



- (51) **International Patent Classification:**
G06F 13/14 (2006.01) *G06F 13/38* (2006.01)
- (21) **International Application Number:**
PCT/US2013/032718
- (22) **International Filing Date:**
16 March 2013 (16.03.2013)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
61/717,091 22 October 2012 (22.10.2012) US
- (72) **Inventors; and**
- (71) **Applicants :** WILLEY, Jeff [US/US]; 6750 Duncan Ct., Timnath, Colorado 90547 (US). BLANKENSHIP, Robert G. [US/US]; 3115 N. 20th St, Tacoma, Washington 98406 (US). SWANSON, Jeffrey C. [US/US]; 792 Trenton Drive, Sunnyvale, California 94087 (US). SAFRANEK,


Robert J. [US/US]; 5816 NW Necanicum Way, Portland, Oregon 97229 (US).

- (74) **Agent:** KOMENDA, J., Kyle; Patent Capital Group, c/o CPA Global, P.O. Box 52050, Minneapolis, MN 55402 (US).

- (81) **Designated States** (*unless otherwise indicated, for every kind of national protection available*): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

[Continued on next page]

- (54) **Title:** MULTISLOT LINK LAYER FLIT



- (57) **Abstract:** Transaction data is identified and a flit is generated to include three or more slots and a floating field to be used as an extension of any one of two or more of the slots. The flit is sent over a serial data link to a device for processing, based at least in part on the three or more slots.



(84) **Designated States** (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

— *of inventorship (Rule 4.17(iv))*

Published:

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

MULTISLOT LINK LAYER FLIT**FIELD**

[0001] The present disclosure relates in general to the field of computer development, and more specifically, to software development involving coordination of mutually-dependent constrained systems.

BACKGROUND

[0002] Advances in semi-conductor processing and logic design have permitted an increase in the amount of logic that may be present on integrated circuit devices. As a corollary, computer system configurations have evolved from a single or multiple integrated circuits in a system to multiple cores, multiple hardware threads, and multiple logical processors present on individual integrated circuits, as well as other interfaces integrated within such processors. A processor or integrated circuit typically comprises a single physical processor die, where the processor die may include any number of cores, hardware threads, logical processors, interfaces, memory, controller hubs, etc.

[0003] As a result of the greater ability to fit more processing power in smaller packages, smaller computing devices have increased in popularity. Smartphones, tablets, ultrathin notebooks, and other user equipment have grown exponentially. However, these smaller devices are reliant on servers both for data storage and complex processing that exceeds the form factor. Consequently, the demand in the high-performance computing market (i.e. server space) has also increased. For instance, in modern servers, there is typically not only a single processor with multiple cores, but also multiple physical processors (also referred to as multiple sockets) to increase the computing power. But as the processing power grows along with the number of devices in a computing system, the communication between sockets and other devices becomes more critical.

[0004] In fact, interconnects have grown from more traditional multi-drop buses that primarily handled electrical communications to full blown interconnect architectures that facilitate fast communication. Unfortunately, as the demand for future processors to consume at even higher-rates corresponding demand is placed on the capabilities of existing interconnect architectures.

BRIEF DESCRIPTION OF THE DRAWINGS

[0005] FIG. 1 illustrates a simplified block diagram of a system including a serial point-to-point interconnect to connect I/O devices in a computer system in accordance with one embodiment;

[0006] FIG. 2 illustrates a simplified block diagram of a layered protocol stack in accordance with one embodiment;

[0007] FIG. 3 illustrates an embodiment of a serial point-to-point link.

[0008] FIG. 4 illustrates embodiments of potential High Performance Interconnect (HPI) system configurations.

[0009] FIG. 5 illustrates an embodiment of a layered protocol stack associated with HPI.

[0010] FIG. 6 illustrates a representation of an example multi-slot flit.

[0011] FIG. 7 illustrates a representation of an example flit sent over an example eight-lane data link.

[0012] FIG. 8 illustrates a representation of an example flit sent over an example eight-lane data link.

[0013] FIG. 9 illustrates a representation of an example flit sent over an example twenty-lane data link.

[0014] FIG. 10 illustrates a representation of use of an example floating payload field of an example multi-slot flit.

[0015] FIG. 11 illustrates an embodiment of a block for an example computing system.

[0016] Like reference numbers and designations in the various drawings indicate like elements.

DETAILED DESCRIPTION

[0017] In the following description, numerous specific details are set forth, such as examples of specific types of processors and system configurations, specific hardware structures, specific architectural and micro architectural details, specific register configurations, specific instruction types, specific system components, specific processor

pipeline stages, specific interconnect layers, specific packet/transaction configurations, specific transaction names, specific protocol exchanges, specific link widths, specific implementations, and operation etc. in order to provide a thorough understanding of the present invention. It may be apparent, however, to one skilled in the art that these specific details need not necessarily be employed to practice the subject matter of the present disclosure. In other instances, well detailed description of known components or methods has been avoided, such as specific and alternative processor architectures, specific logic circuits/code for described algorithms, specific firmware code, low-level interconnect operation, specific logic configurations, specific manufacturing techniques and materials, specific compiler implementations, specific expression of algorithms in code, specific power down and gating techniques/logic and other specific operational details of computer system in order to avoid unnecessarily obscuring the present disclosure.

[0018] Although the following embodiments may be described with reference to energy conservation, energy efficiency, processing efficiency, and so on in specific integrated circuits, such as in computing platforms or microprocessors, other embodiments are applicable to other types of integrated circuits and logic devices. Similar techniques and teachings of embodiments described herein may be applied to other types of circuits or semiconductor devices that may also benefit from such features. For example, the disclosed embodiments are not limited to server computer system, desktop computer systems, laptops, Ultrabooks™, but may be also used in other devices, such as handheld devices, smartphones, tablets, other thin notebooks, systems on a chip (SOC) devices, and embedded applications. Some examples of handheld devices include cellular phones, Internet protocol devices, digital cameras, personal digital assistants (PDAs), and handheld PCs. Here, similar techniques for a high-performance interconnect may be applied to increase performance (or even save power) in a low power interconnect. Embedded applications typically include a microcontroller, a digital signal processor (DSP), a system on a chip, network computers (NetPC), set-top boxes, network hubs, wide area network (WAN) switches, or any other system that can perform the functions and operations taught below. Moreover, the apparatus', methods, and systems described herein are not limited to physical computing devices, but may also relate to software optimizations for energy conservation and efficiency. As may become readily apparent in the description below, the embodiments of methods, apparatus', and systems described herein (whether in reference to hardware, firmware, software, or a combination thereof) may be considered vital to a "green technology" future balanced with performance considerations.

[0019] As computing systems are advancing, the components therein are becoming more complex. The interconnect architecture to couple and communicate between the components has also increased in complexity to ensure bandwidth demand is met for optimal component operation. Furthermore, different market segments demand different aspects of interconnect architectures to suit the respective market. For example, servers require higher performance, while the mobile ecosystem is sometimes able to sacrifice overall performance for power savings. Yet, it is a singular purpose of most fabrics to provide highest possible performance with maximum power saving. Further, a variety of different interconnects can potentially benefit from subject matter described herein. For instance, the Peripheral Component Interconnect (PCI) Express (PCIe) interconnect fabric architecture and QuickPath Interconnect (QPI) fabric architecture, among other examples, can potentially be improved according to one or more principles described herein, among other examples.

[0020] FIG. 1 illustrates one embodiment of a fabric composed of point-to-point Links that interconnect a set of components is illustrated. System 100 includes processor 105 and system memory 110 coupled to controller hub 115. Processor 105 can include any processing element, such as a microprocessor, a host processor, an embedded processor, a co-processor, or other processor. Processor 105 is coupled to controller hub 115 through front-side bus (FSB) 106. In one embodiment, FSB 106 is a serial point-to-point interconnect as described below. In another embodiment, link 106 includes a serial, differential interconnect architecture that is compliant with different interconnect standard.

[0021] System memory 110 includes any memory device, such as random access memory (RAM), non-volatile (NV) memory, or other memory accessible by devices in system 100. System memory 110 is coupled to controller hub 115 through memory interface 116. Examples of a memory interface include a double-data rate (DDR) memory interface, a dual-channel DDR memory interface, and a dynamic RAM (DRAM) memory interface.

[0022] In one embodiment, controller hub 115 can include a root hub, root complex, or root controller, such as in a PCIe interconnection hierarchy. Examples of controller hub 115 include a chipset, a memory controller hub (MCH), a northbridge, an interconnect controller hub (ICH) a southbridge, and a root controller/hub. Often the term chipset refers to two physically separate controller hubs, e.g., a memory controller hub (MCH) coupled to an interconnect controller hub (ICH). Note that current systems often include the MCH integrated with processor 105, while controller 115 is to communicate with I/O devices, in a similar

manner as described below. In some embodiments, peer-to-peer routing is optionally supported through root complex 115.

[0023] Here, controller hub 115 is coupled to switch/bridge 120 through serial link 119. Input/output modules 117 and 121, which may also be referred to as interfaces/ports 117 and 121, can include/implement a layered protocol stack to provide communication between controller hub 115 and switch 120. In one embodiment, multiple devices are capable of being coupled to switch 120.

[0024] Switch/bridge 120 routes packets/messages from device 125 upstream, i.e. up a hierarchy towards a root complex, to controller hub 115 and downstream, i.e. down a hierarchy away from a root controller, from processor 105 or system memory 110 to device 125. Switch 120, in one embodiment, is referred to as a logical assembly of multiple virtual PCI-to-PCI bridge devices. Device 125 includes any internal or external device or component to be coupled to an electronic system, such as an I/O device, a Network Interface Controller (NIC), an add-in card, an audio processor, a network processor, a hard-drive, a storage device, a CD/DVD ROM, a monitor, a printer, a mouse, a keyboard, a router, a portable storage device, a Firewire device, a Universal Serial Bus (USB) device, a scanner, and other input/output devices. Often in the PCIe vernacular, such as device, is referred to as an endpoint. Although not specifically shown, device 125 may include a bridge (e.g., a PCIe to PCI/PCI-X bridge) to support legacy or other versions of devices or interconnect fabrics supported by such devices.

[0025] Graphics accelerator 130 can also be coupled to controller hub 115 through serial link 132. In one embodiment, graphics accelerator 130 is coupled to an MCH, which is coupled to an ICH. Switch 120, and accordingly I/O device 125, is then coupled to the ICH. I/O modules 131 and 118 are also to implement a layered protocol stack and associated logic to communicate between graphics accelerator 130 and controller hub 115. Similar to the MCH discussion above, a graphics controller or the graphics accelerator 130 itself may be integrated in processor 105.

[0026] Turning to FIG. 2 an embodiment of a layered protocol stack is illustrated. Layered protocol stack 200 can include any form of a layered communication stack, such as a QPI stack, a PCIe stack, a next generation high performance computing interconnect (HPI) stack, or other layered stack. In one embodiment, protocol stack 200 can include transaction layer 205, link layer 210, and physical layer 220. An interface, such as interfaces 117, 118, 121, 122, 126, and 131 in FIG. 1, may be represented as communication protocol stack 200.

Representation as a communication protocol stack may also be referred to as a module or interface implementing/including a protocol stack.

[0027] Packets can be used to communicate information between components. Packets can be formed in the Transaction Layer 205 and Data Link Layer 210 to carry the information from the transmitting component to the receiving component. As the transmitted packets flow through the other layers, they are extended with additional information used to handle packets at those layers. At the receiving side the reverse process occurs and packets get transformed from their Physical Layer 220 representation to the Data Link Layer 210 representation and finally (for Transaction Layer Packets) to the form that can be processed by the Transaction Layer 205 of the receiving device.

[0028] In one embodiment, transaction layer 205 can provide an interface between a device's processing core and the interconnect architecture, such as Data Link Layer 210 and Physical Layer 220. In this regard, a primary responsibility of the transaction layer 205 can include the assembly and disassembly of packets (i.e., transaction layer packets, or TLPs). The translation layer 205 can also manage credit-based flow control for TLPs. In some implementations, split transactions can be utilized, i.e., transactions with request and response separated by time, allowing a link to carry other traffic while the target device gathers data for the response, among other examples.

[0029] Credit-based flow control can be used to realize virtual channels and networks utilizing the interconnect fabric. In one example, a device can advertise an initial amount of credits for each of the receive buffers in Transaction Layer 205. An external device at the opposite end of the link, such as controller hub 115 in FIG. 1, can count the number of credits consumed by each TLP. A transaction may be transmitted if the transaction does not exceed a credit limit. Upon receiving a response an amount of credit is restored. One example of an advantage of such a credit scheme is that the latency of credit return does not affect performance, provided that the credit limit is not encountered, among other potential advantages.

[0030] In one embodiment, four transaction address spaces can include a configuration address space, a memory address space, an input/output address space, and a message address space. Memory space transactions include one or more of read requests and write requests to transfer data to/from a memory-mapped location. In one embodiment, memory space transactions are capable of using two different address formats, e.g., a short address format, such as a 32-bit address, or a long address format, such as 64-bit address. Configuration space

transactions can be used to access configuration space of various devices connected to the interconnect. Transactions to the configuration space can include read requests and write requests. Message space transactions (or, simply messages) can also be defined to support in-band communication between interconnect agents. Therefore, in one example embodiment, transaction layer 205 can assemble packet header/payload 206.

