



(51) International Patent Classification:

H04L 25/20 (2006.01) H04L 12/861 (2013.01)
H04L 25/03 (2006.01) H04L 29/10 (2006.01)

(21) International Application Number:

PCT/US2016/033690

(22) International Filing Date:

23 May 2016 (23.05.2016)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

62/184,830 25 June 2015 (25.06.2015) US
15/079,890 24 March 2016 (24.03.2016) US

(71) Applicant: **INTEL CORPORATION** [US/US]; 2200
Mission College Boulevard, Santa Clara, CA 95054-1549
(US).

(72) Inventors: **WIG, Timothy Donald**; 23 Valentine Road,
Northborough, MA 01532-1338 (US). **KROOSWYK,
Steven K.**; 203 Naples Avenue, Cayce, SC 29033 (US).
WELLS, Marc Lawrence; 5550 NE Prescott Street, Port-
land, OR 97218 (US).

(74) Agent: **KOMENDA, J., Kyle**; Patent Capital Group, C/o
CPA Global, 900 Second Avenue South, Suite 600, Min-
neapolis, 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, IR, IS, JP, KE, KG, KN, KP, KR,
KZ, LA, LC, LK, LR, LS, 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, SA, 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.

(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, ST, 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, KM, ML, MR, NE, SN, TD, TG).


Declarations under Rule 4.17:

- as to applicant's entitlement to apply for and be granted a
patent (Rule 4.17(ii))
- as to the applicant's entitlement to claim the priority of the
earlier application (Rule 4.17(iii))

Published:

- with international search report (Art. 21(3))

(54) Title: **SIDEBAND CONDUCTOR RESONANCE MITIGATION**



(57) Abstract: An apparatus comprising includes a first pair of conductors to carry differential signals, at least one ground conductor neighboring the first pair of conductors, the ground conductor to be connected to a ground plane, and at least one particular conduct-
or to carry sideband signals. The particular conductor is to be connected to a ground plane via a resonance mitigation circuit, and the resonance mitigation circuit comprises a resistor.

SIDEBAND CONDUCTOR RESONANCE MITIGATION

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims benefit to and priority to U.S. Provisional Patent Application Serial No. 62/184,830 filed June 25, 2015 entitled “SIDEBAND CONDUCTOR RESONANCE MITIGATION” and U.S. Nonprovisional (Utility) Patent Application Serial No. 15/079,890 filed March 24, 2016 entitled “SIDEBAND CONDUCTOR RESONANCE MITIGATION”, which are incorporated by reference herein in their entireties.

FIELD

[0002] This disclosure pertains to computing system, and in particular (but not exclusively) connections to facilitate electronic communication.

BRIEF DESCRIPTION OF THE DRAWINGS

[0003] FIG. 1 illustrates an embodiment of a computing system including an interconnect architecture.

[0004] FIG. 2 illustrates an embodiment of a interconnect architecture including a layered stack.

[0005] FIG. 3 illustrates an embodiment of a request or packet to be generated or received within an interconnect architecture.

[0006] FIG. 4 illustrates an embodiment of a transmitter and receiver pair for an interconnect architecture.

[0007] FIG. 5 illustrates an example add-in card device.

[0008] FIG. 6 illustrates a representation of an example connector interface including differential signal pairs and ground conductors.

[0009] FIGS. 7A-7G illustrates simplified diagrams of example connectors.

[0010] FIG. 8 illustrates a representation of a pinout of an example connector including one or more sideband conductors.

[0011] FIG. 9 illustrates a graph illustrating crosstalk characteristics of example connectors.

[0012] FIG. 10 illustrates a representation of an example card edge pinfield including one or more sideband conductors.

[0013] FIG. 11 illustrates a representation of an example connector pinfield including one or more sideband conductors.

[0014] FIG. 12 illustrates a representation of an example circuit board connection scheme including one or more sideband conductors.

[0015] FIG. 13 illustrates an embodiment of a block diagram for a computing system including a multicore processor.

[0016] FIG. 14 illustrates another embodiment of a block diagram for a computing system.

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 measurements/heights, specific processor pipeline stages and operation etc. in order to provide a thorough understanding of the present invention. It will be apparent, however, to one skilled in the art that these specific details need not be employed to practice the present invention. In other instances, well known components or methods, such as specific and alternative processor architectures, specific logic circuits/code for described algorithms, specific firmware code, specific 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 haven't been described in detail in order to avoid unnecessarily obscuring the present invention.

