



(19) **United States**
(12) **Patent Application Publication**
Schoinas et al.

(10) **Pub. No.: US 2016/0255008 A1**
(43) **Pub. Date: Sep. 1, 2016**

(54) **SEPARABLE TRANSPORT LAYER IN CACHE COHERENT MULTIPLE COMPONENT MICROELECTRONIC SYSTEMS**

Publication Classification

(71) Applicant: **Intel Corporation**, Santa Clara, CA (US)
(72) Inventors: **Ioannis T. Schoinas**, Portland, OR (US); **Doddaballapur N. Jayasimha**, Sunnyvale, CA (US)
(73) Assignee: **Intel Corporation**, Santa Clara, CA (US)
(21) Appl. No.: **15/060,965**
(22) Filed: **Mar. 4, 2016**

(51) **Int. Cl.**
H04L 12/801 (2006.01)
H04L 12/823 (2006.01)
H04L 1/18 (2006.01)
(52) **U.S. Cl.**
CPC *H04L 47/193* (2013.01); *H04L 1/1803* (2013.01); *H04L 47/34* (2013.01); *H04L 1/1874* (2013.01); *H04L 47/32* (2013.01)

Related U.S. Application Data

(63) Continuation of application No. 11/479,922, filed on Jun. 30, 2006, now Pat. No. 9,304,964.

(57) **ABSTRACT**

A separable Transport Layer is described in the context of cache coherent multiple component micro-electronic systems. In one example, a packet is received from a source component, the packet containing a Protocol Layer. A Transport Layer is attached to the packet and the packet is sent across a component communications interface to a second component, the packet containing the Transport Layer and the Protocol Layer.