[0031] A Link layer 210, also referred to as data link layer 210, can act as an intermediate stage between transaction layer 205 and the physical layer 220. In one embodiment, a responsibility of the data link layer 210 is providing a reliable mechanism for exchanging Transaction Layer Packets (TLPs) between two components on a link. One side of the Data Link Layer 210 accepts TLPs assembled by the Transaction Layer 205, applies packet sequence identifier 211, i.e. an identification number or packet number, calculates and applies an error detection code, i.e. CRC 212, and submits the modified TLPs to the Physical Layer 220 for transmission across a physical to an external device.

[0032] In one example, physical layer 220 includes logical sub block 221 and electrical sub-block 222 to physically transmit a packet to an external device. Here, logical sub-block 221 is responsible for the "digital" functions of Physical Layer 221. In this regard, the logical sub-block can include a transmit section to prepare outgoing information for transmission by physical sub-block 222, and a receiver section to identify and prepare received information before passing it to the Link Layer 210.

[0033] Physical block 222 includes a transmitter and a receiver. The transmitter is supplied by logical sub-block 221 with symbols, which the transmitter serializes and transmits onto to an external device. The receiver is supplied with serialized symbols from an external device and transforms the received signals into a bit-stream. The bit-stream is de-serialized and supplied to logical sub-block 221. In one example embodiment, an 8b/10b transmission code is employed, where ten-bit symbols are transmitted/received. Here, special symbols are used to frame a packet with frames 223. In addition, in one example, the receiver also provides a symbol clock recovered from the incoming serial stream.

[0034] As stated above, although transaction layer 205, link layer 210, and physical layer 220 are discussed in reference to a specific embodiment of a protocol stack (such as a PCIe protocol stack), a layered protocol stack is not so limited. In fact, any layered protocol may be included/implemented and adopt features discussed herein. As an example, a port/interface that is represented as a layered protocol can include: (1) a first layer to assemble packets, i.e. a transaction layer; a second layer to sequence packets, i.e. a link layer; and a third

layer to transmit the packets, i.e. a physical layer. As a specific example, a high performance interconnect layered protocol, as described herein, is utilized.

[0035] Referring next to FIG. 3, an example embodiment of a serial point to point fabric is illustrated. A serial point-to-point link can include any transmission path for transmitting serial data. In the embodiment shown, a link can include two, low-voltage, differentially driven signal pairs: a transmit pair 306/311 and a receive pair 312/307. Accordingly, device 305 includes transmission logic 306 to transmit data to device 310 and receiving logic 307 to receive data from device 310. In other words, two transmitting paths, i.e. paths 316 and 317, and two receiving paths, i.e. paths 318 and 319, are included in some implementations of a link.

[0036] A transmission path refers to any path for transmitting data, such as a transmission line, a copper line, an optical line, a wireless communication channel, an infrared communication link, or other communication path. A connection between two devices, such as device 305 and device 310, is referred to as a link, such as link 315. A link may support one lane – each lane representing a set of differential signal pairs (one pair for transmission, one pair for reception). To scale bandwidth, a link may aggregate multiple lanes denoted by xN , where N is any supported link width, such as 1, 2, 4, 8, 12, 16, 32, 64, or wider.

[0037] A differential pair can refer to two transmission paths, such as lines 316 and 317, to transmit differential signals. As an example, when line 316 toggles from a low voltage level to a high voltage level, i.e. a rising edge, line 317 drives from a high logic level to a low logic level, i.e. a falling edge. Differential signals potentially demonstrate better electrical characteristics, such as better signal integrity, i.e. cross-coupling, voltage overshoot/undershoot, ringing, among other example advantages. This allows for a better timing window, which enables faster transmission frequencies.

[0038] In one embodiment, a new High Performance Interconnect (HPI) is provided. HPI can include a next-generation cache-coherent, link-based interconnect. As one example, HPI may be utilized in high performance computing platforms, such as workstations or servers, including in systems where PCIe or another interconnect protocol is typically used to connect processors, accelerators, I/O devices, and the like. However, HPI is not so limited. Instead, HPI may be utilized in any of the systems or platforms described herein. Furthermore, the individual ideas developed may be applied to other interconnects and platforms, such as PCIe, MIPI, QPI, etc.

[0039] To support multiple devices, in one example implementation, HPI can include an Instruction Set Architecture (ISA) agnostic (i.e. HPI is able to be implemented in multiple different devices). In another scenario, HPI may also be utilized to connect high performance I/O devices, not just processors or accelerators. For example, a high performance PCIe device may be coupled to HPI through an appropriate translation bridge (i.e. HPI to PCIe). Moreover, the HPI links may be utilized by many HPI based devices, such as processors, in various ways (e.g. stars, rings, meshes, etc.). FIG. 4 illustrates example implementations of multiple potential multi-socket configurations. A two-socket configuration 405, as depicted, can include two HPI links; however, in other implementations, one HPI link may be utilized. For larger topologies, any configuration may be utilized as long as an identifier (ID) is assignable and there is some form of virtual path, among other additional or substitute features. As shown, in one example, a four socket configuration 410 has an HPI link from each processor to another. But in the eight socket implementation shown in configuration 415, not every socket is directly connected to each other through an HPI link. However, if a virtual path or channel exists between the processors, the configuration is supported. A range of supported processors includes 2-32 in a native domain. Higher numbers of processors may be reached through use of multiple domains or other interconnects between node controllers, among other examples.

[0040] The HPI architecture includes a definition of a layered protocol architecture, including in some examples, protocol layers (coherent, non-coherent, and, optionally, other memory based protocols), a routing layer, a link layer, and a physical layer including associated I/O logic. Furthermore, HPI can further include enhancements related to power managers (such as power control units (PCUs)), design for test and debug (DFT), fault handling, registers, security, among other examples. FIG. 5 illustrates an embodiment of an example HPI layered protocol stack. In some implementations, at least some of the layers illustrated in FIG. 5 may be optional. Each layer deals with its own level of granularity or quantum of information (the protocol layer 505a,b with packets 530, link layer 510a,b with flits 535, and physical layer 505a,b with phits 540). Note that a packet, in some embodiments, may include partial flits, a single flit, or multiple flits based on the implementation.

[0041] As a first example, a width of a phit 540 includes a 1 to 1 mapping of link width to bits (e.g. 20 bit link width includes a phit of 20 bits, etc.). Flits may have a greater size, such as 184, 192, or 200 bits. Note that if phit 540 is 20 bits wide and the size of flit 535 is 184 bits then it takes a fractional number of phits 540 to transmit one flit 535 (e.g. 9.2 phits at 20 bits to transmit an 184 bit flit 535 or 9.6 at 20 bits to transmit a 192 bit flit, among other

examples). Note that widths of the fundamental link at the physical layer may vary. For example, the number of lanes per direction may include 2, 4, 6, 8, 10, 12, 14, 16, 18, 20, 22, 24, etc. In one embodiment, link layer 510a,b is capable of embedding multiple pieces of different transactions in a single flit, and one or multiple headers (e.g. 1, 2, 3, 4) may be embedded within the flit. In one example, HPI splits the headers into corresponding slots to enable multiple messages in the flit destined for different nodes.

[0042] Physical layer 505a,b, in one embodiment, can be responsible for the fast transfer of information on the physical medium (electrical or optical etc.). The physical link can be point-to-point between two Link layer entities, such as layer 505a and 505b. The Link layer 510a,b can abstract the Physical layer 505a,b from the upper layers and provides the capability to reliably transfer data (as well as requests) and manage flow control between two directly connected entities. The Link Layer can also be responsible for virtualizing the physical channel into multiple virtual channels and message classes. The Protocol layer 520a,b relies on the Link layer 510a,b to map protocol messages into the appropriate message classes and virtual channels before handing them to the Physical layer 505a,b for transfer across the physical links. Link layer 510a,b may support multiple messages, such as a request, snoop, response, writeback, non-coherent data, among other examples.

[0043] The Physical layer 505a,b (or PHY) of HPI can be implemented above the electrical layer (i.e. electrical conductors connecting two components) and below the link layer 510a,b, as illustrated in FIG. 5. The Physical layer and corresponding logic can reside on each agent and connects the link layers on two agents (A and B) separated from each other (e.g. on devices on either side of a link). The local and remote electrical layers are connected by physical media (e.g. wires, conductors, optical, etc.). The Physical layer 505a,b, in one embodiment, has two major phases, initialization and operation. During initialization, the connection is opaque to the link layer and signaling may involve a combination of timed states and handshake events. During operation, the connection is transparent to the link layer and signaling is at a speed, with all lanes operating together as a single link. During the operation phase, the Physical layer transports flits from agent A to agent B and from agent B to agent A. The connection is also referred to as a link and abstracts some physical aspects including media, width and speed from the link layers while exchanging flits and control/status of current configuration (e.g. width) with the link layer. The initialization phase includes minor phases e.g. Polling, Configuration. The operation phase also includes minor phases (e.g. link power management states).

[0044] In one embodiment, Link layer 510a,b can be implemented so as to provide reliable data transfer between two protocol or routing entities. The Link layer can abstract Physical layer 505a,b from the Protocol layer 520a,b, and can be responsible for the flow control between two protocol agents (A, B), and provide virtual channel services to the Protocol layer (Message Classes) and Routing layer (Virtual Networks). The interface between the Protocol layer 520a,b and the Link Layer 510a,b can typically be at the packet level. In one embodiment, the smallest transfer unit at the Link Layer is referred to as a flit which a specified number of bits, such as 192 bits or some other denomination. The Link Layer 510a,b relies on the Physical layer 505a,b to frame the Physical layer's 505a,b unit of transfer (phit) into the Link Layer's 510a,b unit of transfer (flit). In addition, the Link Layer 510a,b may be logically broken into two parts, a sender and a receiver. A sender/receiver pair on one entity may be connected to a receiver/sender pair on another entity. Flow Control is often performed on both a flit and a packet basis. Error detection and correction is also potentially performed on a flit level basis.

[0045] In one embodiment, Routing layer 515a,b can provide a flexible and distributed method to route HPI transactions from a source to a destination. The scheme is flexible since routing algorithms for multiple topologies may be specified through programmable routing tables at each router (the programming in one embodiment is performed by firmware, software, or a combination thereof). The routing functionality may be distributed; the routing may be done through a series of routing steps, with each routing step being defined through a lookup of a table at either the source, intermediate, or destination routers. The lookup at a source may be used to inject a HPI packet into the HPI fabric. The lookup at an intermediate router may be used to route an HPI packet from an input port to an output port. The lookup at a destination port may be used to target the destination HPI protocol agent. Note that the Routing layer, in some implementations, can be thin since the routing tables, and, hence the routing algorithms, are not specifically defined by specification. This allows for flexibility and a variety of usage models, including flexible platform architectural topologies to be defined by the system implementation. The Routing layer 515a,b relies on the Link layer 510a,b for providing the use of up to three (or more) virtual networks (VNs) — in one example, two deadlock-free VNs, VN0 and VN1 with several message classes defined in each virtual network. A shared adaptive virtual network (VNA) may be defined in the Link layer, but this adaptive network may not be exposed directly in routing concepts, since each message class and virtual network may have dedicated resources and guaranteed forward progress, among other features and examples.

[0046] In one embodiment, HPI can include a Coherence Protocol layer 520a,b is support agents caching lines of data from memory. An agent wishing to cache memory data may use the coherence protocol to read the line of data to load into its cache. An agent wishing to modify a line of data in its cache may use the coherence protocol to acquire ownership of the line before modifying the data. After modifying a line, an agent may follow protocol requirements of keeping it in its cache until it either writes the line back to memory or includes the line in a response to an external request. Lastly, an agent may fulfill external requests to invalidate a line in its cache. The protocol ensures coherency of the data by dictating the rules all caching agents may follow. It also provides the means for agents without caches to coherently read and write memory data.

[0047] Two conditions may be enforced to support transactions utilizing the HPI Coherence Protocol. First, the protocol can maintain data consistency, as an example, on a per-address basis, among data in agents' caches and between those data and the data in memory. Informally, data consistency may refer to each valid line of data in an agent's cache representing a most up-to-date value of the data and data transmitted in a coherence protocol packet can represent the most up-to-date value of the data at the time it was sent. When no valid copy of the data exists in caches or in transmission, the protocol may ensure the most up-to-date value of the data resides in memory. Second, the protocol can provide well-defined commitment points for requests. Commitment points for reads may indicate when the data is usable; and for writes they may indicate when the written data is globally observable and will be loaded by subsequent reads. The protocol may support these commitment points for both cacheable and uncacheable (UC) requests in the coherent memory space.