[0018] Although the following embodiments may be described with reference to energy conservation and energy efficiency 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 better energy efficiency and energy conservation. For example, the disclosed embodiments are not limited to desktop computer systems or Ultrabooks™. And may be also used in other devices, such as handheld devices, 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. 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.

[0019] As computing systems are advancing, the components therein are becoming more complex. As a result, the interconnect architecture to couple and communicate between the components is also increasing in complexity to ensure bandwidth requirements are met for optimal component operation. Furthermore, different market segments demand different aspects of interconnect architectures to suit the market's needs. For example, servers require higher performance, while the mobile ecosystem is sometimes able to sacrifice overall performance for power savings. Yet, it's a singular purpose of most fabrics to provide highest possible performance with maximum power saving. Below, a number of interconnects are discussed, which would potentially benefit from aspects of the invention described herein.

[0020] One interconnect fabric architecture includes the Peripheral Component Interconnect (PCI) Express (PCIe) architecture. A primary goal of PCIe is to enable components and devices from different vendors to inter-operate in an open architecture, spanning multiple market segments; Clients (Desktops and Mobile), Servers (Standard and Enterprise), and Embedded and Communication devices. PCI Express is a high performance, general purpose I/O interconnect defined for a wide variety of future computing and communication platforms. Some PCI attributes, such as its usage model, load-store architecture, and software interfaces, have been maintained through its revisions, whereas previous parallel bus implementations have been replaced by a highly scalable, fully serial interface. The more recent versions of PCI Express take advantage of advances in point-to-

point interconnects, Switch-based technology, and packetized protocol to deliver new levels of performance and features. Power Management, Quality Of Service (QoS), Hot-Plug/Hot-Swap support, Data Integrity, and Error Handling are among some of the advanced features supported by PCI Express.

[0021] Referring to FIG. 1, an 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 includes 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.

[0022] 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.

[0023] In one embodiment, controller hub 115 is a root hub, root complex, or root controller in a Peripheral Component Interconnect Express (PCIe or 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, i.e. 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.

[0024] 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, 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.

[0025] 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 PCIe to PCI/PCI-X bridge to support legacy or other version PCI devices. Endpoint devices in PCIe are often classified as legacy, PCIe, or root complex integrated endpoints.

[0026] Graphics accelerator 130 is also 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 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.

[0027] Turning to FIG. 2 an embodiment of a layered protocol stack is illustrated. Layered protocol stack 150 includes any form of a layered communication stack, such as a Quick Path Interconnect (QPI) stack, a PCIe stack, a next generation high performance computing interconnect stack, or other layered stack. Although the discussion immediately below in reference to Figures 1-4 are in relation to a PCIe stack, the same concepts may be applied to other interconnect stacks. In one embodiment, protocol stack 150 is a PCIe protocol stack including transaction layer 155, link layer 210, and physical layer 220. An interface, such as interfaces 117, 118, 121, 122, 126, and 131 in Figure 1, may be represented as communication protocol stack 150. Representation as a communication protocol stack may also be referred to as a module or interface implementing/including a protocol stack.

[0028] PCI Express uses packets to communicate information between components. Packets are formed in the Transaction Layer 155 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 necessary 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 155 of the receiving device.

Transaction Layer

[0029] In one embodiment, transaction layer 155 is to 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 155 is the assembly and disassembly of packets (i.e., transaction layer packets, or TLPs). The translation layer 155 typically manages credit-base flow control for TLPs. PCIe implements split transactions, 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.

[0030] In addition PCIe utilizes credit-based flow control. In this scheme, a device advertises an initial amount of credit for each of the receive buffers in Transaction Layer 155. An external device at the opposite end of the link, such as controller hub 115 in Figure 1, counts 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. An advantage of a credit scheme is that the latency of credit return does not affect performance, provided that the credit limit is not encountered.

[0031] In one embodiment, four transaction address spaces 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 are used to access configuration space of the PCIe devices. Transactions to the configuration space include read requests and write requests. Message transactions are defined to support in-band communication between PCIe agents.

[0032] Therefore, in one embodiment, transaction layer 155 assembles packet header/payload 156. Format for current packet headers/payloads may be found in the PCIe specification at the PCIe specification website.