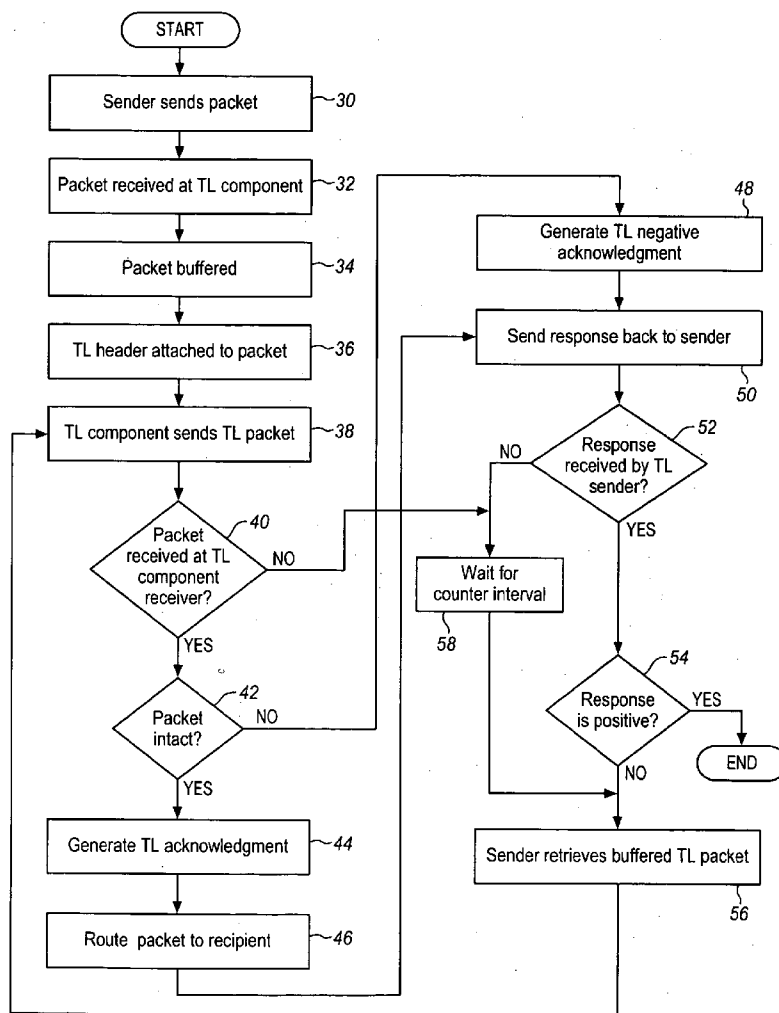


FIG. 1

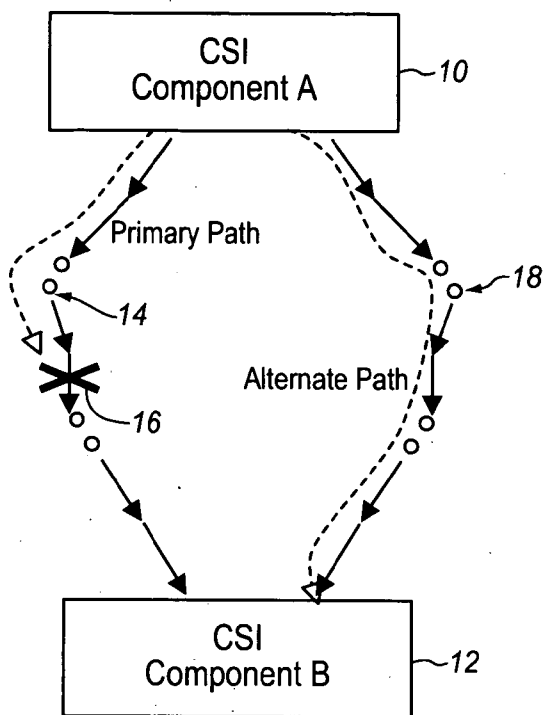


FIG. 2

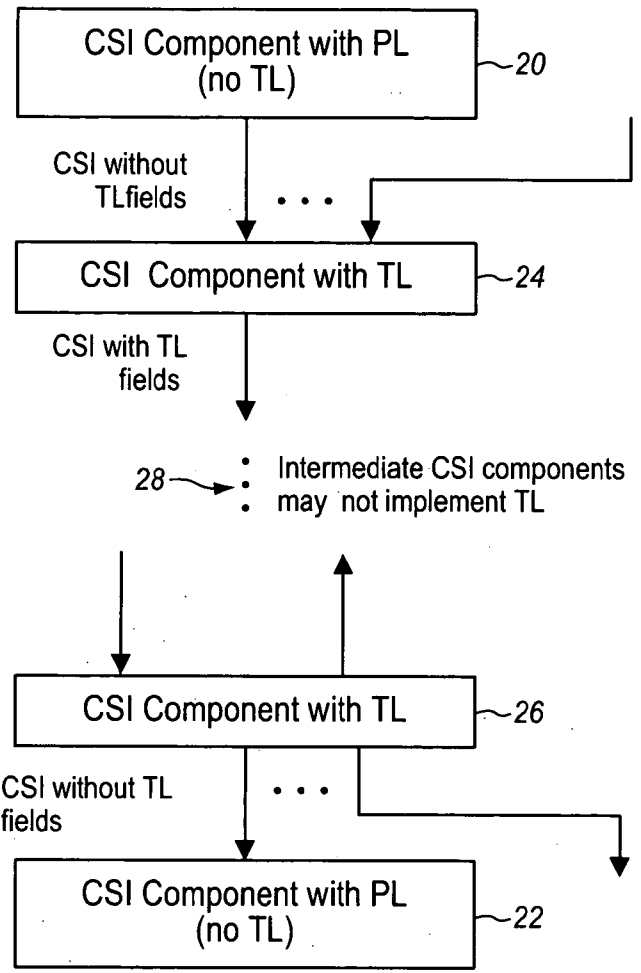
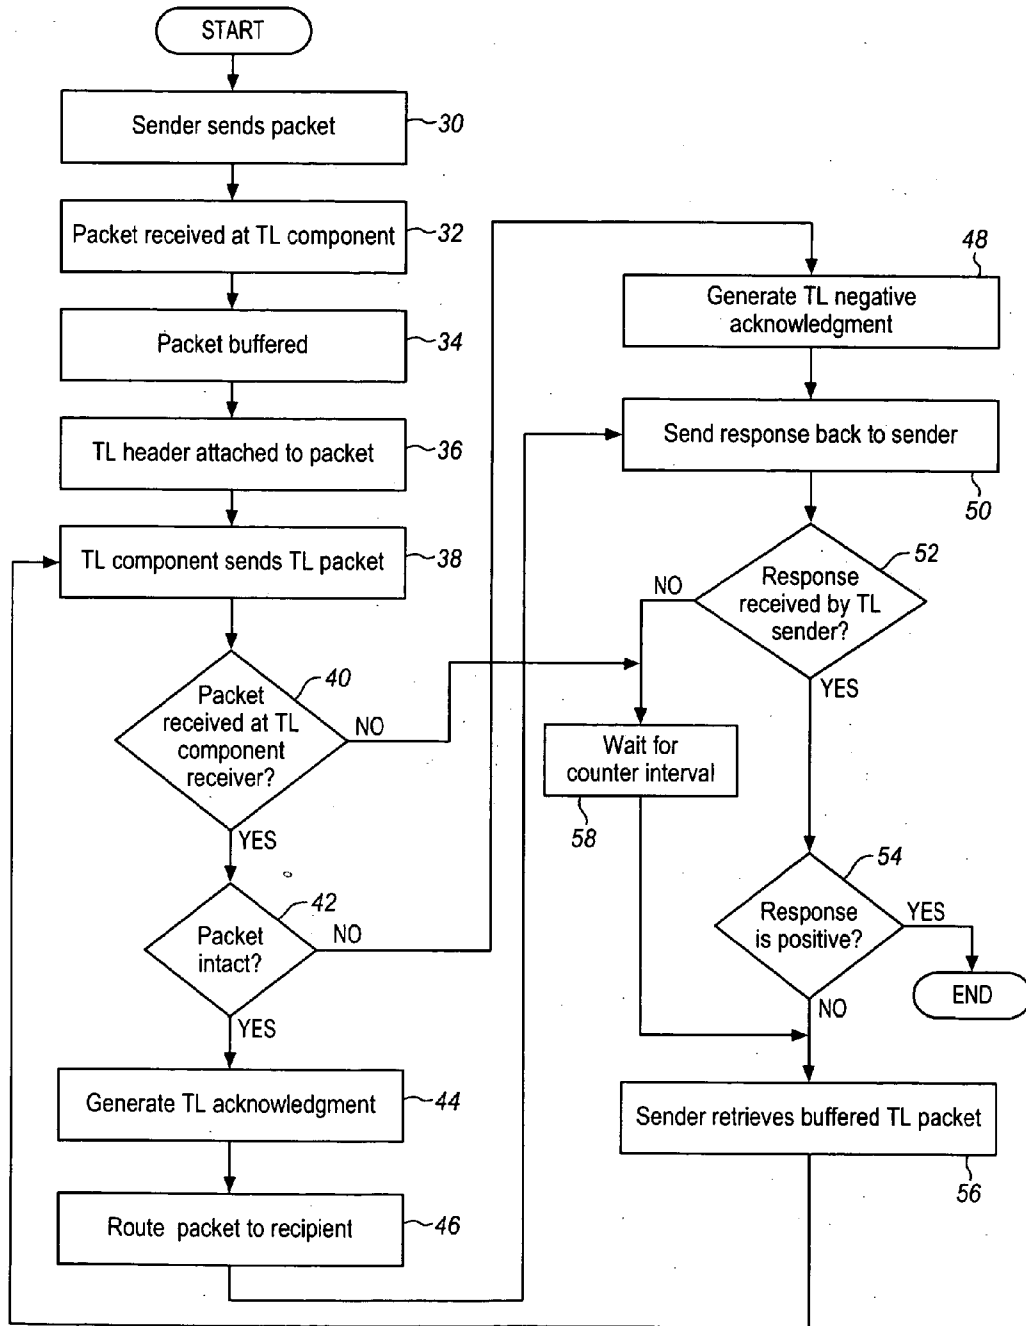