[0048] The HPI Coherence Protocol also may ensure the forward progress of coherence requests made by an agent to an address in the coherent memory space. Certainly, transactions may eventually be satisfied and retired for proper system operation. The HPI Coherence Protocol, in some embodiments, may have no notion of retry for resolving resource allocation conflicts. Thus, the protocol itself may be defined to contain no circular resource dependencies, and implementations may take care in their designs not to introduce dependencies that can result in deadlocks. Additionally, the protocol may indicate where designs are able to provide fair access to protocol resources.

[0049] Logically, the HPI Coherence Protocol, in one embodiment, can include three items: coherence (or caching) agents, home agents, and the HPI interconnect fabric connecting the agents. Coherence agents and home agents can work together to achieve data consistency

by exchanging messages over the interconnect. The link layer 510a,b and its related description can provide the details of the interconnect fabric including how it adheres to the coherence protocol's requirements, discussed herein. (It may be noted that the division into coherence agents and home agents is for clarity. A design may contain multiple agents of both types within a socket or even combine agents behaviors into a single design unit, among other examples.)

[0050] In some implementations, HPI can utilize an embedded clock. A clock signal can be embedded in data transmitted using the interconnect. With the clock signal embedded in the data, distinct and dedicated clock lanes can be omitted. This can be useful, for instance, as it can allow more pins of a device to be dedicated to data transfer, particularly in systems where space for pins is at a premium.

[0051] The Link layer can guarantee reliable data transfer between two protocol or routing entities. The Link layer can abstract the Physical layer from the Protocol layer, handle flow control between two protocol agents, and provide virtual channel services to the Protocol layer (Message Classes) and Routing layer (Virtual Networks).

[0052] In some implementations, the Link layer can deal with a fixed quantum of information, termed a flit. In one example, the flit can be defined to be 192 bits in length. However, any range of bits, such as 81-256 (or more) may be utilized in different variations. A large flit size, such as 192 bits, may include format, cyclic redundancy check (CRC), and other changes. For instance, a larger flit length can also permit the CRC field to be expanded (e.g., to 16 bits) to handle the larger flit payload. The number of phits or unit intervals (UI) (e.g., the time used to transfer a single bit or phit, etc.) to transfer a single flit can vary with link width. For instance, a 20 lane or bit link width can transfer a single 192 bit flit in 9.6 UI, while an 8 lane link width transfers the same flit in 24 UI, among other potential examples. The link layer crediting and protocol packetizing can also be based on a flit.

[0053] FIG. 6 illustrates a representation 600 of a generalized flit for an 8 lane link width. Each column of the representation 600 can symbolize a link lane and each row a respective UI. In some implementations, a single flit can be subdivided into two or more slots. Distinct messages or link layer headers can be included in each slot, allowing multiple distinct, and in some cases, independent messages corresponding to potentially different transactions to be sent in a single flit. Further, the multiple messages included in slots of a single flit may also be destined to different destination nodes, among other examples. For instance, the example of

FIG. 6 illustrates a flit format with three slots. The shaded portions can represent the portion of the flit included in a respective slot.

[0054] In the example of FIG. 6, three slots, Slots 0, 1, and 2, are provided. Slot 0 can be provided 72 bits of flit space, of which 22 bits are dedicated to message header fields and 50 bits to message payload space. Slot 1 can be provided with 70 bits of flit space, of which 20 bits are dedicated to message header fields and 50 bits to message payload space. The difference in message header field space between can be optimized to provide that certain message types will be designated for inclusion in Slot 0 (e.g., where more message header encoding is utilized). A third slot, Slot 2, can be provided that occupies substantially less space than Slots 0 and 1, in this case utilizing 18 bits of flit space. Slot 2 can be optimized to handle those messages, such as acknowledges, credit returns, and the like that do not utilize larger message payloads. Additionally, a floating payload field can be provided that allows an additional 11 bits to be alternatively applied to supplement the payload field of either Slot 0 or Slot 1.

[0055] Continuing with the specific example of FIG. 6, other fields can be global to a flit (i.e., apply across the flit and not to a particular slot). For instance, a header bit can be provided together with a 4-bit flit control field that can be used to designate such information as a virtual network of the flit, identify how the flit is to be encoded, among other examples. Additionally, error control functionality can be provided, such as through a 16-bit cyclic CRC field, among other potential examples.

[0056] A flit format can be defined so as to optimize throughput of messages on the Link layer. Some traditional protocols have utilized unslotted, smaller flits. For instance, in QPI an 80-bit flit was utilized. While the flit throughput of a larger (e.g., 192-bit flit) may be lower, message or packet throughput can be increased by optimizing use of the flit data. For instance, in some implementations of QPI, the entire 80-bit flit space was utilized regardless of the message size or type. By subdividing a larger flit into slots of predetermined lengths and fields, the 192 flit length can be optimized realizing higher efficiency even in instances when one or more of the available slots are sometimes unused. Indeed, Link layer traffic can be assumed to include many different types of messages and traffic, including messages and packets with varying header lengths and fields. The respective lengths and organization of slots defined in a flit can be defined so as to correspond with the statistical or expected frequency of various messages and the needs of these messages. For instance, two larger slots can be defined for every small slot, to accommodate an expected statistical frequency of

messaging using these larger message types and header lengths, among other example. Further, flexibility can also be provided to further accommodate the varied traffic, such as through a floating payload field, as in the example of FIG. 6. In some instances, a flit format can be fixed, including the bits dedicated to particular slots in the flit.

[0057] In the example of FIG. 6, a “Hdr” field can be provided for the flit generally and represent a header indication for the flit. In some instances, the Hdr field can indicate whether the flit is a header flit or a data flit. In data flits, the flit can still remain slotted, but omit or replace the use of certain fields with payload data. In some cases, data fields may include an opcode and payload data. In the case of header flits, a variety of header fields can be provided. In the example of FIG. 6, “Oc” fields can be provided for each slot, the Oc field representing an opcode. Similarly, one or more slots can have a corresponding “msg” field representing a message type of the corresponding packet to be included in the slot, provided the slot is designed to handle such packet types, etc. “DNID” fields can represent a Destination Node ID, a “TID” field can represent a transaction ID, a “RHTID” field can represent either a requestor node ID or a home tracker ID, among other potential fields. Further, one or more slots can be provided with payload fields. Additionally, a CRC field can be included within a flit to provide a CRC value for the flit, among other examples.

[0058] In some implementations, link width can vary during the life of the link. For instance, the Physical layer can transition between link width states, such as to and from a full or original lane width and a different or partial lane width. For example, in some implementations, a link can be initialized to transfer data over 20 lanes. Later, the link can transition to a partial width transmitting state where only 8 lanes are actively used, among many other potential examples. Such lane width transitions can be utilized, for instance, in connection with power management tasks governed by one or more power control units (PCU) among other examples.

[0059] As noted above, link width can influence flit throughput rate. FIG. 7 is a representation of an example 192-bit flit sent over an 8 lane link, resulting in throughput of the flit at 24UI. Further, as shown in the example of FIG. 7, bits of the flit can be sent out of order in some instances, for example, to send more time-sensitive fields earlier in the transfer (e.g., flit type fields (e.g., data or header flit), opcodes, etc.), preserve or facilitate particular error detection or other functionality embodied in the flit, among other examples. For instance, in the example of FIG. 7, bits 191, 167, 143, 119, 95, 71, 47, and 23 are sent in parallel on lanes L7 through L0 during a first UI (i.e., UI0) of transfer, while bits 168, 144, 120, 96, 72, 48, 24,

and 0 are sent during the 24th (or final) UI of the flit transfer (i.e., UI23). It should be appreciated that other ordering schemes, flit lengths, lane widths, etc. can be utilized in other implementations and examples.

[0060] In some instances, the length of the flit can be a multiple of the number of active lanes. In such instances, the flit can be transmitted evenly on all active lanes and transfer of the flit can end substantially simultaneously at a clean (i.e., non-overlapping) boundary. For example, as shown in the representation of FIG. 8, bits of a flit can be considered to be transmitted in consecutive groupings of 4 bits, or “nibbles.” In this example, a 192 bit flit is to be transferred over an 8 lane link. As 192 is a multiple of 8, the entire flit can be cleanly transferred over the 8 lane link in 24 UI. In other instances, the flit width may not be a multiple of the number of active lanes. For instance, FIG. 9 shows another representation of an example 192 bit transferred over 20 lanes. As 192 is not evenly divisible by 20, transfer of the full flit would require a non-integer number of intervals (e.g., 9.6 UI). In such cases, rather than wasting “extra” lanes not utilized during the 10th UI of transfer, a second overlapping flit can be transferred with the final bits of a preceding flit. Such overlapping, or swizzling, of the flits can result in jagged flit boundaries and flit bits sent out of order in some implementations. The pattern utilized for the transfer can be configured to allow more time-sensitive fields of the flit to be transferred earlier in the flit, preservation of error detection and correction, among other considerations. Logic can be provided in one or both of the Physical and Link layers to transfer flit bits according to such patterns and dynamically change between patterns based on the current link width. Further logic can be provided to re-order and re-construct flits from such swizzled or ordered bit streams, among other examples.

[0061] In some implementations, flits can be characterized as header flits (e.g., bearing packet header data) or data flits (e.g., bearing packet payload data). Returning to FIG. 6, a flit format can be defined that includes three (3) distinct slots (e.g., 0, 1, and 2), allowing up to three headers to be transferred in a single flit (e.g., one header in each slot). Accordingly, each slot can have both control fields and a payload field. In addition to these, payload fields can be defined for each header (and slot). Further, a floating payload field can be defined that can be flexibly used as extra payload length for two or more of the slots (e.g., by either slot 0 or slot 1), based on the header types in these slots. The floating field can enable, in one implementation, 11 extra bits of payload for either Slot 0 or Slot 1. Note in implementations defining a larger flit more floating bits may be used and in smaller flits less floating bits may be provided.

[0062] In some implementations, by allowing a field to float between the two slots, extra bits can be provided as needed for certain messages while still staying within a predefined flit length (e.g., 192 bits) and maximizing the utilization of the bandwidth. Turning to the examples of FIG. 10, two instances 1005, 1010 of an example 192-bit flit are shown on an 8 lane data link. In one instance, a flit (e.g., 1005) can include three slots, Slots 0, 1, and 2. Each of Slots 0 and 1 can include 50-bit payload fields. The floating field can be provided to alternatively extend the payload field of the either Slot 0 or Slot 1 by the field length (e.g., 11 bits) of the floating field. The use of a floating field can further extend the efficiency gains provided through a defined, multi-slot flit format. The sizing of the slots within the flit, and the types of messages that can be placed in each slot, can potentially provide increased bandwidth even with a reduced flit rate.

[0063] In the particular example of FIG. 6, the messages that can use Slots 1 and 2 can be optimized, reducing the number of bits to be set aside to encode these slots' opcodes. When a header having more bits than Slot 0 can provide enters the Link layer, slotting algorithms can be provided to allow it to take over Slot 1 payload bits for additional space. Special control (e.g. LLCTRL) flits may also be provided that consume all three slots worth of bits for their needs. Slotting algorithms may also exist to allow individual slots to be utilized while other slots carry no information, for cases where the link is partially busy.

[0064] In the particular example of FIG. 10, example use of a floating flit field is shown. For instance, in the case of Standard Address Snoop (SA-S) Headers, only a single SA-S message (and header) may be permitted to be sent in the same flit (e.g., to prevent conflicts or where the SA-S payload utilizes a larger than 50-bit payload, etc.). Consequently, in such examples, a SA-S may only be sent in either Slot 0 or Slot 1 of the same flit in such instances. In the example of flit 1005, an SA-S header is included in Slot 0 and is to make use of the floating field. Consequently, in the example of flit 1005, the use of the floating field is dedicated to extend the payload of Slot 0's payload. In another example, of flit 1010, the SA-S header is to occupy Slot 1. In the example of flit 1010, the floating field is instead dedicated to extend the payload of Slot 1. Other potential examples can also make use of the flexibility provided through a floating payload field of a slotted flit utilizing principles illustrated in the particular examples of FIGS. 6 and 10.

[0065] In one embodiment, such as that illustrated in connection with FIG. 6, two slots, Slot 0 and 1, can be defined as having equally sized payload fields, while Slot 2 has a much smaller payload field for use by a particular subset of headers that lack the use of such larger

payload fields, for instance. Further, in one example, Slot 1 and 2 control fields may not carry full Message Class encodings (unlike Slot 0), and Slot 2 may not carry a full opcode encoding, among other potential implementations.

[0066] As noted above, in some implementations, Slots 1 and 2 may not carry full Message Class encodings, as not all bits are utilized due to slotting restrictions. Slot 1 can carries a Message Class bit 0. Here, request (REQ) and snoop (SNP) packets are allowed. In this implementation, REQ and SNP Message Class encodings are differentiated by bit 0. As a result, if a designer wanted to allow different message classes in partial message class field, they could either select a different bit position (i.e. an upper bit that differentiates two different types of messages) or assign different message types to the lower order bit. However, here the upper two bits are implied as 0's with the lower bit distinguishing between a REQ and a SNP. In this example, Slot 2 carries no Message Class bits, as only response (RSP) (encoding 2) packets are allowed in. Therefore, the Message Class encoding for Slot 2 is a RSP-2. Slot 2 can also carry a partial opcode. As above, one or more of the opcode bits can be assumed to be 0. As a result, partial message class fields and partial operation code fields may be utilized that define a subset of messages and op codes that may be utilized. Note that multiple sets of opcodes and messages classes may be defined. Here, if a lower order bit of the message class is used, then a subset of message types (i.e. MSG type 1/MSG type 2) is available. However, if 2 bits are used, then a larger subset is provided (e.g. Message Type 1/Message Type 2/Message Type 3/Message Type 4), among other examples.