[0033] Quickly referring to Figure 3, an embodiment of a PCIe transaction descriptor is illustrated. In one embodiment, transaction descriptor 300 is a mechanism for carrying transaction information. In this regard, transaction descriptor 300 supports identification of transactions in a system. Other potential uses include tracking modifications of default transaction ordering and association of transaction with channels.

[0034] Transaction descriptor 300 includes global identifier field 302, attributes field 304 and channel identifier field 306. In the illustrated example, global identifier field 302 is depicted comprising local transaction identifier field 308 and source identifier field 310. In one embodiment, global transaction identifier 302 is unique for all outstanding requests.

[0035] According to one implementation, local transaction identifier field 308 is a field generated by a requesting agent, and it is unique for all outstanding requests that require a completion for that requesting agent. Furthermore, in this example, source identifier 310 uniquely identifies the requestor agent within a PCIe hierarchy. Accordingly, together with source ID 310, local transaction identifier 308 field provides global identification of a transaction within a hierarchy domain.

[0036] Attributes field 304 specifies characteristics and relationships of the transaction. In this regard, attributes field 304 is potentially used to provide additional information that allows modification of the default handling of transactions. In one embodiment, attributes field 304 includes priority field 312, reserved field 314, ordering field 316, and no-snoop field 318. Here, priority sub-field 312 may be modified by an initiator to assign a priority to the transaction. Reserved attribute field 314 is left reserved for future, or vendor-defined usage. Possible usage models using priority or security attributes may be implemented using the reserved attribute field.

[0037] In this example, ordering attribute field 316 is used to supply optional information conveying the type of ordering that may modify default ordering rules. According to one example implementation, an ordering attribute of "0" denotes default ordering rules are to apply, wherein an ordering attribute of "1" denotes relaxed ordering, wherein writes can pass writes in the same direction, and read completions can pass writes in the same direction. Snoop attribute field 318 is utilized to determine if transactions are

snooped. As shown, channel ID Field 306 identifies a channel that a transaction is associated with.

Link Layer

[0038] Link layer 210, also referred to as data link layer 210, acts as an intermediate stage between transaction layer 155 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 a link. One side of the Data Link Layer 210 accepts TLPs assembled by the Transaction Layer 155, 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.

Physical Layer

[0039] In one embodiment, 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 includes 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.

[0040] 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 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.

[0041] As stated above, although transaction layer 155, link layer 210, and physical layer 220 are discussed in reference to a specific embodiment of a PCIe protocol stack, a layered protocol stack is not so limited. In fact, any layered protocol may be included/implemented. As an example, an port/interface that is represented as a layered protocol includes: (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 common standard interface (CSI) layered protocol is utilized.

[0042] Referring next to Figure 4, an embodiment of a PCIe serial point to point fabric is illustrated. Although an embodiment of a PCIe serial point-to-point link is illustrated, a serial point-to-point link is not so limited, as it includes any transmission path for transmitting serial data. In the embodiment shown, a basic PCIe link includes two, low-voltage, differentially driven signal pairs: a transmit pair 406/411 and a receive pair 412/407. Accordingly, device 405 includes transmission logic 406 to transmit data to device 410 and receiving logic 407 to receive data from device 410. In other words, two transmitting paths, i.e. paths 416 and 417, and two receiving paths, i.e. paths 418 and 419, are included in a PCIe link.

[0043] 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 405 and device 410, is referred to as a link, such as link 415. 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.

[0044] A differential pair refers to two transmission paths, such as lines 416 and 417, to transmit differential signals. As an example, when line 416 toggles from a low voltage level to a high voltage level, i.e. a rising edge, line 417 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, etc. This allows for better timing window, which enables faster transmission frequencies.