FIG. 3






FIG. 4




FIG. 5

**SEPARABLE TRANSPORT LAYER IN CACHE
COHERENT MULTIPLE COMPONENT
MICROELECTRONIC SYSTEMS**

FIELD

[0001] The present description relates to coherent interconnection links in multiple processor communication systems, and in particular to an additional layer to provide explicit confirmation of received data packets.

BACKGROUND

[0002] For years, computer systems with more than one microprocessor would use a host bus that is shared by the processors and the supporting chipset. The processors and the chipset, for example an input/output controller hub or memory controller hub would all connect to the same bus and the bus might also include cache memory or maybe even graphics. Increasingly, computer systems are being built with point-to-point links between the processors and with the chipset instead of the shared connections. This is done not only in high power workstation computers, but also in multi-processor server computers. With a point-to-point link, a coherent interface may connect a microprocessor directly to another microprocessor. Another separate link may connect the microprocessor to the input/output or memory controller hub. Another link may connect the microprocessor to a third microprocessor. Where there is a small number of processors, such as 2 or 3, a point-to-point link can allow the link from each processor to be directly connected to another processor's link. The link may be built into each processor and already configured to be installed into and operate within a multiple processor computer system.

[0003] If, on the other hand, there are more than a very few processors, such a point-to-point link to each component becomes cumbersome. Also, in large scale workstation and server systems, where a very high degree of reliability is desired, a special chipset and a unique, typically proprietary, interface link definition is often used. The unique chipset and link add additional costs to the computer systems. While the additional layers required for higher reliability could be added to all microprocessors and chipsets, spreading the costs over a much larger volume of components, this would add cost and complexity to the smaller scale systems. It might also slow communications over the interface in systems that have very little extra margin for speed losses.

BRIEF DESCRIPTION OF THE FIGURES

[0004] The invention may be more fully appreciated in connection with the following detailed description taken in conjunction with the accompanying drawings, in which like reference numerals refer to corresponding parts throughout the several views of the drawings, and in which:

[0005] FIG. 1 is a block diagram of two components communicating through primary and alternate paths according to an embodiment of the invention;

[0006] FIG. 2 is a block diagram of two components communicating through Transport Layer-enabled components according to an embodiment of the invention;

[0007] FIG. 3 is a timing diagram of enter entering a quiescent state on a communications interface according to an embodiment of the invention; and

[0008] FIG. 3 is a process flow diagram of reliably sending packets on a communications interface according to an embodiment of the invention;

[0009] FIG. 4 is a block diagram of a computer system with communications buses suitable for implementing embodiments of the present invention.

[0010] FIG. 5 is a block diagram of another computer system with communications buses suitable for implementing embodiments of the present invention.

DETAILED DESCRIPTION

[0011] An added degree of reliability may be added to a small scale computer system's processor interface link by adding an optional Transport Layer to, for example, a point-to-point interconnect. OEMs (Original Equipment Manufacturers) that desire a higher degree of reliability may use a point-to-point interconnect with the Transport Layer in their components. OEMs that do not desire the higher reliability or the extra processing and communications overhead of the Transport Layer may ignore it. As described below, components that implement the Transport Layer may work seamlessly with components that do not. In a further development, components that implement the Transport Layer and those that don't may exist in the system simultaneously.

[0012] The Transport Layer provides support for end-to-end reliable transmission between two interconnected components each implementing this layer. It relies on the services provided by the Routing Layer below it, while, in turn, providing reliable transmission support to the Protocol Layer above it. The Transport Layer may be particularly valuable, for example, for inter-board traffic carried by cables. For intra-board signals carried on traces, the Transport Layer may also be used but, depending on the application, may not be necessary or desired.

[0013] FIG. 1 shows two components A 10 and B 12 which both implement the Transport Layer (FIG. 1). These components have a point-to-point interconnect to allow packets to be communicated between them. Embodiments of the present invention may be applied to many different component interconnect interfaces and protocols. The components may be processors, controllers, memory devices, or communications hubs.

[0014] If both components A and B implement the Transport Layer, then reliable end-to-end transmission from the sender component A to the destination component B is maintained without any intermediate connected components in between. The path from A to B to implement the Transport Layer, may be direct. The granularity for reliable transmission is a packet.

[0015] Consider at least two disjoint routable paths from A to B, a primary path 14 to the left and an alternate path 18 to the right. In a simple example, there are no failures along one of these paths, the Transport Layer mechanism guarantees the delivery of a packet originating at A and destined to B. In other words, the Transport Layer provides a mechanism against one or more failures along one of the disjoint paths. In general, between any two Transport Layer entities, the Transport Layer mechanism, described below, provides protection against a single point of failure.