[0067] Message class encodings can correspond to particular header types to be included (or to utilize) one or more defined slots in a flit. For instance, a header may have multiple sizes. In one example, a three slot flit can be defined to support potentially four sizes of header, based on header type. Table 1 includes an exemplary listing of potential header formats and associated sizes:

TABLE 1

Header Format	Header Size	Description
SA	Single Slot	Request
SA-S	Single Slot	Snoops (incorporates floating payload field)
SA-D	Single Slot	Data header
SR-U	Small Slot	Completion without data

SR-O	Single Slot	Ordering
SR-C	Single Slot	Conflict resolution
SR-D	Single Slot	Data header
PW	Dual Slot	Partial write
PR	Dual Slot	Partial read
P2P	Dual Slot	Peer-to-peer
NCM	Dual Slot	Non-coherent messaging
Slot-NULL	Single Slot (or Opcode only)	Control flit
LLCRD	Small Slot	Control flit
LLCTRL	Full Flit	Control flit

[0068] Small (or single) slot headers can be for those message small enough to fit in Slot 2, and that don't have protocol ordering requirements forcing them into Slot 0. A small slot header can also be placed in Slot 0, if the slotting restrictions for the flit call for it. The single slot header can be for those messages with payload that can fit in Slot 0 or Slot 1. Some single slot headers may also make use of the floating payload field. For instance, Standard Address Snoop (SA-S) Headers, in one embodiment, may not be sent in both slot 0 and slot 1 of the same flit in the example where only one HTID or floating field exists. Certain single slot headers may use Slot 0 based on protocol ordering requirements. The dual slot header can be for those messages large enough that they are to consume both the Slot 0 and Slot 1 payload fields, in addition to the floating payload field, among other examples.

[0069] A slot NULL opcode may include a special opcode that can be used, in one example, in either Slot 0 or Slot 1. For Slot 0, Slot_NULL may be used when the link layer has no header to transmit in Slot 0, but it does have a header to transmit in Slot 1 or 2. When Slot_NULL is used in Slot 0, the Slot 0 payload is considered reserved (RSVD), among other examples. In some implementations, Slot_NULL can be utilized in Slot 1 potentially under two conditions. First, when Slot 0 is encoding a dual slot or special control header, and thus consuming the Slot 1 payload. In such instances, the Slot 1 opcode can be set to Slot_NULL. The second condition is when the link layer has nothing to send in Slot 1, but does have a valid Single Slot header for Slot 0 or Small Slot Header for Slot 2. Under this condition, the Slot 1 opcode can be set to Slot_NULL and the Slot 1 payload can be considered Reserved, among other potential examples.

[0070] In some implementations, the small Slot 2, may include a reduced number of opcode bits. When the link layer has nothing to send in Slot 2, it may send an “Implicit NULL” by encoding a specific opcode, such as a link layer credit opcode and setting the Slot 2 payload field to all zeros. The receiver of this Slot 2 encoding can process it as a link layer credit message (except in the case of the special control flits), but the all zeros encoding will have no effect on the Credit and Acknowledge state. In the case of special control flits, because they can consume the entire flit, the Slot 2 payload can be considered RSVD and the Implicit NULL will be ignored. Where the link layer has nothing to send in any of the three slots and the CRD/ACK fields, the link layer may transmit a special control null message, among other examples.

[0071] Slotting restrictions can be defined for one or more of the defined slots of a flit. In one embodiment, dual slot headers may only have their Message Class and Opcode placed in Slot 0. When Slot 0 contains a Dual Slot Header, Slot 1 may encode a Slot_NULL opcode, as the Slot 1 Payload field will be consumed by the Slot 0 header. When Slot 0 contains a Slot_NULL, single slot, or small slot header, Slots 1 and 2 may both encode a non-NULL header. Only small slot headers are allowed in Slot 2 in this particular example (e.g., illustrated in FIG. 6). When both Slot 0 and Slot 1 contain single slot headers, one may be of a type that consumes the floating payload field. If neither Slot 0 or Slot 1 contain a header type that consumes the floating payload field, the field may be considered RSVD.

[0072] Additionally, in some implementations, the Link layer can utilize multiple different types of virtual network or virtual channel credits. In one example, pooled virtual network adaptive (VNA) credits can be supported and a VNA field can be provided. In one example implementation, when the VNA field indicates a non-VNA flit (e.g., a flit that utilizes a different credit pool), the header may be designated to be placed in Slot 0. Further, the Slot 2 opcode may include a Slot_2 credit in this case. Further, when Slot 0 encodes a special control Header, both Slot 1 and Slot 2 control fields may be set to fixed values, and no headers may be placed in these slots, among other potential implementations.

[0073] As noted above, in header flits, a variety of different fields can be provided to be incorporated in corresponding flit slots, such as illustrated in the particular example of FIG. 6. Note that the fields illustrated and described are provided by way of example and additional or substitute fields can also be incorporated. Indeed, some of the fields described may be optional and be omitted in some implementations, among other examples.

[0074] In one example, a message class (MC) field can be provided, as well as other fields. In some examples, the Protocol layer can use the Message Class field to define the Protocol Class which also acts as the Major Opcode field. The Link layer can use the Message Class field as part of the virtual channel (VC) definition. Some Protocol Classes/VC can use multiple Message Class encodings due to the number of opcodes that are to be encoded, among other examples. For instance, Requests (REQ), Snoops (SNP), Response (RSP), writeback, non-coherent bypass, and non-coherent standard types can be supported. If each type encoded sixteen operations, then there would be an opcode space of 96 operations. And if another mode bit or other opcode space was defined for each type, then another 96 operations could be provided; and so on.

[0075] In one example, an Opcode field can additionally be provided. The Protocol layer may use the opcode in conjunction with the Message Class to form a complete opcode (i.e. define the message class type and the operation within). As an example, the same opcode with a REQ message type may define a first request operation, while the same opcode with a SNP message class may define a second, different SNP operation, among other examples. The Link Layer may use the opcode to distinguish, for instance, between a Home Agent target or a Caching Agent target for packets when a Home Agent and a Caching Agent share the same NodeID. Additionally, the Link Layer may also use the opcode to determine packet size, among other potential uses.

[0076] As noted above, flit headers can further include a Virtual network Adaptive (VNA) field. In one example, when a VNA field is set to a first value, the field can indicate that the flit is using VNA credits. When set to a second value, the flit is using VN0 or VN1 credits, among other potential implementations. In one embodiment, a value may indicate the flit is a single slot flit and slots 1 and 2 codes can be defined as NULL.

[0077] A Virtual Network (VN) field can also be provided and indicate for a flit if the header(s) in the flit are utilizing a particular virtual network, such as a virtual network VN0 or VN1. This may be used for both crediting purposes and to indicate which virtual network a message should drain to if using VNA. If one VN bit is provided for the entire flit, any VNA flit that contains multiple headers can ensure that all of them are draining to VN0 or all of them are draining to VN1. Alternatively, multiple VN bits may be provided. For non VNA flits, only Slot 0 may be allowed to have a non-control opcode, so the VN may indicate that header's network.

[0078] In some implementations, slots in a flit can be used for small payload messages such as credit returns, ACKs, NAKs, among others. In one example, a channel field can be provided that can be encoded for use in credit returns. This encoding, in combination with the Virtual Network field, may provide the Virtual Channel that a credit return maps to. Where a Message Class has multiple encodings, they may all map to a single Channel value for crediting. When the credit return type is VNA, the Channel value can be ignored. Use of RSVD encodings may be treated as an error by the receiving component. Table 2 includes examples of different Channel options that can be encoded. Note that any combination of bits (or bits representing a hexadecimal value) may be utilized. As an example, a lower order of 3 bits can be used for encoding.

TABLE 2

Channel
REQ: Request
SNP: Snoop
RSP: Response
RSVD: Reserved
WB: Write back
NCB: Non-coherent bypass
NCS: Non-coherent standard

[0079] Acknowledgement, or ACK, fields can also be provided as header fields to be included in a flit slot. An ACK field may be used by the Link layer to communicate from a receiver to a sender error free receipt of flits. ACK having a first value indicates that a number of flits, such as 4, 8, or 12, have been received without error. When a sender receives an ACK it may deallocate the corresponding flits from the Link Layer Retry Queue. Ack and Ack fields can be used in credit return control flits (e.g., LLCRD), with the total number of Acknowledges being returned determined by creating the full acknowledge return value (Acknowledge first portion, ACK, Acknowledge second portion), among other examples.

[0080] As noted above, a Header indication bit (Hdr) can also be provided in some implementations and can be used for one or more purposes. For instance, a Hdr packet can identify whether the packet is a header or data flit, can indicate that the flit is the start of a new packet, as well as indicate the start of an interleaved Link Layer Control flit. The Hdr can be set for the first flit of all packets. Further, an Address field can be provided to identify a global

system address. All coherent transactions may be a number of byte aligned and may return the number of bytes of data, eliminating the need for some portion of the Address bits (e.g. at 64 bytes, the lower 6 bits may be omitted). For certain other packets, a full byte level address is to be utilized. A Length field can be provided in some examples to indicate a length of the requested data in bytes for any transaction that is doing a partial read. The partial read specifies the offset (e.g. the lower portion of the address bits omitted above) and the Length. Valid lengths are 0 to the number of bytes that the transactions are aligned to less one, among other examples.

[0081] Additional fields can be included. A Byte Enable field can be provided in some instances to indicate the valid bytes for any transaction doing a partial write. A Byte Enable field may have any number 0 to the number of bytes that the transactions are aligned to less one. A Request TID (RTID) field can be used to uniquely identify the different requests from a single Protocol Agent. A Home tracker ID (HTID) field can be used in Snoop packets and Snoop Response packets to indicate the Home Tracker ID of the transaction the snoop and its response are to be associated with. An RHTID field can also be provided in some implementations and flexibly embody an RTID or an HTID, depending on the opcode. For instance, for a snoop, RHTID can be interpreted as RTID, as snoops have an explicit HTID field. For response packets, on the other hand, targeting a home agent, RHTID can be interpreted as HTID. Additionally, for response packets targeting a cache agent, RHTID can be interpreted as RTID for opcodes except FwdCnfltO, among other examples. In some implementations, other message types can default to being interpreted as RTID.

[0082] In some implementations, additional fields can be provided such as a Destination Node ID (DNID) field, Requestor Node ID (RNID) field, Conflict Node ID (CNID) field, and Source Node ID (SNID) field. The DNID can identify the destination of a packet. It can be supplied by the Protocol Layer and used by the Link and Routing layers to guide packets to their destinations. The RNID field can identify the original requester/initiator of a transaction and can be supplied by the Protocol Layer. The CNID can be used in RspCnflt packets to indicate the node ID (NID) of the agent where the snoop experienced a conflict and the FwdCnfltO should be sent. A SNID field can be used in SR-D packets to indicate the Node ID of the agent transmitting the SR-D packet.

[0083] Additionally, a Prior Cache Line State (PCLS) field can be provided to encode the state of the cache line where it was found at either a peer caching agent or in a home node. For example, if the cache line was supplied by a peer node in the F state, the field should be set

to a first value. If the cache line was sent by a home node, the home node should set the field to reflect either the I state or S state depending on the snoop responses it received. If an agent does not support this field it should always be encoded as a default value. Note the PCLS field may be used for performance monitoring/tuning. A Non-Coherent Protected field can denote whether a request is to Normal or Protected space. See the table below for the encodings.

[0084] In some implementations, HPI Link layer can support a multi-slot flit with explicit fields, such as those described above, as well as implicit fields. For instance, slot message encoding and opcodes can be regarded as implicit. For instance, Slots 1 and 2 may not carry full Message Class encodings, as not all bits are required, in some instances, due to slotting restrictions. Slot 1 carries only Message Class bit 0, and only REQ and SNP packets may be allowed in this slot. REQ and SNP Message Class encodings can be differentiated by bit 0, and the upper two bits can be implied as 0's. Slot 2 may not carry Message Class bits, as only RSP (encoding 2) packets are allowed in this Slot. Therefore the Message Class encoding for Slot 2 may be RSP-2. Slot 2 can also only carry a portion of an opcode, with a second portion of an opcode being assumed to be a default value. This means that RSP-2 packets with the second portion holding the default value are allowed in Slot 2. Further, the Complete opcode field, in one embodiment, can be created by combining the full Message Class with the full Opcode field, forming a Complete Opcode.