[0045] FIG. 5 illustrates an example add-in card that can be mounted to a baseboard 500 so as to communicatively couple to the baseboard by one of a set of conductors. The add-card can be configured for a specific purpose, so as to extend functionality of the computer to which it is connected (e.g., through the computer's baseboard). In one example implementation, the add-in card comprises a device compatible with a PCI-based interconnect protocol, such as PCIe. FIG. 6 illustrates the use of differential pairs in PCIe

connector 600, shown as an add-in card (AIC) mounted to a baseboard. All contacts can be electrically isolated from one another by airspace and a plastic connector shell (not shown). Conductors 610 and 612 can be utilized to represent a first differential pair, capable of carrying equal and opposite currents to balance the signal integrity effects. Likewise, in this example, conductors 620 and 622 represent a second differential pair. Additional conductors can be provided and utilized for ground, sideband, and other terminations. For instance, ground conductor pairs 630 and 632, 640 and 642, and 650 and 652 can be provided, which are joined only at the printed circuit board (PCB) level (baseboard and AIC). In some implementations, a ground conductor (e.g., 642) adjacent to a signal conductor (e.g., 610) of a differential pair may balance some of the return current (in addition to 612). (It should be noted that in other implementations the signals are not differential. This, however, does not necessarily limit the applicability of the concepts described herein, as the above concepts can be applied to other systems, including signals that are non-differential signals.)

[0046] Continuing with the example of FIG. 6, among the conductors, several sideband pins may be provided among the high speed differential pairs and grounds. Sideband pins can be utilized, in some implementations, for low speed or DC signals. In some cases, sideband pins may be unused. In some instances, such as in some PCIe-compliant connectors, sideband conductors can exhibit a resonance due to the signaling on nearby high speed lanes (or conductors). This resonance can generate a strong crosstalk peak, such as a peak falling at about 5 GHz, in the middle of the 0-8 GHz bandwidth of Gen4. This resonance can also cause a corresponding dropout in insertion loss at the same frequency. The resonance is a result of electromagnetic coupling among the conducting pinfield vias, connector pins and contacts, and add-in card (AIC) edge fingers, combined with the fact that these conductors are typically unterminated, or terminated in a highly mismatched impedance. This resonance may be present in both surface mount and thru-hole versions of the PCIe connector, and in x1, x4, x8, and x16 widths, among other examples.

[0047] In some implementations, a mechanism can be provided for reducing crosstalk in connectors. More particularly, in some implementations of high speed interconnects, such as the PCI Express (PCIe) connector interface, crosstalk is manifest stemming from a resonance in unterminated or poorly terminated sideband pins. This crosstalk can be mitigated by attaching an AC-coupled termination network that will effectively damp the resonance without disrupting the assigned electrical function of the connector pin. Such

circuitry may be applied to other connectors and, potentially, socket pins, among other examples.

[0048] In earlier, lower-bandwidth generations of PCIe (and other interconnects) these crosstalk issues were less troublesome. However, with the higher bandwidth requirements of more recent versions (e.g., 16GT/s PCIe Gen4 channels), increased sensitivity to these crosstalk mechanisms is observed. While only a subset of the connector's high speed lanes may be affected by this resonance, these vulnerable lanes serve to limit the overall link bandwidth. Further, in the example of PCIe, these resonance issues are present in most PCIe connector interfaces, including the x1, x4, x8, and x16 styles. While degraded channel performance can be partially offset using other methods, such as shortening the channel length, or replacing connectors with much more expensive backplane connectors, such alternatives can limit the scope of the channel's utility and may not be compatible with some applications. For instance, replacing edge connectors with higher performance backplane connectors can make it impossible to realize interconnects that guarantee backward compatibility (e.g., backward compatibility among PCIe Generation 1-2-3-4 baseboards and add-in cards), among other examples.

[0049] In one implementation, a passive series resistor-capacitor (RC) termination network can be applied between the PCIe connector sideband pins and ground. While the present disclosure discusses several examples where the network only includes a simple termination R-C network, it should be appreciated that more complex AC termination networks are possible, using more than a single resistor and capacitor, for example. Further, it may be observed that termination to DC network other than Ground, such as a 3.3 Volt DC plane, may be possible, among other alternative implementations. Regardless, such networks can permit absorption of coupled energy to provide damping of resonant energy at higher frequencies (generally >100MHz) on sideband channels without perturbing the operation of the sideband signals, which would otherwise result from excessive low frequency loading. In one implementation, the capacitor can be sized appropriately (e.g., 1pF) to provide a capacitive reactance sufficiently low to pass the undesired coupled energy, while largely blocking low frequency and DC traffic. Such circuit enhancements, among other example advantages, may minimize crosstalk stemming from resonance manifesting on high speed links. Such solutions might also improve other metrics of channel quality such as insertion loss and return loss, among other examples.