[0016] If there is a failure along one of the paths, that is if a packet from A does not successfully arrive at B, then the Transport Layer may provide for retransmission. Transmission and retransmission may be based on a notion of Transport Layer retry (TLR). If component A determines that a

packet that it sent to component B along the primary path may have been lost, then it may retransmit the packet to B until it succeeds or until, after a certain number of tries, it gives up.

[0017] A threshold may be used to limit the number of retries. Component A may then repeatedly resend the packet and keep track of the number of retries until the threshold is reached. At that point, the retries stop. In FIG. 1, the packet is first sent along the primary path 14, but there is a break in the path at some point 16. This may occur because the packet is routed on a path through a failed component. Packets may fail for many other reasons. The packet is then resent along an alternate path 18.

[0018] To allow for the retry mechanism, a timer may be used between retries to give enough time for the packet to be received and acknowledged. The time between two such transmissions may be referred to as a retry time out (rTTO). The sending component A may also have a buffer to store the packet for retransmission until A is guaranteed that the packet has been received at B.

[0019] The other side of the Transport Layer is an acknowledgment. The destination component B may acknowledge the receipt of each packet by sending a special acknowledgement response to A.

[0020] FIG. 1 shows two completely independent paths between components A and B. In that example, there are no single points of failure in the interconnect, because the Transport Layer agent is able to send a request down alternate paths. While these two independent paths may enhance reliability, only one path is needed to implement a Transport Layer. In a complete system, there may be some components that are connected with more than one path and other component interconnections that have only one path. The routing and configuration of the paths may be designed to suit a variety of different applications depending on the reliability desired as well as other factors such as packaging, cost constraints, performance demands, etc.

[0021] A packet in a point-to-point interconnect may contain a variety of different fields, depending on the particular message and the application to which the packet format has been adapted. The fields may include addresses, commands, parameters, data, operation codes, transaction requests and identifiers, and profile definitions. The particular type and arrangement of the fields is not important to the present description. In addition, while the present description is presented in the context of adding a Transport Layer to packets, the Transport Layer may be adapted to a variety of different types of packet formats and communications protocols.

[0022] The Transport Layer adds some fields to the communications protocol packet. The Transport Layer may include a 1) sequence number, 2) a transaction response field (Transport Layer Response or TLR), 3) a sender node identification (NodeID), and 4) a time-out field. The structure of the Transport Layer may be adjusted to suit a particular application, packet format, or communications protocol. More or fewer fields may be used and the positions of the fields may be moved to suit particular applications.

[0023] The sequence number field allows a unique sequence number to be associated with each packet sent from a source component to a destination component. For each source-destination pair, the sequence number may be unique as long as the packet is alive from the Transport Layer's point of view. Once the sender has received a confirmation that the packet has been consumed at the receiver, the sequence number may be retired and can be reused.

[0024] The sequence number identifies a packet in a sequence of packets, without counting the number of retries. The first transmission of a packet in a sequence may be labeled as 0. The next packet in the sequence may then be identified as 1, and the next packet with 2, and so on. When the same packet is resent, it may be sent with the same sequence number each time. This allows the recipient to identify the duplicate packets as such in the Transport Layer. Any number of other packet numbering schemes may be used as an alternative to this scheme depending on the implementation. For example, the first packet may not be labeled at all and the next packet labeled as 0. Various coded or nonsequential numbering schemes may also be used.

[0025] The sequence number field may have any number of bits depending on the particular application. The field width may be selected based on different design considerations. One consideration is to provide enough unique sequence numbers to prevent transactions from backpressuring into the Protocol Layer when unique sequence numbers run out. Another consideration is to provide enough sequence numbers to prevent wrap around within a packet's life time in the Transport Layer. In other words, there is a round trip time for the retry process that includes at least the time necessary for: a) a packet to be sent from the sender to the receiver, b) the sender to wait for and not receive an acknowledgment, c) the sender to determine that no acknowledgment is coming, d) the sender to retry the packet and e) the receiver to receive the retry. During this round trip time, additional packets are sent with additional sequence numbers. From the receiver's perspective, providing enough sequence numbers to accommodate the round trip time will help the Transport Layer at the receiver to recognize the retry as a retry and not as a duplicate of a later sent packet. Stated another way, the retry packet should be clearly recognizable as a never received packet and not as a duplicate of a later sent packet.

[0026] A time-out field may also be used in the packet header. Such an explicit time-out may be used to facilitate dropping "aged" packets in the interconnect. However, a time-out field may add additional overhead. Instead of an explicit time-out field, the support mechanisms to drop packets in order to prevent sequence number wraparound as described below may be used.

[0027] The wraparound constraint described above is implicit in the Transport Layer protocol. However, for greater flexibility it may be made explicit using another field or by embedding a constraint in another message. A wraparound constraint may be used to prevent a sequence number wraparound that could result in a destination component being unable to distinguish between a duplicate packet and a new packet. It may be permitted for the Transport Layer to stop accepting transactions from the Protocol Layer.

[0028] The transaction response field may be used for acknowledgments and similar messages as described in more detail below. The sender Node ID, indicates the source of the message and may be used as an address for sending acknowledgments. For some interfaces, the information in the NodeID field may be present elsewhere in the packet. Placing it in the Transport Layer allows it to be accessed within the Transport Layer which may be simpler in some embodiments. Other information may be accessed from the packet, as appropriate, depending on the circumstances.