[0085] Additional examples of implicit fields can include packet length, which can be implied by the opcode. Further, the globally Unique Transaction ID (UTID) may be formed by combining Requester NodeID with Requester Transaction ID. Note, that there may be an overlap in the RTID space between P2P and non-P2P transactions. For instance, the globally P2P Unique Transaction ID (P2PUTID) may be formed by combining Requester NodeID with Requester Transaction ID.

[0086] In some implementations, such as that illustrated in the examples of FIG. 6, the structure of the flit can permit Transaction IDs (TIDs) that utilize 11 bits of flit space. As a result, pre-allocation and the enabling of distributed home agents may be removed. Furthermore, use of 11 bits, in some implementations, allows for the TID to be used without having use for an extended TID mode.

[0087] Link layer logic can be provided on each agent on each side of a link. A transmitter of an agent or device can receive data from higher layers (e.g., a Protocol or Routing layer) and generate one or more flits to transfer the data to a receiver of a remote

agent. The agent can generate a flit with two or more slots. In some instances, the agent can attempt to combine multiple messages or packets within a single flit utilizing the defined slots.

[0088] Link layer logic can include, in some implementations, dedicated paths corresponding to each defined slot. The paths can be embodied in either or both hardware and software. A receiver of an agent can receive a flit (as re-constructed using the Physical layer) and Link layer logic can identify each of the slots and process the slots using each slot's respective path. The Link layer can process the flits, and the data included in each slot, according to one or more encoded fields of the flit, such as a control field, header field, CRC field, etc.

[0089] In one illustrative example, a transmitter can receive a write request associated with a first transaction, a snoop request associated with another second transaction, and one or more acknowledges or credit returns that it can send to (or through) another device. The transmitter can send a single flit to the other device over a serial data link of an interconnect, the single flit to include headers for each of the write request, the snoop, and an acknowledge (e.g., a completion), each header occupying a respective flit slot (such as in the 3-slot implementation illustrated in the example of FIG. 6). The transmitter can buffer data it receives and identify opportunities to send multiple messages in a single flit. The receiver can receive the compiled flit and process each slot in parallel to identify and process each of the three messages, among many other potential examples.

[0090] In some implementations, multiple headers can be included in a multi-slot flit so as to send multiple messages using a single flit. In some examples, the respective headers can each relate to fully independent transactions. In some implementations, the flexibility of the flit can be constrained such that each flit only contains messages directed to a particular virtual network. Other implementations may abstain from such a condition.

[0091] In instances where slot messages are to all apply to a common one of a plurality of virtual networks, bits that would have traditionally been reserved for identification of a slot's respective virtual network can be dedicated to other uses, that in some implementations, further increases efficiency gains introduced by the flit format, among potentially other benefits. In one example, all slots in a multi-slot header flit may be aligned to a single virtual network such as only VNA, only VN0, or only VN1, etc. By enforcing this, per slot bits indicating virtual network can be removed. This increases the efficiency of flit bit utilization and potentially enables such other features, as expanding from 10 bit TIDs to 11 bit

TIDs, among other examples. Expanding to an 11 bit TID can, in some implementations, allow for the TID to be used without having use for an extended TID mode.

[0092] HPI can be incorporated in any variety of computing devices and systems, including mainframes, server systems, personal computers, mobile computers (such as tablets, smartphones, personal digital systems, etc.), smart appliances, gaming or entertainment consoles and set top boxes, among other examples. For instance, FIG. 11 illustrates an example computer system 1100 in accordance with some implementations. As shown in FIG. 11, multiprocessor system 1100 is a point-to-point interconnect system, and includes a first processor 1170 and a second processor 1180 coupled via a point-to-point interconnect 1150. Each of processors 1170 and 1180 may be some version of a processor. In one embodiment, 1152 and 1154 are part of a serial, point-to-point coherent interconnect fabric, such as a high-performance architecture. As a result, the invention may be implemented within the QPI architecture.

[0093] While shown with only two processors 1170, 1180, it is to be understood that the scope of the present invention is not so limited. In other embodiments, one or more additional processors may be present in a given processor.

[0094] Processors 1170 and 1180 are shown including integrated memory controller units 1172 and 1182, respectively. Processor 1170 also includes as part of its bus controller units point-to-point (P-P) interfaces 1176 and 1178; similarly, second processor 1180 includes P-P interfaces 1186 and 1188. Processors 1170, 1180 may exchange information via a point-to-point (P-P) interface 1150 using P-P interface circuits 1178, 1188. As shown in FIG. 11, IMCs 1172 and 1182 couple the processors to respective memories, namely a memory 1132 and a memory 1134, which may be portions of main memory locally attached to the respective processors.

[0095] Processors 1170, 1180 each exchange information with a chipset 1190 via individual P-P interfaces 1152, 1154 using point to point interface circuits 1176, 1194, 1186, 1198. Chipset 1190 also exchanges information with a high-performance graphics circuit 1138 via an interface circuit 1192 along a high-performance graphics interconnect 1139.

[0096] A shared cache (not shown) may be included in either processor or outside of both processors; yet connected with the processors via P-P interconnect, such that either or both processors' local cache information may be stored in the shared cache if a processor is placed into a low power mode.

[0097] Chipset 1190 may be coupled to a first bus 1116 via an interface 1196. In one embodiment, first bus 1116 may be a Peripheral Component Interconnect (PCI) bus, or a bus such as a PCI Express bus or another third generation I/O interconnect bus, although the scope of the present invention is not so limited.

[0098] As shown in FIG. 11, various I/O devices 1114 are coupled to first bus 1116, along with a bus bridge 1118 which couples first bus 1116 to a second bus 1120. In one embodiment, second bus 1120 includes a low pin count (LPC) bus. Various devices are coupled to second bus 1120 including, for example, a keyboard and/or mouse 1122, communication devices 1127 and a storage unit 1128 such as a disk drive or other mass storage device which often includes instructions/code and data 1130, in one embodiment. Further, an audio I/O 1124 is shown coupled to second bus 1120. Note that other architectures are possible, where the included components and interconnect architectures vary. For example, instead of the point-to-point architecture of FIG. 11, a system may implement a multi-drop bus or other such architecture.

[0100] While the present invention has been described with respect to a limited number of embodiments, those skilled in the art will appreciate numerous modifications and variations therefrom. It is intended that the appended claims cover all such modifications and variations as fall within the true spirit and scope of this present invention.

[0101] A design may go through various stages, from creation to simulation to fabrication. Data representing a design may represent the design in a number of manners. First, as is useful in simulations, the hardware may be represented using a hardware description language or another functional description language. Additionally, a circuit level model with logic and/or transistor gates may be produced at some stages of the design process. Furthermore, most designs, at some stage, reach a level of data representing the physical placement of various devices in the hardware model. In the case where conventional semiconductor fabrication techniques are used, the data representing the hardware model may be the data specifying the presence or absence of various features on different mask layers for masks used to produce the integrated circuit. In any representation of the design, the data may be stored in any form of a machine readable medium. A memory or a magnetic or optical storage such as a disc may be the machine readable medium to store information transmitted via optical or electrical wave modulated or otherwise generated to transmit such information. When an electrical carrier wave indicating or carrying the code or design is transmitted, to the extent that copying, buffering, or re-transmission of the electrical signal is performed, a new

copy is made. Thus, a communication provider or a network provider may store on a tangible, machine-readable medium, at least temporarily, an article, such as information encoded into a carrier wave, embodying techniques of embodiments of the present invention.

[0102] A module as used herein refers to any combination of hardware, software, and/or firmware. As an example, a module includes hardware, such as a micro-controller, associated with a non-transitory medium to store code adapted to be executed by the micro-controller. Therefore, reference to a module, in one embodiment, refers to the hardware, which is specifically configured to recognize and/or execute the code to be held on a non-transitory medium. Furthermore, in another embodiment, use of a module refers to the non-transitory medium including the code, which is specifically adapted to be executed by the microcontroller to perform predetermined operations. And as can be inferred, in yet another embodiment, the term module (in this example) may refer to the combination of the microcontroller and the non-transitory medium. Often module boundaries that are illustrated as separate commonly vary and potentially overlap. For example, a first and a second module may share hardware, software, firmware, or a combination thereof, while potentially retaining some independent hardware, software, or firmware. In one embodiment, use of the term logic includes hardware, such as transistors, registers, or other hardware, such as programmable logic devices.

[0103] Use of the phrase ‘configured to,’ in one embodiment, refers to arranging, putting together, manufacturing, offering to sell, importing and/or designing an apparatus, hardware, logic, or element to perform a designated or determined task. In this example, an apparatus or element thereof that is not operating is still ‘configured to’ perform a designated task if it is designed, coupled, and/or interconnected to perform said designated task. As a purely illustrative example, a logic gate may provide a 0 or a 1 during operation. But a logic gate ‘configured to’ provide an enable signal to a clock does not include every potential logic gate that may provide a 1 or 0. Instead, the logic gate is one coupled in some manner that during operation the 1 or 0 output is to enable the clock. Note once again that use of the term ‘configured to’ does not require operation, but instead focus on the latent state of an apparatus, hardware, and/or element, where in the latent state the apparatus, hardware, and/or element is designed to perform a particular task when the apparatus, hardware, and/or element is operating.

[0104] Furthermore, use of the phrases ‘to,’ ‘capable of/to,’ and or ‘operable to,’ in one embodiment, refers to some apparatus, logic, hardware, and/or element designed in such a way to enable use of the apparatus, logic, hardware, and/or element in a specified manner.

Note as above that use of to, capable to, or operable to, in one embodiment, refers to the latent state of an apparatus, logic, hardware, and/or element, where the apparatus, logic, hardware, and/or element is not operating but is designed in such a manner to enable use of an apparatus in a specified manner.

[0105] A value, as used herein, includes any known representation of a number, a state, a logical state, or a binary logical state. Often, the use of logic levels, logic values, or logical values is also referred to as 1's and 0's, which simply represents binary logic states. For example, a 1 refers to a high logic level and 0 refers to a low logic level. In one embodiment, a storage cell, such as a transistor or flash cell, may be capable of holding a single logical value or multiple logical values. However, other representations of values in computer systems have been used. For example the decimal number ten may also be represented as a binary value of 1010 and a hexadecimal letter A. Therefore, a value includes any representation of information capable of being held in a computer system.

[0106] Moreover, states may be represented by values or portions of values. As an example, a first value, such as a logical one, may represent a default or initial state, while a second value, such as a logical zero, may represent a non-default state. In addition, the terms reset and set, in one embodiment, refer to a default and an updated value or state, respectively. For example, a default value potentially includes a high logical value, i.e. reset, while an updated value potentially includes a low logical value, i.e. set. Note that any combination of values may be utilized to represent any number of states.

[0107] The embodiments of methods, hardware, software, firmware or code set forth above may be implemented via instructions or code stored on a machine-accessible, machine readable, computer accessible, or computer readable medium which are executable by a processing element. A non-transitory machine-accessible/readable medium includes any mechanism that provides (i.e., stores and/or transmits) information in a form readable by a machine, such as a computer or electronic system. For example, a non-transitory machine-accessible medium includes random-access memory (RAM), such as static RAM (SRAM) or dynamic RAM (DRAM); ROM; magnetic or optical storage medium; flash memory devices; electrical storage devices; optical storage devices; acoustical storage devices; other form of storage devices for holding information received from transitory (propagated) signals (e.g., carrier waves, infrared signals, digital signals); etc, which are to be distinguished from the non-transitory mediums that may receive information there from.

[0108] Instructions used to program logic to perform embodiments of the invention may be stored within a memory in the system, such as DRAM, cache, flash memory, or other storage. Furthermore, the instructions can be distributed via a network or by way of other computer readable media. Thus a machine-readable medium may include any mechanism for storing or transmitting information in a form readable by a machine (e.g., a computer), but is not limited to, floppy diskettes, optical disks, Compact Disc, Read-Only Memory (CD-ROMs), and magneto-optical disks, Read-Only Memory (ROMs), Random Access Memory (RAM), Erasable Programmable Read-Only Memory (EPROM), Electrically Erasable Programmable Read-Only Memory (EEPROM), magnetic or optical cards, flash memory, or a tangible, machine-readable storage used in the transmission of information over the Internet via electrical, optical, acoustical or other forms of propagated signals (e.g., carrier waves, infrared signals, digital signals, etc.). Accordingly, the computer-readable medium includes any type of tangible machine-readable medium suitable for storing or transmitting electronic instructions or information in a form readable by a machine (e.g., a computer).

[0109] The following examples pertain to embodiments in accordance with this Specification. One or more embodiments may provide an apparatus, a system, a machine readable storage, a machine readable medium, and a method to identify transaction data, generate a flit to include three or more slots and a floating field to be used as an extension of any one of two or more of the slots, and send the flit to transmit the flit.

[0110] In at least one example, I/O logic comprises a layered stack comprising physical layer logic, link layer logic, and protocol layer logic

[0111] In at least one example, the three or more slots consist of three defined slots.

[0112] In at least one example, the flit comprises 192 bits.