[0050] FIGS. 7A-7G illustrate block diagrams representing example implementations of a multi-conductor connector that includes high speed differential pairs, grounds, and one or more other conductors that may be used to carry sideband signals (i.e., “sideband conductors”). For instance, FIG. 7A is a simplified block diagram 700a illustrating conductors of a conventional connector. In this example, the connector can include Differential Pair 0 and Differential Pair 1, each designated for carrying differential, high speed data signals. In FIG. 7A, the terminals of the sideband conductor are open-circuited permitting multiple reflections that lead to a resonance, whose frequency may be dictated by the length of the floating conductor. As noted above, this resonance can, in turn, manifest as crosstalk in neighboring differential conductors (e.g., conductors of Differential Pair 0 and Differential Pair 1), even when a ground conductor is positioned at least partially between the sideband conductor and high speed signaling conductors (e.g., Differential Pair 0 and Differential Pair 1).

[0051] Turning to FIG. 7B, a simplified block diagram 700b is shown illustrating an implementation including circuitry 705b to remediate potential crosstalk resonance at the sideband conductors, wherein the sideband conductor is resistively terminated to ground to damp the resonance. In this example, the implementation of crosstalk remediation circuitry 705b effectively provides DC coupling. In some cases, such a DC-coupled solution may be limited in its application (e.g., cannot be applied to PCIe sideband signals). Accordingly, the simplified block diagram 700c of FIG. 7C illustrates another implementation including circuitry 705c to remediate potential crosstalk resonance at the sideband conductors. For instance, in the implementation of FIG. 7C, series capacitance is added to the solution of FIG. 7B to provide both AC decoupling and resistive termination through an RC remediation circuit.

[0052] Remediation circuitry can effectively damp resonance even where terminations of the sideband conductors are connected to ground through remediation circuitry at only one end of a connection. FIGS. 7D-7F are simplified block diagrams 700d-f illustrating additional implementations of sideband conductors provided with remediation circuitry, such as introduced in the examples of FIGS. 7B-7C. For instance, FIG. 7D shows a first device (e.g., an add-in card) connected to a second device (e.g., a board) via a connector 710. The first device can include one or more sideband conductors terminated to ground via remediation circuitry 705d that includes a resistor and capacitor in series, such as illustrated in

FIG. 7D. The second device can also be configured to terminate a corresponding sideband conductor to ground via a resistor and capacitor-based remediation circuit (e.g., 715d). This can represent an ideal termination of sideband conductors, with remediation circuitry (705d, 715d) at both ends of a connection.

[0053] As introduced above, remediation circuitry can be effective even when only applied to a sideband conductor at one end of a connection. For instance, FIG. 7E represents an alternative implementation of the example of FIG. 7D where remediation circuitry (e.g., 715e) is instead applied to a sideband conductor at a second device, but not at the corresponding sideband conductor of the first device to which the second device is connected (e.g., via a connector). Alternatively, the first device can apply the remediation circuitry (e.g., 705f) while the second device does not, as illustrated in FIG. 7F.

[0054] Turning to the example FIG. 7G, yet another example implementation is illustrated of remediation circuitry for provision on one or both terminations of a sideband conductor. In the particular example of FIG. 7G, at one side of a connection (i.e., at a second device), the remediation circuitry 720 for a sideband conductor includes a tee-connection, allowing the sideband conductor to implement an assigned circuit function, while remediation is provided through a resistor or resistor-capacitor termination forming an additional branch to the existing circuit routing that does not substantially interfere with the assigned function. The assigned function may involve termination at another component or device, among other examples. As further illustrated in the example of FIG. 7G, different types of remediation circuitry can be provided at each end of a sideband conductor, with remediation circuitry 705g (e.g., not with a tee-connection) provided at the termination of a first device, among other examples. In some cases, the impedance match achieved with the addition of AC termination networks will not be ideal. It is enough that the AC terminations substantially damp each reflection; even with an imperfect match, the presence of the AC termination networks (e.g., 705b-d,f,g, 715d-e, 720) will quickly damp any resonance stemming from energy exchanged with other coupled conductors.