[0029] In one embodiment, every packet has a sequence number field and a NodeID field. The TLR field may be optional. The TLR field is able to support a variety of different

transaction response types, for example a Transport Layer Acknowledgment (TL_ACK) and a Transport Layer Negative Acknowledgment (TL_NACK).

[0030] These fields and transactions may be configured so that they are not visible to the Protocol Layer. In one embodiment, the Transport Layer ACKs and NACKs may be piggy-backed along with the regular packets. The overall savings in interconnect bandwidth comes at the cost of increasing the packet header to accommodate the additional sequence number field and the ACK/NACK.

[0031] For routing, each Transport Layer component may have a set of unique protocol agents in its domain. The protocol agents associated with the Transport Layer component may be configured to launch Transport Layer transactions only from the one Transport Layer component and no other Transport Layer components in the system. Consequently, any packets targeted to a member of the set of unique protocol agents that use the Transport Layer are always routed through the one Transport Layer component. Such a Transport Layer component may be integrated into the same chip with its protocol agent or agents or provided in a separate package.

[0032] The TLR field may carry different Transport Layer responses. TL_ACK and TL_NACK may be treated as no data responses and routed in that message class with the underlying protocol. All other packets may be routed in the same manner as they would be without the Transport Layer.

[0033] In one embodiment, each Transport Layer request is acknowledged by the receiving component. This maximizes reliability at the expense of additional traffic overhead. When an acknowledgment is not received within tRTO (Retry Time Out), the packet is presumably lost and the sender retransmits the packet with the same sequence number as before (TLR). The number of retries may be selected to suit a particular application. As an alternative, only some packets are acknowledged. The determination of which packets to acknowledge may be set, for example, by the sender using a TLR field in the Transport Layer of the sent packet.

[0034] In some instances, a packet may be successfully received but the return acknowledgment may fail. The sender will nevertheless resend the packet for failure to receive the acknowledgment. This may result in a duplicate packet at the receiver. In one example, the receiver drops the second packet at the Transport Layer and does not pass it on to the Protocol Layer. The receiver may use the sequence number, for example, to identify the duplicate. When the sender retries the packet, it may be sent with the same sequence number as the original packet which was assumed to be lost.

[0035] For simplicity, the Transport Layer may be implemented with no explicit Transport Layer request transactions. Protocol Layer requests and responses may be used as implicit Transport Layer requests. As described above, at least two new response transactions unique to the Transport Layer may be defined.

[0036] The Transport Layer Acknowledgment (TL_ACK) is a response sent by the Transport Layer component acknowledging the error-free receipt of a Transport Layer request. The request may be considered to be implicit in the sending of a packet with a Transport Layer sequence number and NodeID. The response acknowledges the receipt of the request and identifies the request by the sequence number. It is directed to the component identified by the sender NodeID in the request.

[0037] In one embodiment, several Transport Layer requests may be acknowledged with one TL_ACK packet. By

not acknowledging each Transport Layer request individually, communication resources are conserved for other packets. A range of requests may be acknowledged with a single response by tracking the sequence numbers and NodeIDs for a group of packets and then sending a single packet that indicates the range.

[0038] In one embodiment, for a particular source-destination pair, that is, a sender-receiver combination, an acknowledgement may be sent with the highest sequence number received since the last acknowledgement. The receiver will know if there were any gaps in the sequence based on gaps in the sequence of sequence numbers. So, for example, assume that a request with a sequence number S11 was previously received and acknowledged and then packets with sequence numbers S12-S32 are all successfully received.

[0039] The receiver may send a TL_ACK response with sequence number S32 directed to the source in order to acknowledge the receipt of all of the requests with sequence numbers in the range from S 12-S32. More generally, if the sender receives a TL_ACK for Si and then a TL_ACK for Sn, then this serves as an acknowledgement for packets from S1+1 to Sn(modulo 2k), where k is the width of the sequence number field in the packet. Acknowledging a range may cause the time between sending a packet and receiving an acknowledgement to increase. The width of the sequence number and any constraints on acknowledging ranges may be selected to avoid errors due to wraparound.

[0040] A Transport Layer Negative Acknowledgment (TL_NACK) may be used by a receiver to acknowledge an erroneous receipt of a Transport Layer request. The response “negatively acknowledges” the request identified by the sequence number and is directed to the component identified by the sender NodeID in the request. Such a transaction is not necessary for the Transport Layer because a sender may be configured to retransmit a message if it does not receive a TL_ACK response within the set tRTO. A TL_NACK message may be used, however, to cause the retries to happen sooner. If a packet is lost along the path, then the retry time out may be used, while if the packet is received at the receiver but somehow in error, then a TL_NACK may be used to stimulate the retry.

[0041] The NodeID may be implemented in different ways to suit a particular transaction. A direct NodeID may be inserted into the Transport Layer field supplied in the packet header to explicitly identify the sending component.

[0042] An indirect NodeID may be used if a protocol agent uses the sender NodeID. The Transport Layer component need not then be identified explicitly. The sender NodeID may refer to the NodeID of the originating Protocol Layer request or response corresponding to this Transport Layer request. Specific responsibilities may be imposed on both the sender and the receiver TL entities to avoid errors or misidentifications.

[0043] FIG. 2 shows an example of how packets may be communicated through a coherent inter-component interface with a Transport Layer. In FIG. 2, a first component 20 is shown as a sender without any Transport Layer functionality. A second component 22 is shown as a receiver without any Transport Layer functionality. These two components may communicate with each other in the usual manner without benefit of a Transport Layer. For example, the sending component may use routing tables and may use a sender NodeID field in the sent packets. In, some point-to-point interconnect

protocols, a sender NodeID field is provided as an optional feature and may be used to send packets to the receiver component 22.

