By way example, in the case of an IP link, by using the User Datagram Protocol (UDP), the memory fabric protocol can be encapsulated in UDP messages between node routers and inter-node routers and between inter-node routers, within standard IP network routers and switches. The memory fabric node routers and inter-node routers are responsible for memory fabric protocol and error checking. Thus standard IP network routers and switches properly route memory fabric UDP packets without any knowledge of the object-based memory fabric or memory fabric protocol. Now the links between node router and inter-node router and different inter-node routers are point to point in a virtual sense, but utilize the shared IP network to provide distributed virtual connectivity of memory fabric nodes, within a data center, across data centers, distributed memory fabric nodes or mobile memory fabric nodes.

**[0412]** By using such a protocol and the coherency mechanisms described above, the underlying ability for the memory fabric to cache data locally with neighborhoods of equal distant nodes as described herein can be exploited to more aggressively cache and duplicate objects that are utilized in multiple physical locations. The chaining and memory fabric instruction set enable this data movement to be initiated at any node in the memory fabric hierarchy on a per object basis. Once objects are duplicated only the data that changes moves.

**[0413]** The memory fabric packets that are sent between the inter-node router and/or node routers can be referenced by OA (Object Address) and based in a conventional memory type protocol. A link level reliability protocol with packet numbering and acknowledgement can be included and can reduce reliance on the relatively inefficient IP protocol.

**[0414]** The present disclosure, in various aspects, embodiments, and/or configurations, includes components, methods, processes, systems, and/or apparatus substantially as depicted and described herein, including various aspects, embodiments, configurations embodiments, sub combinations, and/or subsets thereof. Those of skill in the art will understand how to make and use the disclosed aspects, embodiments, and/or configurations after understanding the present disclosure. The present disclosure, in various aspects, embodiments, and/or configurations, includes providing devices and processes in the absence of items not depicted and/or described herein or in various aspects, embodiments, and/or configurations hereof, including in the absence of such items as may have been used in previous devices or processes, e.g., for improving performance, achieving ease and/or reducing cost of implementation.

**[0415]** The foregoing discussion has been presented for purposes of illustration and description. The foregoing is not intended to limit the disclosure to the form or forms disclosed herein. In the foregoing Detailed Description for example, various features of the disclosure are grouped together in one or more aspects, embodiments, and/or configurations for the purpose of streamlining the disclosure.

The features of the aspects, embodiments, and/or configurations of the disclosure may be combined in alternate aspects, embodiments, and/or configurations other than those discussed above. This method of disclosure is not to be interpreted as reflecting an intention that the claims require more features than are expressly recited in each claim. Rather, as the following claims reflect, inventive aspects lie in less than all features of a single foregoing disclosed aspect, embodiment, and/or configuration. Thus, the following claims are hereby incorporated into this Detailed Description, with each claim standing on its own as a separate preferred embodiment of the disclosure.

**[0416]** Moreover, though the description has included description of one or more aspects, embodiments, and/or configurations and certain variations and modifications, other variations, combinations, and modifications are within the scope of the disclosure, e.g., as may be within the skill and knowledge of those in the art, after understanding the present disclosure. It is intended to obtain rights which include alternative aspects, embodiments, and/or configurations to the extent permitted, including alternate, interchangeable and/or equivalent structures, functions, ranges or steps to those claimed, whether or not such alternate, interchangeable and/or equivalent structures, functions, ranges or steps are disclosed herein, and without intending to publicly dedicate any patentable subject matter.

What is claimed is:

1. A hardware-based processing node of a plurality of hardware-based processing nodes in an object memory fabric, the hardware-based processing node comprising:

a memory module storing and managing a plurality of memory objects in a hierarchy of the object memory fabric, wherein:

each memory object is created natively within the memory module,

each memory object is accessed using a single memory reference instruction without Input/Output (I/O) instructions,

each memory object is managed by the memory module at a single memory layer, and

the object memory fabric distributes and tracks the memory objects across the hierarchy of the object memory fabric and the plurality of hardware-based processing nodes on a per-object basis.

2. The hardware-based processing node of claim 1, wherein distributing the memory objects across the hierarchy of the object memory fabric and the plurality of hardware-based processing nodes comprises storing, on a per-object basis, each memory object on two or more nodes of the plurality of hardware-based processing nodes of the object memory fabric.

3. The hardware-based processing node of claim 2, wherein the two or more nodes of the plurality of hardware-based processing nodes are remote from each other in the hierarchy of the object memory fabric.

4. The hardware-based processing node of claim 2, wherein the two or more nodes of the plurality of hardware-based processing nodes are in different physical locations.

5. The hardware-based processing node of claim 1, wherein tracking the memory objects across the hierarchy of the object memory fabric comprises tracking, by the hardware-based processing node, the memory objects and blocks within the memory objects that are present on the hardware-based processing node.

6. The hardware-based processing node of claim 1, wherein tracking the memory objects across the hierarchy of the object memory fabric comprises tracking the memory objects and blocks within the memory objects that are present on each level of the hierarchy of the object memory fabric and branches from each level of the hierarchy towards leaves of the hierarchy of the object memory fabric.

7. The hardware-based processing node of claim 1, wherein the hardware-based processing node utilizes a same algorithm for object management as each other node of the plurality of hardware-based processing nodes.

8. The hardware-based processing node of claim 7, wherein the algorithm for object management comprises an algorithm for object creation.

9. The hardware-based processing node of claim 7, wherein the algorithm for object management comprises an algorithm for block allocation.

10. The hardware-based processing node of claim 7, wherein the algorithm for object management comprises an algorithm for block coherency.