[0113] In at least one example, the first of the three slots comprises 72 bits, the second of the three slots comprises 70 bits, and third slot comprises 18 bits

[0114] In at least one example, the first slot and second slot each include a respective 50 bit payload field.

[0115] In at least one example, the floating field is to extend the payload field of either the first slot or the second slot by eleven bits.

[0116] In at least one example, the third slot is adapted to be encoded with one or more of acknowledgements and credit returns.

[0117] In at least one example, the flit further comprises a 16-bit cyclic redundancy check (CRC) field.

[0118] In at least one example, the flit further comprises an 11-bit transaction identifier (TID) field.

[0119] In at least one example, each slot is to include a header of a distinct message.

[0120] In at least one example, each message is associated with a respective transaction within a particular virtual network.

[0121] In at least one example, the flit further comprises a virtual network identifier to identify the particular virtual network.

[0122] In at least one example, wherein message headers associated with transactions in different virtual networks are to be included in distinct flits.

[0123] One or more embodiments may provide an apparatus, a system, a machine readable storage, a machine readable medium, and a method to receive a flit, wherein the flit is to include three or more slots to be included in the flit and a floating field to be used as an extension of any one of two or more of the slots, and process each slot to identify one or more headers relating to one or more transactions.

[0124] In at least one example, the one or more headers comprise three or more headers.

[0125] In at least one example, each of the headers corresponds to a respective message associated with a different, respective transaction.

[0126] In at least one example, each of the transactions is included in a particular virtual network.

[0127] In at least one example, it can be identified which of the first and second slots the floating field is to extend.

[0128] In at least one example, the third slot is adapted to be encoded with one or more of acknowledgements and credit returns.

[0129] In at least one example, the flit can be sent from a first device to a second device over the data link. The first second devices can include microprocessors, graphics accelerators, and other devices.

[0130] One or more examples can further provide a layered protocol stack including at least a link layer and a physical layer, the layered protocol stack configured to be coupled to a serial, differential link, wherein the layered protocol stack is further configured to transmit a 192-bit flit over the serial, differential link.

[0131] In at least one example, the 192-bit flit includes a 16 bit CRC

[0132] One or more examples can further provide a layered protocol stack including at least a link layer and a physical layer, the layered protocol stack configured to be coupled to a serial, differential link, wherein the layered protocol stack is further configured to transmit a flit over the serial, differential link, the flit to include an 11-bit transaction identifier field.

[0133] One or more examples can further provide a layered protocol stack including at least a link layer and a physical layer, the layered protocol stack configured to be coupled to a serial, differential link, wherein the layered protocol stack is further configured to assemble a header flit including a plurality of slots.

[0134] In at least one example, the plurality of payload slots include 3 slots.

[0135] In at least one example, the first and second slots of the 3 slots are equal in size and the third slot of the 3 slots is smaller than the first slot.

[0136] In at least one example, special control flits may consume all 3 slots.

[0137] In at least one example, the flit includes a 16 bit CRC.

[0138] Reference throughout this specification to “one embodiment” or “an embodiment” means that a particular feature, structure, or characteristic described in connection with the embodiment is included in at least one embodiment of the present invention. Thus, the appearances of the phrases “in one embodiment” or “in an embodiment” in various places throughout this specification are not necessarily all referring to the same embodiment. Furthermore, the particular features, structures, or characteristics may be combined in any suitable manner in one or more embodiments.

[0139] In the foregoing specification, a detailed description has been given with reference to specific exemplary embodiments. It will, however, be evident that various modifications and changes may be made thereto without departing from the broader spirit and scope of the invention as set forth in the appended claims. The specification and drawings are, accordingly, to be regarded in an illustrative sense rather than a restrictive sense. Furthermore, the foregoing use of embodiment and other exemplarily language does not necessarily refer to the same embodiment or the same example, but may refer to different and distinct embodiments, as well as potentially the same embodiment.

What is claimed is:

1. An apparatus comprising:
I/O logic to:
 identify transaction data;
 generate a flit to include three or more slots and a floating field to be used as an extension of any one of two or more of the slots; and
 transmit the flit.
2. The apparatus of Claim 1, wherein the I/O logic comprises a layered stack comprising physical layer logic, link layer logic, and protocol layer logic
3. The apparatus of Claim 1, wherein the three or more slots consist of three defined slots.
4. The apparatus of Claim 3, wherein the flit comprises 192 bits.
5. The apparatus of Claim 4, wherein the first of the three slots comprises 72 bits, the second of the three slots comprises 70 bits, and third slot comprises 18 bits
6. The apparatus of Claim 5, wherein the first slot and second slot each include a respective 50 bit payload field.
7. The apparatus of Claim 6, wherein the floating field is to extend the payload field of either the first slot or the second slot by eleven bits.
8. The apparatus of Claim 5, wherein the third slot is adapted to be encoded with one or more of acknowledgements and credit returns.
9. The apparatus of Claim 5, wherein the flit further comprises a 16-bit cyclic redundancy check (CRC) field.

10. The apparatus of Claim 5, wherein the flit further comprises an 11-bit transaction identifier (TID) field.
11. The apparatus of Claim 1, wherein each slot is to include a header of a distinct message.
12. The apparatus of Claim 11, wherein each message is associated with a respective transaction within a particular virtual network.
13. The apparatus of Claim 12, wherein the flit further comprises a virtual network identifier to identify the particular virtual network.
14. The apparatus of Claim 12, wherein message headers associated with transactions in different virtual networks are to be included in distinct flits.
15. An apparatus comprising:
I/O logic to:
receive a flit, wherein the flit is to include three or more slots to be included in the flit and a floating field to be used as an extension of any one of two or more of the slots; and
process each slot to identify one or more headers relating to one or more transactions.
16. The apparatus of Claim 15, wherein the one or more headers comprise three or more headers.
17. The apparatus of Claim 16, wherein each of the headers corresponds to a respective message associated with a different, respective transaction.
18. The apparatus of Claim 17, wherein each of the transactions is included in a particular virtual network.

19. The apparatus of Claim 15, wherein the three or more slots consist of three defined slots.
20. The apparatus of Claim 19, wherein the flit comprises 192 bits.
21. The apparatus of Claim 20, wherein the first of the three slots comprises 72 bits, the second of the three slots comprises 70 bits, and third slot comprises 18 bits.
22. The apparatus of Claim 21, wherein the floating field is to extend the payload field of either the first slot or the second slot.
23. The apparatus of Claim 22, wherein the I/O logic is further to identify which of the first and second slots the floating field is to extend.
24. The apparatus of Claim 21, wherein the third slot is adapted to be encoded with one or more of acknowledgements and credit returns.
25. A method comprising:
 - identifying transaction data to be sent to a device on a serial data link;
 - generating, from the transaction data, one or more flits, each flit including three or more slots and a floating field to be used as an extension of any one of two or more of the slots; and
 - sending the flit to the device over the serial data link.
26. The method of Claim 25, wherein the flit comprises 192 bits, a first of the slots comprises 72 bits, a second of the slots comprises 70 bits, and a third of the slots comprises 18 bits.
27. The method of Claim 25, wherein the floating field comprises eleven bits.

28. A system comprising:
a first device; and
a second device communicatively coupled to the first device using a serial data link, the second device including a link layer module executed by at least one processor to:
identify transaction data to be sent to the first device on the data link;
generate, from the transaction data, one or more flits, wherein each flit is to include three or more slots to be included in the flit and a floating field to be used as an extension of any one of two or more of the slots; and
send the flit to the first device over the data link.
29. The system of Claim 28, wherein the first device comprises a microprocessor.
30. The system of Claim 29, wherein the second device comprises a second microprocessor.
31. The system of Claim 29, wherein the second device comprises a graphics accelerator.
32. The system of Claim 28, wherein the first device includes link layer logic to:
receive the flits over the data link; and
process each slot to identify one or more headers relating to one or more transactions.
33. An apparatus comprising
a layered protocol stack including at least a link layer and a physical layer, the layered protocol stack configured to be coupled to a serial, differential link, wherein the layered protocol stack is further configured to transmit a 192-bit flit over the serial, differential link.
34. The apparatus of Claim 33, wherein the 192-bit flit includes a 16 bit CRC
35. An apparatus comprising
a layered protocol stack including at least a link layer and a physical layer, the layered protocol stack configured to be coupled to a serial, differential link, wherein the layered protocol stack is further configured to transmit a flit over the serial, differential link, the flit to include an 11-bit transaction identifier field.

36. An apparatus comprising a layered protocol stack including at least a link layer and a physical layer, the layered protocol stack configured to be coupled to a serial, differential link, wherein the layered protocol stack is further configured to assemble a header flit including a plurality of slots.

37. The apparatus of Claim 36, wherein the plurality of payload slots include 3 slots.

38. The apparatus of any of Claims 36-37, wherein the first and second slots of the 3 slots are equal in size and the third slot of the 3 slots is smaller than the first slot.