[0044] FIG. 2 further shows two further components, a first relay component 24 near the sender 20 and a second relay component 26 near the receiver 22. In contrast to the first two components, these implement the Transport Layer (TL) described above. To provide reliable end-to-end transmission between the first two components, these additional TL components may be interposed in the communication path between the first two. Such an interposition may be made seamlessly in that neither of the first two components 20, 22 require any modification in structure or operation in order to interoperate with the TL components. The routing tables, for example, do not change. Additional components 28 with and without a Transport Layer functionality may also be interposed between the first two components, between the second two components and between the components that use a Transport Layer and those that do not.

[0045] When the sender 20 sends a packet across the coherent interface to the receiver, it will be intercepted at the first TL component 24. At the first TL component, the packet, whether it be a Protocol Layer request or a response, may be converted into a Transport Layer request or response. The packet may also be buffered at this first relay component. To do this, the first TL component 24 may generate the TL header information by generating the information necessary for each field and adding the header to the intercepted packet. The NodeID field may be defined by a protocol agent or taken from explicit NodeIDs elsewhere in the packet. The sequence number may be generated using the first TL component's own sequence series or a sequence based on the sending component's 20 sequence. The transaction response field is not necessary for sent packets, however, it may be used for some applications.

[0046] The first TL component 24 then sends the packet with the Transport Layer attached across the coherent interface to the receiver 22. The packet may then be intercepted by the second TL component 26 that is near the receiver. If there are additional components 28 along the path, they may intercept the packet too, for a successful transmission, the packet will continue on its course unaffected. In one embodiment, there may be intermediate components that support the Protocol Layer but not the Transport Layer. These intermediate components may be configured to pass the packet header Transport Layer fields unmodified from the Transport Layer source to Transport Layer destination.

[0047] For the receiver without a Transport Layer to benefit from the reliability improvements offered by the Transport Layer, the second TL component is located sufficiently close to the receiver. The physical positioning of the various components may be adapted to suit any particular implementation.

[0048] The second TL component 26 determines whether the packet has arrived intact, then looks at the Transport Layer of the packet and generates a Transport Layer acknowledgment. If the packet does not arrive intact, then the acknowledgment may be a negative acknowledgment. The receiver then sends a TL_ACK or TL_NACK back to the sender based on the NodeID in the Transport Layer. The receiving TL component 26 then routes the packet to the non-TL equipped and intended recipient, the first receiver 22.

[0049] The TL-equipped receiver may perform some additional tests to see whether the packet satisfies other condi-

tions. The particular tests may depend upon the particular application to which the TL is applied. The receiving TL component or node may, for example, check the sequence number field of the Transport Layer to determine whether the packet is a duplicate.

[0050] The sender NodeID, in the example above, corresponds to the original sender's 20 NodeID, not the TL enabled sender's 22 NodeID. However, the routing may be modified so that the sender NodeID is for the TL components and the ultimate sender and receiver are identified in another way. For example, if the sender NodeID is not defined at the Protocol Layer component, then each Transport Layer component may have a unique NodeID which is then used as the sender NodeID in the outgoing Transport Layer request. In such a case, TL_ACKs and TL_NACKs are targeted explicitly to the TL components 24, 26, and not to the ultimate sender and receiver 20, 22.

[0051] In the present example, the route between the ultimate sender and receiver is configured to pass through the TL enabled sender and receiver. The TL enabled components also are configured to act as proxies for their respective ultimate senders and receivers. The TL enabled sender can then receive, for example the TL_ACKs and TL_NACKs and generate an appropriate response based on packets stored in its buffer.

[0052] FIG. 3 shows an example process flow of using a Transport Layer as described with respect to FIGS. 1 and 2 above. First, at block 30, the sender 20 sends a packet across the coherent interface in the direction of a designated receiver 22. The packet is received at a Transport Layer component at block 32. The TL component may be part of the same component as the sender or a separate discrete component. At the TL components at block 34, the packet is buffered and at block 36, a TL header is generated and attached to the packet, converting it into a Transport Layer request or response. The TL header as described above may include a transaction response field, a NodeID field and a sequence# field. Other fields may also be used and neither of these fields is required. The Transport Layer enabled component 24 then sends the packet at block 38 with the Transport Layer attached across the coherent interface toward the designated recipient.

[0053] If the packet is successful, at block 40, it is received at a TL component receiver 26. The TL component receiver determines at block 42 whether the packet has arrived intact. If so, then it generates a Transport Layer acknowledgment at block 44 and routes the packet to the non-TL equipped and intended recipient, the first receiver 22, at block 46. The TL component may remove the Transport Layer header or not before sending it to the recipient, depending on the implementation. As with the sending TL component, the receiving TL component may be a part of the receiver or a discrete component.

[0054] If the packet does not arrive intact, then the TL component generates a negative acknowledgment at block 48. The receiver then sends the response back to the sender at block 50, based on the NodeID in the Transport Layer. The response is received by the TL sender component at block 52. If the response is a positive acknowledgment, then at block 54, the Transport Layer process ends. If the response is a negative acknowledgment, then the sender retrieves the corresponding buffered packet at block 56 and resends it including the appropriate TL header as at block 38. In the described embodiment, the sender resends exactly the same packet with exactly the same header, but the packet or its header may be

modified for some applications. The resent packet is treated in the same way as the original packet at block 38. As mentioned above, a retry counter may be activated to limit the number of attempts to send the packet.