[0055] Note that the simplified planar representations illustrated in FIGS. 7A-7G may not adequately represent the actual connector geometry; the baseboard, connector, and add-in card geometry suffers greater sideband conductor coupling than the 2D illustrations might imply. Further, as noted above, the principles described above can be applied to connections, devices, and systems compliant with a PCIe-based specification. For instance, in the case of

PCIe Gen4, data rates may meet or exceed 16GT/s. To support PCIe Gen4 traffic the frequency bandwidth may generally be considered to span the 0-8GHz range. This exceeds the bandwidth of previous PCIe versions' connector interface, when the PCIe Gen3 Card Electromechanical Spec (CEM Spec) and common PCIe Gen 1-2-3 PCB design techniques are applied. In the pursuit of Gen4 readiness, numerous channel impairments that limit the performance of the connector have been identified (e.g., through electromagnetic component simulation as well as simulation at the channel level, and subsequently confirmed in hardware, using vector network analysis). Many of these effects have either been previously unknown or not addressed due in large part to the fact that many of these signal integrity problems did not affect Gen1-2-3 channels very strongly.

[0056] FIG. 8 illustrates a representation of at least a portion 800 of a pin field layout of a an example (e.g., PCIe) connector including one or more sideband conductors (e.g., 805a-f). Multiple high speed differential signaling pairs (e.g., Rx0, Rx1, Rx2, Tx0, Tx1, Tx2, etc.) can also be provided, as well as ground conductors (e.g., 810a-e). A variety of pinfield layouts can be provided on a connector. In some examples, such as that illustrated in FIG. 8, ground conductors (e.g., 810a-e) can be positioned relative to different signaling pairs and sideband conductors to assist in absorbing signals from neighboring conductors (e.g., high speed differential signaling pairs and sideband conductors) to assist in preventing crosstalk. For instance, ground conductors may be provided between two potentially competing conductors, such as ground conductor 810a between Tx0 and sideband 805b, or ground conductor 810e between Tx3 and sideband conductors 805e and 805f, etc. Resonance appearing on unterminated sideband conductors (e.g., 805e) may nonetheless cause crosstalk to manifest on a nearby high speed signaling conductor (e.g., in Tx3), despite the presence (and partial assistance) of ground conductor. Provisioning one or more (or all) of the sideband conductors (e.g., 805a-f) with remediation circuitry at the sidebands' terminations (such as described above) can remediate such resonance and protect signaling on nearby conductors from interference emanating on the sideband conductor, all while allowing (e.g., lower frequency or DC) sideband signals to be successfully communicated using the sideband conductors.

[0057] As noted above, the principles described above can be applied in PCIe-based connectors and interconnects to improve signal quality. For instance, high speed PCIe Gen4 traffic may be particular susceptible to channel impediments caused by sideband channels of

PCIe connectors. Indeed, sideband conductor resonance can be a prominent interference mechanism, manifested as a strong crosstalk peak that is a consequence of having unterminated sideband pins. For instance, FIG. 9 shows a graph 900 illustrating a comparison of the crosstalk among high speed pairs adjacent to sideband pins that are terminated either with an open circuit or a matched resistance. The crosstalk peak of the dashed trace 905 near 5 GHz is a consequence of the undamped resonance in the unterminated baseline configuration. The solid line 910 illustrates the crosstalk reduction when AC termination is applied (such as with the example remediation circuitry discussed above), demonstrating the effectiveness of such sideband termination solutions.

[0058] As discussed above, a remediation circuit can include either a resistor (introducing impedance to the sideband conductor) alone or a resistor and capacitor in series. In the case of a remediation circuit with a resistor only, a simple resistive termination may disrupt the assigned function of the sideband pins, which prompted this invention. If left unconnected, the sideband pins can function as undamped, high-Q resonators that absorb and emit energy from/to neighboring pins. A resistor connected to one or both ends of the unterminated networks will absorb (or damp) the energy that would otherwise be re-emitted, to manifest as crosstalk peaks or insertion loss dropouts. The sideband pins that are currently assigned carry only very low frequency or DC traffic, so an appropriately sized termination resistor alone (perhaps 50Ω) would heavily load these networks, and would interfere with the function of the pins. Adding a capacitor in series with the resistor can cure this disruption while preserving the remedial properties of the circuit.

[0059] A series RC circuit network can prevent low frequency loading that would occur with the resistor alone. For instance, for the add-in card, the termination may be placed adjacent to the gold edge finger region that engages the baseboard connector. There is ample space for placement of a surface mount resistor and capacitor in this region, as the signals fan out to meet the relatively widely spaced edge fingers. For instance, as illustrated in the representation 1000 of a printed circuit board (PCB) diagram shown in FIG. 10, it can be seen that sufficient PCB area exists to add a capacitor to an existing resistive termination network, for a PCIe add-in card. Similarly, for a PCIe baseboard, the termination may be placed adjacent to the pinfield, similar to the termination shown in the example layout illustrated in FIG. 11, among other examples.