39. The apparatus of any of Claims 36-38, wherein special control flits may consume all 3 slots.

40. The apparatus of any of Claims 36-39, wherein the flit includes a 16 bit CRC.

1 / 11

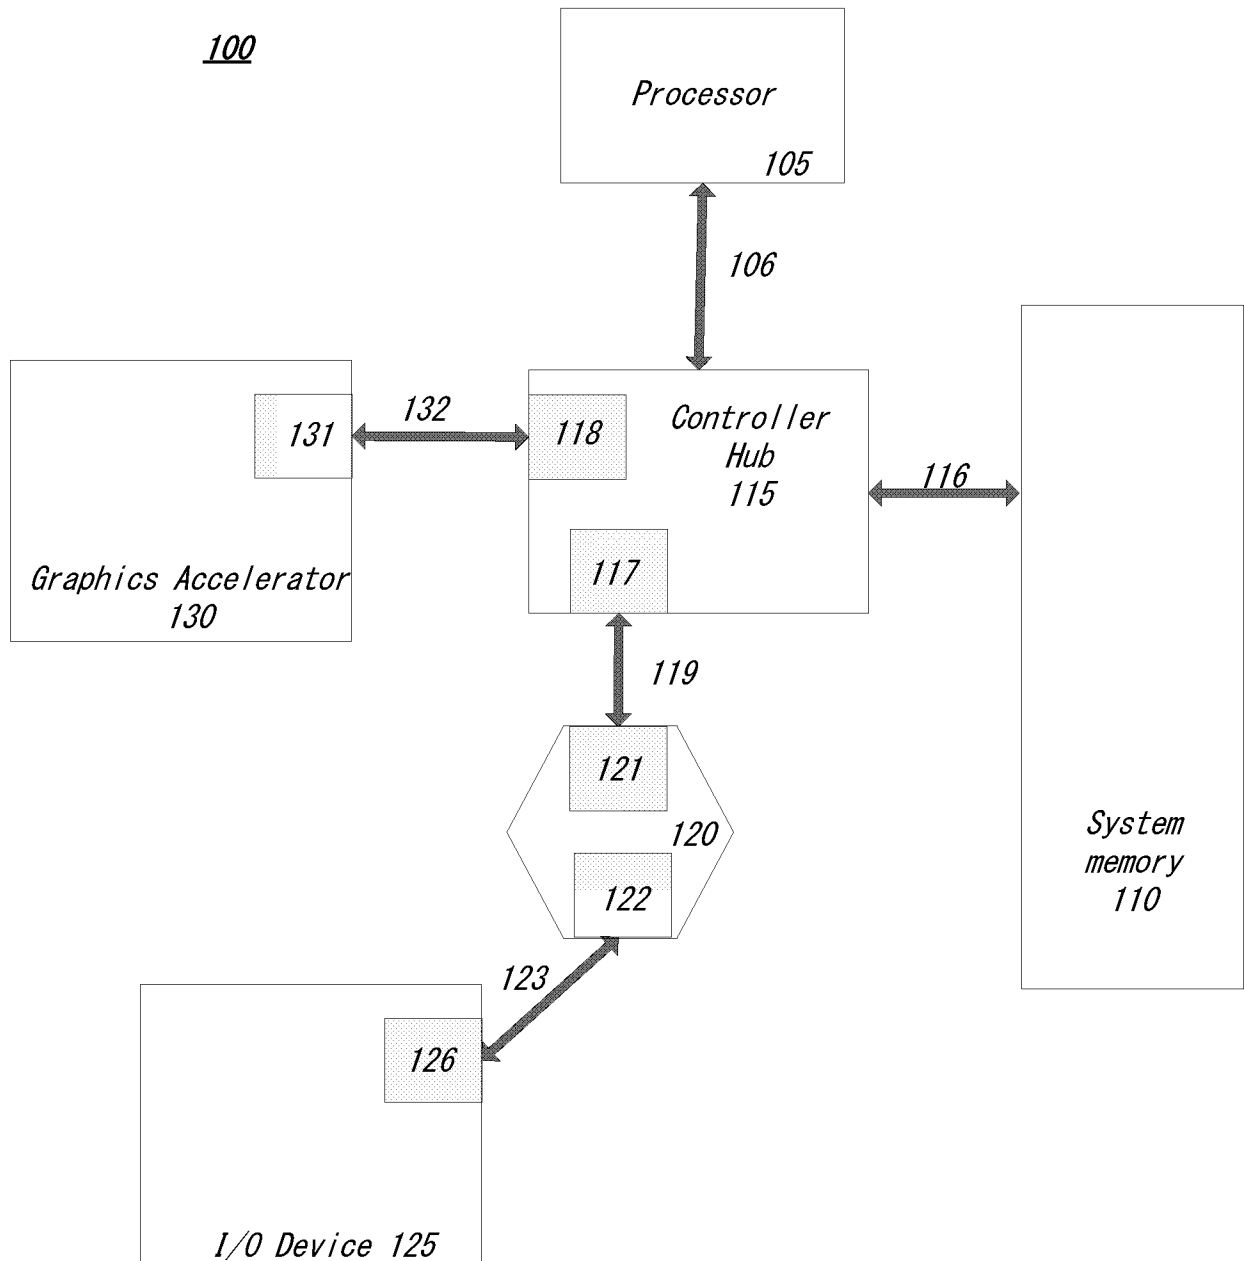


FIG. 1

2 / 11

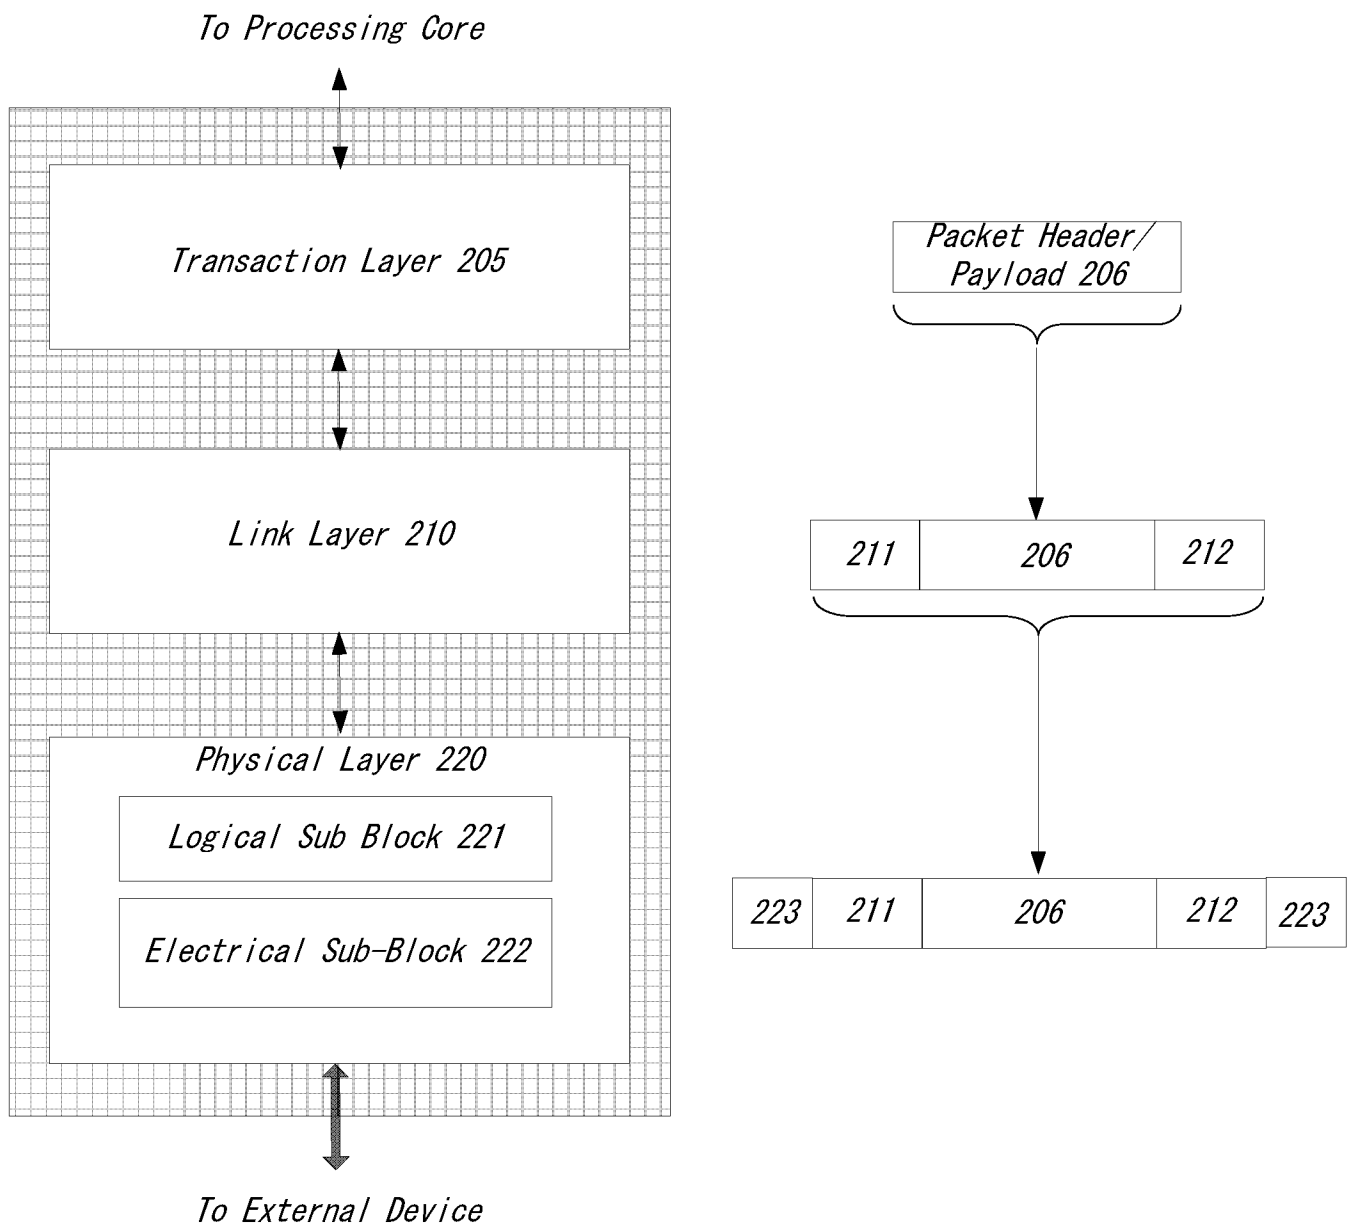
Layered Protocol Stack 200

FIG. 2

300





FIG. 3

4 / 11






Fig. 5

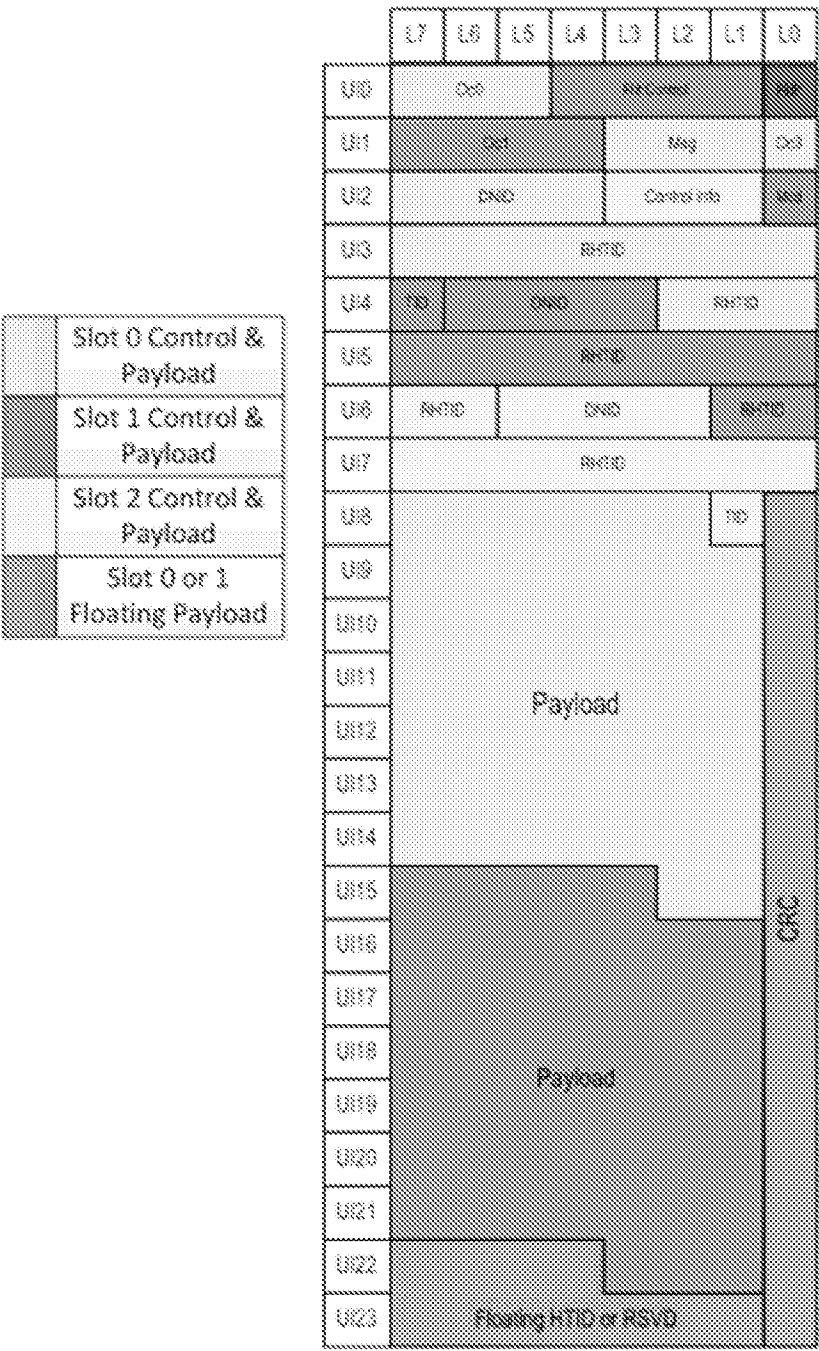


Fig. 6

7/11

	L7	L6	L5	L4	L3	L2	L1	L0
UI0	191	167	143	119	95	71	47	23
UI1	190	166	142	118	94	70	46	22
UI2	189	165	141	117	93	69	45	21
UI3	188	164	140	116	92	68	44	20
UI4	187	163	139	115	91	67	43	19
UI5	186	162	138	114	90	66	42	18
UI6	185	161	137	113	89	65	41	17
UI7	184	160	136	112	88	64	40	16
UI8	183	159	135	111	87	63	39	15
UI9	182	158	134	110	86	62	38	14
UI10	181	157	133	109	85	61	37	13
UI11	180	156	132	108	84	60	36	12
UI12	179	155	131	107	83	59	35	11
UI13	178	154	130	106	82	58	34	10
UI14	177	153	129	105	81	57	33	9
UI15	176	152	128	104	80	56	32	8
UI16	175	151	127	103	79	55	31	7
UI17	174	150	126	102	78	54	30	6
UI18	173	149	125	101	77	53	29	5
UI19	172	148	124	100	76	52	28	4
UI20	171	147	123	99	75	51	27	3
UI21	170	146	122	98	74	50	26	2
UI22	169	145	121	97	73	49	25	1
UI23	168	144	120	96	72	48	24	0

Fig. 7

UI	7	6	5	4	3	2	1	0
0,1,2,3	47	41	35	29	23	17	11	5
4,5,6,7	46	40	34	28	22	16	10	4
8,9,10,11	45	39	33	27	21	15	9	3
12,13,14,15	44	38	32	26	20	14	8	2
16,17,18,19	43	37	31	25	19	13	7	1
20,21,22,23	42	36	30	24	18	12	6	0

Fig. 8

UI	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0,1,2,3	47	44	42	39	37	34	32	29	27	24	22	19	17	14	12	9	7	5	3	1
4,5,6,7	46	43	41	38	36	33	31	28	26	23	21	18	16	13	11	8	6	4	2	0
8,9,10,11	45	45	40	40	35	35	30	30	25	25	20	20	15	16	10	12	9	7	4	2
12,13,14,15	47	44	42	39	37	34	32	29	27	24	22	19	17	15	13	11	8	6	3	1
16,17,18,19	46	43	41	38	36	33	31	28	26	23	21	18	19	14	14	10	9	5	4	0
20,21,22,23	47	44	42	39	37	34	32	29	27	25	23	21	18	16	13	11	8	6	3	1
24,25,26,27	46	43	41	38	36	33	31	28	26	24	22	20	17	15	12	10	7	5	2	0
28,29,30,31	45	45	40	40	35	36	30	32	29	27	24	22	19	17	14	12	9	7	4	2
32,33,34,35	47	44	42	39	37	35	33	31	28	26	23	21	18	16	13	11	8	6	3	1
36,37,38,39	46	43	41	38	39	34	34	30	29	25	24	20	19	15	14	10	9	5	4	0
40,41,42,43	47	45	43	41	38	36	33	31	28	26	23	21	18	16	13	11	8	6	3	1
44,45,46,47	46	44	42	40	37	35	32	30	27	25	22	20	17	15	12	10	7	5	2	0