[0055] Until the sender receives the response, it continues to operate a counter. After the appropriate amount of time, at block 58, if no response is received from the TL component receiver, then the TL component sender will return to block 56 to retrieve the buffered packet and then resend the packet. As mentioned above, the number of retries may be limited using a counter or other mechanism.

[0056] FIG. 4 shows an example of a multiple component system using TL components and components without a Transport Layer as an example of possible arrangements for communicating components. In FIG. 4, there are four printed circuit boards (PCB) 62-0, 62-1, 62-2, 62-3. Each PCB carries a group of four interconnected microprocessors 64-0, 64-1, 64-2, 64-3, 66-0, 66-1, 66-2, 66-3, 68-0, 68-1, 68-2, 68-3, 70-0, 70-1, 70-2, 70-3. While microprocessors are shown, any other type of microelectronic device may be used. In addition, the devices are not necessarily all of the same type or kind, so that different types of devices may be carried by each PCB or by a single PCB. The microprocessors are each connected to each other microprocessor by a group of dedicated communication channels. In the illustrated example, each microprocessor has a discrete link to each other microprocessor on the board. However, connections may be made through other components as may be desired for a particular application. The microprocessors are also shown as being connected to other components, indicated generally as X. These components may be memory, hubs, I/O devices or any other component or components.

[0057] Each PCB also has two cable buffers 72-0, 72-1, 74-0, 74-1, 76-0, 76-1, 78-0, 78-1. The buffers connect to data communications cables 80 that connect the PCBs to each other. The cables allow microprocessors on each PCB to communicate with the microprocessors on each other PCB. As shown in FIG. 4, each buffer connects to two other PCBs and each PCB has two buffers so that there is a path from each PCB to each other PCB. The buffers also connect to each microprocessor using circuit board traces or another connector (not shown).

[0058] In the example of FIG. 4, the buffers may operate as Transport Layer enabled relay components similar to the TL components 24, 26 of FIG. 2, while the microprocessors operate without a Transport Layer similar to the sending and receiving components 20, 22 of FIG. 2. In this configuration, the microprocessors can communicate with other microprocessors on the same PCB without the additional overhead of a Transport Layer. For communications through the cables, however, the extra reliability of the Transport Layer may be used. Alternatively, the Transport Layer may be used for some cables but not for others. The configuration of FIG. 4 is provided as an example. A variety of modifications may be made depending on the particular application.

[0059] FIG. 5 shows an example of a computer system suitable for implementing the present invention. An IOH (Input/Output Hub), north bridge, or host controller 363 interfaces one or more CPUs (central processing unit) with memory and I/O devices and may provide a wide range of other features such as increased performance, reliability, availability and serviceability, and system management. It may include I/O clusters, a memory controller, snoop filters, and a wide range of logic for handling transactions. While the

example of FIG. 5, includes a microprocessor coupled to an IOH and an ICH (Input/Output Controller Hub), either the IOH or the ICH or both or any of the functions of these chips may be incorporated into any one or more of the microprocessors. The IOH and the ICH may also be combined, in whole or in part, inside of or outside of the microprocessor.

[0060] In the example of FIG. 5, the IOH 363 has a point-to-point interconnect link 309 for each of three CPUs or processor cores 311, 313, 315. More or less than one IOH, three processor cores and links may be used. As shown in the diagram, CPU1 has a direct link to CPU 2 and the IOH, but CPU 1 and CPU 3 communicate only through CPU 2 or the IOH. An additional link may be provided to allow CPU 1 and CPU 3 to communicate with each other directly. Alternatively, links may be removed so that all of the CPUs communicate through one of the other CPUs or the IOH.

[0061] The IOH provides additional connectivity to other devices. There is an interface to system memory 367, such as DIMMs (Dual In-line Memory Modules) in which instructions and data may be stored, and a high speed interface, such as PCI (peripheral component interconnect) Express. The PCI Express interface may be used to couple to a variety of different high and low speed devices. In the example of FIG. 5, six PCI Express x4 lanes are shown. Two lanes connect to a TCP/IP (Transmission Control Protocol/Internet Protocol) Offload Engine 317 which may connect to network or TCP/IP devices such as a Gigabit Ethernet controllers 339. Two lanes connect to an I/O Processor node 319 which can support storage devices 321 using SCSI (Small Computer System Interface), RAID (Redundant Array of Independent Disks) or other interfaces. Two more lanes connect to a PCI translator hub 323 which may support interfaces to connect PCI-X 325 and PCI 327 devices. Two, four or more lanes of PCI Express couple to a graphics controller 341 to render images or video on a display 337. The PCI Express interface may support more or fewer devices than are shown here. In addition, the IOH may be adapted to support other protocols and interfaces instead of, or in addition to those described.

[0062] The IOH may also be coupled, using PCI Express or another bus to an ICH. The ICH 365 offers possible connectivity to a wide range of different devices. Well-established conventions and protocols may be used for these connections. Alternatively, these connections may be provided using the PCI interface 327 or another interface. The connections may include a SIO (Super Input/Output) port 375, a USB hub 371, and a local BIOS (Basic Input/Output System) flash memory 373. The SIO (Super Input/Output) port 375 may provide connectivity for a front panel 377 with buttons and a display, a keyboard 379, a mouse 381, and infrared devices 385, such as IR blasters or remote control sensors. The I/O port may also support floppy disk, parallel port, and serial port connections 383. Alternatively, any one or more of these devices may be supported from a USB, PCI or any other type of bus or interconnect. Wireless interfaces such as Bluetooth and WiFi may also be supported from any one or more of these busses.