[0060] In some implementations, such as PCIe Gen4, the sideband signals that would benefit from termination lie in the high speed side of the edge connector, e.g., at pin positions A12/B12 and above. Other sideband pins, such as those assigned to JTAG, in pin positions A1 to A11 and B1 to B11, may not benefit as strongly from the RC termination. In one example, sideband pins that are provided with AC matched termination (e.g., using implementations of the remediation circuitry described herein) may include such PCIe sideband channels as Reserved, Present# (Prsnt), Clock Request, Power Brake, etc. (such as shown in FIG. 8). There were originally seven Reserved pins, but the definition of Power Brake and Clock Request leave five remaining Reserved pins. A reasonable concern would be the forward compatibility of AC termination of the undefined Reserved pins, since their bandwidth requirements may extend into the range where the AC termination provides stronger resistive loading. It has been observed that the insertion loss and crosstalk performance of the sideband signal paths would limit the bandwidth to perhaps 100MHz, with relatively low edge rates. Based upon this, it can be concluded that an AC termination scheme may not disturb any future, as-yet undefined, sideband signals (e.g., for Reserved pins), among other example considerations and examples.

[0061] Turning to FIG. 12, a simplified block diagram 1200 is shown illustrating an example PCI express x16 add-in card, depicting the B side of the card only. All sideband signals in this example are shown terminated with AC matched termination. The low speed and power pins B1-B11 are shown without additional AC matched termination (e.g., using implementations of the remediation circuitry described herein). For example, the CLKREQ# (Clock Request) pin is shown connected as a Tee connection, one PRESENT2# pin, B81, is also shown as a Tee connection, as well as the PWRBRK# (Power Brake) pin, B30 (similar to that illustrated in the example of FIG. 7G). The pinfield on the “A” side of the card, comprising pins A1 to A82, are not shown, but it is understood that the A side sideband pins would be terminated in a similar fashion (or using another implementation of remediation circuitry), among other examples.

[0062] Note that the apparatus', methods', and systems described above may be implemented in any electronic device or system as aforementioned. As specific illustrations, the figures below provide exemplary systems for utilizing the invention as described herein. As the systems below are described in more detail, a number of different interconnects are disclosed, described, and revisited from the discussion above. And as is readily apparent, the

advances described above may be applied to any of those interconnects, fabrics, or architectures.

[0063] Referring to FIG. 13, an embodiment of a block diagram for a computing system including a multicore processor is depicted. Processor 1300 includes any processor or processing device, such as a microprocessor, an embedded processor, a digital signal processor (DSP), a network processor, a handheld processor, an application processor, a co-processor, a system on a chip (SOC), or other device to execute code. Processor 1300, in one embodiment, includes at least two cores—core 1301 and 1302, which may include asymmetric cores or symmetric cores (the illustrated embodiment). However, processor 1300 may include any number of processing elements that may be symmetric or asymmetric.

[0064] In one embodiment, a processing element refers to hardware or logic to support a software thread. Examples of hardware processing elements include: a thread unit, a thread slot, a thread, a process unit, a context, a context unit, a logical processor, a hardware thread, a core, and/or any other element, which is capable of holding a state for a processor, such as an execution state or architectural state. In other words, a processing element, in one embodiment, refers to any hardware capable of being independently associated with code, such as a software thread, operating system, application, or other code. A physical processor (or processor socket) typically refers to an integrated circuit, which potentially includes any number of other processing elements, such as cores or hardware threads.

[0065] A core often refers to logic located on an integrated circuit capable of maintaining an independent architectural state, wherein each independently maintained architectural state is associated with at least some dedicated execution resources. In contrast to cores, a hardware thread typically refers to any logic located on an integrated circuit capable of maintaining an independent architectural state, wherein the independently maintained architectural states share access to execution resources. As can be seen, when certain resources are shared and others are dedicated to an architectural state, the line between the nomenclature of a hardware thread and core overlaps. Yet often, a core and a hardware thread are viewed by an operating system as individual logical processors, where the operating system is able to individually schedule operations on each logical processor.