Fig. 9

10/11

1005

1010





Fig. 10

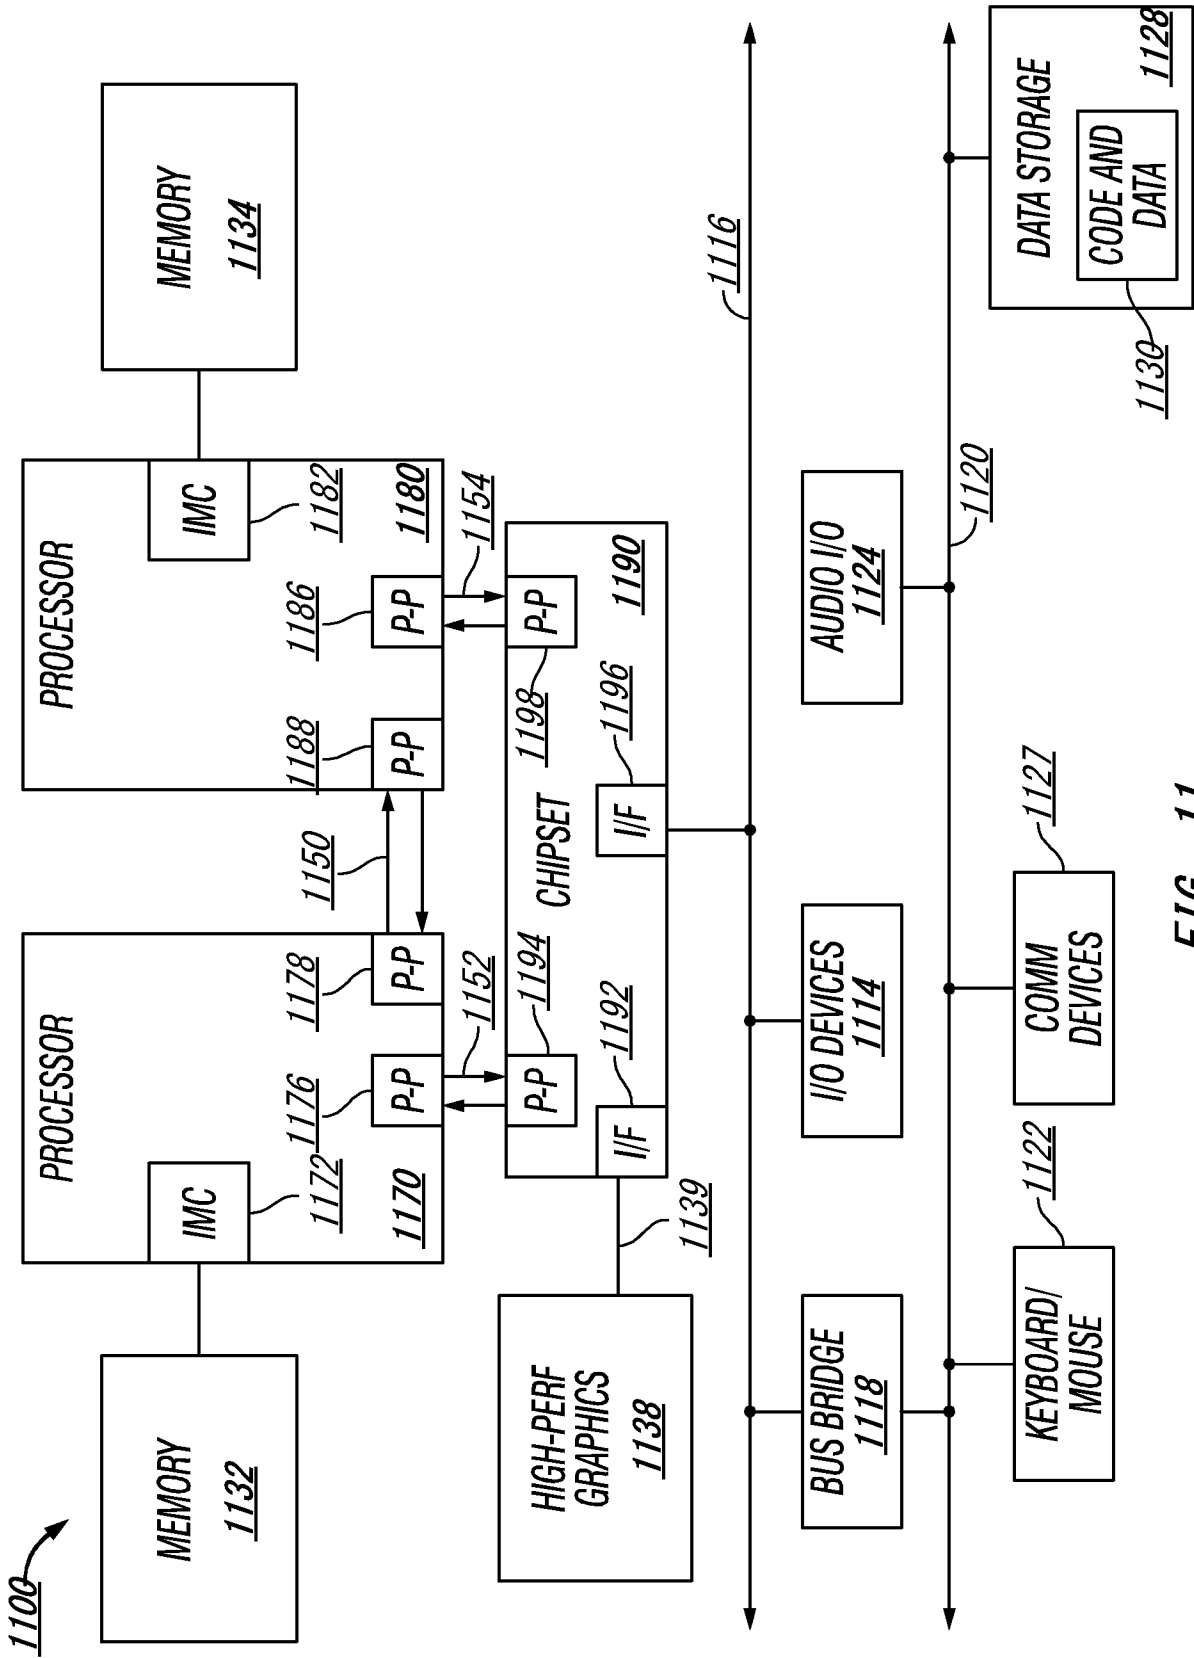


FIG. 11

A. CLASSIFICATION OF SUBJECT MATTER**G06F 13/14(2006.01)i, G06F 13/38(2006.01)i**

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G06F 13/14; G06F 13/40; G06F 15/16; H04J 3/02; H04J 3/04; G06F 13/42; G06F 15/167; G06F 13/00; G06F 12/08; G06F 13/38

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

Korean utility models and applications for utility models

Japanese utility models and applications for utility models

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

eKOMPASS(KIPO internal) & Keywords: flit, slot, format, link, packet, phit

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 7957428 B2 (MAURICE B. STEINMAN et al.) 07 June 2011 See column 2, line 9; column 3, lines 31-32; column 8, claim 1; column 9, claim 8; column 10, claim 19.	1-4, 25, 28-30, 33-34
A		5-24, 26-27, 31-32, 35-38
Y	US 8037224 B2 (JULIANNE JIANG ZHU et al.) 11 October 2011 See column 1, lines 35-38; column 7, lines 16-18, 19-21, 25, 38-39; column 29, claim 1.	1-4, 25, 28-30, 33-34
A	US 2012-0079156 A1 (ROBERT J. SAFRANEK et al.) 29 March 2012 See paragraphs [0028], [0036]; and figures 1, 3.	1-38
A	US 2004-0095967 A1 (DEBENDRA DAS SHARMA) 20 May 2004 See paragraphs [0028], [0031]; and figures 2A, 3.	1-38
A	US 2011-0238778 A1 (PHANINDRA K. MANNAVA et al.) 29 September 2011 See paragraphs [0029], [0032]; and figures 3, 4.	1-38



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

20 August 2013 (20.08.2013)

Date of mailing of the international search report

21 August 2013 (21.08.2013)

Name and mailing address of the ISA/KR

Korean Intellectual Property Office
189 Cheongsa-ro, Seo-gu, Daejeon Metropolitan City,
302-701, Republic of Korea

Facsimile No. +82-42-472-7140

Authorized officer

BYUN Sung Cheal

Telephone No. +82-42-481-8262



INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2013/032718

Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:

2. ☐ Claims Nos.:
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:

3. ☒ Claims Nos.: 39-40
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

1. ☐ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. ☐ As all searchable claims could be searched without effort justifying an additional fee, this Authority did not invite payment of any additional fee.
3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:

4. ☐ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

Remark on Protest

- ☐ The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.
- ☐ The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.
- ☐ No protest accompanied the payment of additional search fees.

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2013/032718

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 7957428 B2	07/06/2011	AT 464607 T	15/04/2010
		CN 103034605 A	10/04/2013
		CN 1700700 A	23/11/2005
		DE 602004026555 D1	27/05/2010
		EP 1598745 A2	23/11/2005
		EP 1598745 A3	02/11/2006
		EP 1598745 B1	14/04/2010
		JP 04035532 B2	23/01/2008
		JP 2005-332359 A	02/12/2005
		RU 2004130346 A	20/03/2006
		RU 2288542 C2	27/11/2006
		TW I311252 B	21/06/2009
		US 2005-0259599 A1	24/11/2005
		US 2005-0259696 A1	24/11/2005
		US 8204067 B2	19/06/2012
US 8037224 B2	11/10/2011	CN 100498757 C	10/06/2009
		CN 101878475 A	03/11/2010
		CN 1842781 A0	04/10/2006
		EP 2174229 A1	14/04/2010
		EP 2174229 A4	15/02/2012
		JP 04498356 B2	07/07/2010
		JP 2007-500886 A	18/01/2007
		JP 2009-026320 A	05/02/2009
		JP 2010-079921 A	08/04/2010
		KR 10-1279473 B1	30/07/2013
		US 2004-0103248 A1	27/05/2004
		US 2005-0027793 A1	03/02/2005
		US 2005-0033831 A1	10/02/2005
		US 2005-0033832 A1	10/02/2005
		US 2005-0033889 A1	10/02/2005
		US 2005-0041651 A1	24/02/2005
		US 2005-0041666 A1	24/02/2005
		US 2005-0044308 A1	24/02/2005
		US 2005-0044323 A1	24/02/2005
		US 2005-0044324 A1	24/02/2005
		US 2005-0055502 A1	10/03/2005
		US 2005-0055503 A1	10/03/2005
		US 2005-0055504 A1	10/03/2005
		US 2005-0055510 A1	10/03/2005
		US 2005-0055540 A1	10/03/2005
		US 2005-0086361 A1	21/04/2005
		US 2007-0204130 A1	30/08/2007
		US 2008-0062927 A1	13/03/2008
		US 2008-0126709 A1	29/05/2008
		US 2008-0140956 A1	12/06/2008
		US 2008-0184008 A1	31/07/2008
		US 2008-0216074 A1	04/09/2008
		US 2010-0042785 A1	18/02/2010

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2013/032718

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
		US 2010-0077150 A1	25/03/2010
		US 2010-0318703 A1	16/12/2010
		US 2012-089762 A1	12/04/2012
		US 7334086 B2	19/02/2008
		US 7346757 B2	18/03/2008
		US 7461213 B2	02/12/2008
		US 7467243 B2	16/12/2008
		US 7509462 B2	24/03/2009
		US 7509476 B2	24/03/2009
		US 7627717 B2	01/12/2009
		US 7627721 B2	01/12/2009
		US 7864760 B2	04/01/2011
		US 7924828 B2	12/04/2011
		US 7941603 B2	10/05/2011
		US 7961723 B2	14/06/2011
		US 7984268 B2	19/07/2011
		US 7991977 B2	02/08/2011
		US 8015567 B2	06/09/2011
		US 8065456 B2	22/11/2011
		US 8176298 B2	08/05/2012
		WO 2005-013061 A2	10/02/2005
		WO 2005-013061 A3	08/12/2005
		WO 2009-017668 A1	05/02/2009
US 2012-0079156 A1	29/03/2012	CN 103119908 A	22/05/2013
		TW 201220064 A	16/05/2012
		WO 2012-040648 A2	29/03/2012
		WO 2012-040648 A3	28/06/2012
US 2004-0095967 A1	20/05/2004	JP 04925559 B2	25/04/2012
		JP 2004-173278 A	17/06/2004
		US 7505486 B2	17/03/2009
US 2011-0238778 A1	29/09/2011	CN 102209104 A	05/10/2011
		JP 05085758 B2	28/11/2012
		JP 2011-211711 A	20/10/2011
		US 2013-103783 A1	25/04/2013
		US 8473567 B2	25/06/2013