11. The hardware-based processing node of claim 1, wherein the hardware-based processing node comprises a Dual In-line Memory Module (DIMM) card.

12. The hardware-based processing node of claim 1, wherein the hardware-based processing node comprises a commodity server and wherein the memory module comprises a Dual In-line Memory Module (DIMM) card installed within the commodity server.

13. The hardware-based processing node of claim 1, wherein the hardware-based processing node comprises a mobile computing device.

14. The hardware-based processing node of claim 1, wherein the hardware-based processing node comprises a single chip.

15. An object memory fabric comprising:

a plurality of hardware-based processing nodes, each hardware-based processing node comprising:

a memory module storing and managing a plurality of memory objects in a hierarchy of the object memory fabric, wherein each memory object is created natively within the memory module, each memory object is accessed using a single memory reference instruction without Input/Output (I/O) instructions,

each memory object is managed by the memory module at a single memory layer, and the object memory fabric distributes and tracks the memory objects across the hierarchy of the object memory fabric and the plurality of hardware-based processing nodes on a per-object basis, and

a node router communicatively coupled with each of the one or more memory modules of the node and adapted to route memory objects or portions of memory objects between the one or more memory modules of the node; and

one or more inter-node routers communicatively coupled with each node router, wherein each of the plurality of nodes of the object memory fabric is communicatively coupled with at least one of the inter-node routers and adapted to route memory objects or portions of memory objects between the plurality of nodes.

16. The object memory fabric of claim 15, wherein distributing the memory objects across the hierarchy of the object memory fabric and the plurality of hardware-based processing nodes comprises storing, on a per-object basis,



each memory object on two or more nodes of the plurality of hardware-based processing nodes of the object memory fabric.

17. The object memory fabric of claim 16, wherein the two or more nodes of the plurality of hardware-based processing nodes are remote from each other in the hierarchy of the object memory fabric.

18. The object memory fabric of claim 16, wherein the two or more nodes of the plurality of hardware-based processing nodes are in different physical locations.

19. The object memory fabric of claim 15, wherein tracking the memory objects across the hierarchy of the object memory fabric comprises tracking, by each hardware-based processing node, the memory objects and blocks within the memory objects that are present on the hardware-based processing node.

20. The object memory fabric of claim 15, wherein tracking the memory objects across the hierarchy of the object memory fabric comprises tracking the memory objects and blocks within the memory objects that are present on each level of the hierarchy of the object memory fabric and branches from each level of the hierarchy towards leaves of the hierarchy of the object memory fabric.

21. The object memory fabric of claim 15, wherein each hardware-based processing node utilizes a same algorithm for object management.

22. The object memory fabric of claim 21, wherein the algorithm for object management comprises an algorithm for object creation.

23. The object memory fabric of claim 21, wherein the algorithm for object management comprises an algorithm for block allocation.

24. The object memory fabric of claim 21, wherein the algorithm for object management comprises an algorithm for block coherency.

25. The object memory fabric of claim 15, wherein at least one of the plurality of hardware-based processing nodes comprises a Dual In-line Memory Module (DIMM) card.

26. The object memory fabric of claim 15, wherein at least one of the plurality of hardware-based processing nodes comprises a commodity server and wherein the memory module comprises a Dual In-line Memory Module (DIMM) card installed within the commodity server.

27. The object memory fabric of claim 15, wherein at least one of the plurality of hardware-based processing nodes comprises a mobile computing device.

28. The object memory fabric of claim 15, wherein at least one of the plurality of hardware-based processing nodes comprises a single chip.

29. A method for providing coherency and fault tolerance in an object memory fabric comprising a plurality of hardware-based processing nodes, the method comprising:

creating, by the hardware-based processing nodes of the object-based memory fabric, each memory object natively within a memory module of the hardware-based processing node;

accessing, by the hardware-based processing nodes, each memory object using a single memory reference instruction without Input/Output (I/O) instructions;

managing, by the hardware-based processing nodes, each memory object within the memory module at a single memory layer; and

distributing the memory objects across a hierarchy of the object memory fabric and the plurality of hardware-based processing nodes on a per-object basis; and

tracking the memory objects across the hierarchy of the object memory fabric and the plurality of hardware-based processing nodes on the per-object basis.

30. The method of claim 29, wherein distributing the memory objects across the hierarchy of the object memory fabric and the plurality of hardware-based processing nodes comprises storing, on a per-object basis, each memory object on two or more nodes of the plurality of hardware-based processing nodes of the object memory fabric.

31. The method of claim 30, wherein the two or more nodes of the plurality of hardware-based processing nodes are remote from each other in the hierarchy of the object memory fabric.

32. The method of claim 30, wherein the two or more nodes of the plurality of hardware-based processing nodes are in different physical locations.

33. The method of claim 29, wherein tracking the memory objects across the hierarchy of the object memory fabric comprises tracking, by the hardware-based processing node, the memory objects and blocks within the memory objects that are present on the hardware-based processing node.

34. The method of claim 29, wherein tracking the memory objects across the hierarchy of the object memory fabric comprises tracking the memory objects and blocks within the memory objects that are present on each level of the hierarchy of the object memory fabric and branches from each level of the hierarchy towards leaves of the hierarchy of the object memory fabric.

35. The method of claim 29, wherein the hardware-based processing node utilizes a same algorithm for object management as each other node of the plurality of hardware-based processing nodes.

36. The method of claim 35, wherein the algorithm for object management comprises an algorithm for object creation.

37. The method of claim 35, wherein the algorithm for object management comprises an algorithm for block allocation.

38. The method of claim 35, wherein the algorithm for object management comprises an algorithm for block coherency.

\* \* \* \* \*