[0063] The particular nature of any attached devices may be adapted to the intended use of the device. Any one or more of the devices, buses, or interconnects may be eliminated from this system and others may be added. For example, video may be provided on the PCI bus, on an AGP bus, through the PCI Express bus or through an integrated graphics portion of the host controller or a processing core.

[0064] It is to be appreciated that a lesser or more equipped component communication protocol, Transport Layer mes-

sage structure, retry protocol, communications bus, and computer environment than the examples described above may be preferred for certain implementations. Therefore, the configuration of the signaling system and the computer system will vary from implementation to implementation depending upon numerous factors, such as price constraints, performance requirements, technological improvements, or other circumstances. Embodiments of the invention may also be applied to other types of software-driven systems that use different hardware architectures than those shown in the Figures.

[0065] While embodiments of the invention have been described in the context of an input/output hub chip coupled to microprocessors, memory, and other interconnects, embodiments of the invention may also be applied to a wide range of other devices capable of communicating over a point-to-point interconnect bus. In addition, embodiments of the invention may be applied to any device communications interface that transfers data bidirectionally with a communications interface. Embodiments of the invention may also be applied to a wide variety of components with simpler communications interfaces or test procedures

[0066] In the description above, for purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of the present invention. It will be apparent, however, to one skilled in the art that the present invention may be practiced without some of these specific details. In other instances, well-known structures and devices are shown in block diagram form.

[0067] The present invention may include various steps. The steps of the present invention may be performed by hardware components, such as those shown in the Figures, or may be embodied in machine-executable instructions, which may be used to cause general-purpose or special-purpose processor or logic circuits programmed with the instructions to perform the steps. Alternatively, the steps may be performed by a combination of hardware and software.

[0068] The present invention may be provided as a computer program product which may include a machine-readable medium having stored thereon instructions which may be used to program an agent or a computer system to perform a process according to the present invention. The machine-readable medium may include, but is not limited to, floppy diskettes, optical disks, CD-ROMs, and magneto-optical disks, ROMs, RAMs, EPROMs, EEPROMs, magnet or optical cards, flash memory, or other type of machine-readable media suitable for storing electronic instructions. Moreover, the present invention may also be downloaded as a computer program product, wherein the program may be transferred from a remote computer to a requesting computer by way of data signals embodied in a carrier wave or other propagation medium via a communication link (e.g., a modem or network connection).

[0069] Many of the methods and apparatus are described in their most basic form but steps may be added to or deleted from any of the methods and components may be added or subtracted from any of the described apparatus without departing from the basic scope of the present invention. It will be apparent to those skilled in the art that many further modifications and adaptations may be made. The particular embodiments are not provided to limit the invention but to illustrate it. The scope of the present invention is not to be determined by the specific examples provided above but only by the claims below.

1. A method comprising:
 - receiving a packet from a source component, the packet containing Protocol Layer;
 - attaching a Transport Layer to the packet; and
 - sending the packet across a component communications interface to a second component, the packet containing the Transport Layer and the Protocol Layer.
2. The method of claim 1, further comprising:
 - waiting to receive an acknowledgment of the sent packet across the component communications interface Transport Layer; and
 - if no acknowledgment is received after a specific time interval, then resending the packet.
3. The method of claim 1, further comprising if a negative acknowledgment is received across the component communications interface Transport Layer, then resending the packet.
4. The method of claim 1, further comprising, after resending the packet, repeating waiting to receive an acknowledgment of the resent packet and resending the packet until either an acknowledgement is received or the packet has been resent a specific number of times.
5. The method of claim 1, wherein attaching a Transport Layer comprises attaching an identifier of the source of the packet.
6. The method of claim 1, wherein attaching a Transport Layer comprises attaching a sequence number for the packet.
7. The method of claim 6, wherein the sequence number is based on the sequence of packets received from the source component.
8. The method of claim 1, further comprising buffering the received packet.
9. The method of claim 8, further comprising receiving an acknowledgement of the sent packet and then removing the packet from the buffer without resending.
10. A machine-readable medium comprising data that when operated on by the machine cause the machine to perform operations comprising:
 - receiving a packet from a source component, the packet having destination address;
 - determining if there is a Transport Layer enabled component in as path to the destination address;
 - if there is a Transport Layer enabled component, then attaching a Transport Layer to the packet; and
 - sending the packet to the destination address on the path through the Transport Layer enabled component.
11. The medium of claim 10, wherein the operations further comprise buffering the packet until an acknowledgment is received from the Transport Layer-enabled device.
12. The medium of claim 10, wherein the Transport Layer comprises a sequence number and an identification of the source component.
13. A method comprising:
 - receiving a packet across a component communications interface, the packet containing Transport Layer and a Protocol Layer;
 - determining whether the packet is valid;
 - if the packet is valid, then using the Transport Layer, sending an acknowledgment to the sender, the acknowledgment being in a Transport Layer of the acknowledgment;
 - if the packet is valid, then removing the Transport Layer and forwarding the packet to a destination component.

14. The method of claim **13**, further comprising the packet is not valid, then sending a negative acknowledgement to the sender, the negative acknowledgment being in a Transport layer of the acknowledgment.

15. The method of claim **13**, further comprising checking a sequence number of the packet Transport Layer to determine if the packet is a duplicate and if the packet is duplicate, then dropping the packet.

16. The method of claim **13**, wherein using the Transport Layer comprises determining an address for the sender using the Transport Layer.

17-20. (canceled)

* * * * *