[0066] Physical processor 1300, as illustrated in FIG. 13, includes two cores—core 1301 and 1302. Here, core 1301 and 1302 are considered symmetric cores, i.e. cores with the same configurations, functional units, and/or logic. In another embodiment, core 1301

includes an out-of-order processor core, while core 1302 includes an in-order processor core. However, cores 1301 and 1302 may be individually selected from any type of core, such as a native core, a software managed core, a core adapted to execute a native Instruction Set Architecture (ISA), a core adapted to execute a translated Instruction Set Architecture (ISA), a co-designed core, or other known core. In a heterogeneous core environment (i.e. asymmetric cores), some form of translation, such as a binary translation, may be utilized to schedule or execute code on one or both cores. Yet to further the discussion, the functional units illustrated in core 1301 are described in further detail below, as the units in core 1302 operate in a similar manner in the depicted embodiment.

[0067] As depicted, core 1301 includes two hardware threads 1301a and 1301b, which may also be referred to as hardware thread slots 1301a and 1301b. Therefore, software entities, such as an operating system, in one embodiment potentially view processor 1300 as four separate processors, i.e., four logical processors or processing elements capable of executing four software threads concurrently. As alluded to above, a first thread is associated with architecture state registers 1301a, a second thread is associated with architecture state registers 1301b, a third thread may be associated with architecture state registers 1302a, and a fourth thread may be associated with architecture state registers 1302b. Here, each of the architecture state registers (1301a, 1301b, 1302a, and 1302b) may be referred to as processing elements, thread slots, or thread units, as described above. As illustrated, architecture state registers 1301a are replicated in architecture state registers 1301b, so individual architecture states/contexts are capable of being stored for logical processor 1301a and logical processor 1301b. In core 1301, other smaller resources, such as instruction pointers and renaming logic in allocator and renamer block 1330 may also be replicated for threads 1301a and 1301b. Some resources, such as re-order buffers in reorder/retirement unit 1335, ILTB 1320, load/store buffers, and queues may be shared through partitioning. Other resources, such as general purpose internal registers, page-table base register(s), low-level data-cache and data-TLB 1315, execution unit(s) 1340, and portions of out-of-order unit 1335 are potentially fully shared.

[0068] Processor 1300 often includes other resources, which may be fully shared, shared through partitioning, or dedicated by/to processing elements. In FIG. 13, an embodiment of a purely exemplary processor with illustrative logical units/resources of a processor is illustrated. Note that a processor may include, or omit, any of these functional

units, as well as include any other known functional units, logic, or firmware not depicted. As illustrated, core 1301 includes a simplified, representative out-of-order (OOO) processor core. But an in-order processor may be utilized in different embodiments. The OOO core includes a branch target buffer 1320 to predict branches to be executed/taken and an instruction-translation buffer (I-TLB) 1320 to store address translation entries for instructions.

[0069] Core 1301 further includes decode module 1325 coupled to fetch unit 1320 to decode fetched elements. Fetch logic, in one embodiment, includes individual sequencers associated with thread slots 1301a, 1301b, respectively. Usually core 1301 is associated with a first ISA, which defines/specifies instructions executable on processor 1300. Often machine code instructions that are part of the first ISA include a portion of the instruction (referred to as an opcode), which references/specifies an instruction or operation to be performed. Decode logic 1325 includes circuitry that recognizes these instructions from their opcodes and passes the decoded instructions on in the pipeline for processing as defined by the first ISA. For example, as discussed in more detail below decoders 1325, in one embodiment, include logic designed or adapted to recognize specific instructions, such as transactional instruction. As a result of the recognition by decoders 1325, the architecture or core 1301 takes specific, predefined actions to perform tasks associated with the appropriate instruction. It is important to note that any of the tasks, blocks, operations, and methods described herein may be performed in response to a single or multiple instructions; some of which may be new or old instructions. Note decoders 1326, in one embodiment, recognize the same ISA (or a subset thereof). Alternatively, in a heterogeneous core environment, decoders 1326 recognize a second ISA (either a subset of the first ISA or a distinct ISA).

[0070] In one example, allocator and renamer block 1330 includes an allocator to reserve resources, such as register files to store instruction processing results. However, threads 1301a and 1301b are potentially capable of out-of-order execution, where allocator and renamer block 1330 also reserves other resources, such as reorder buffers to track instruction results. Unit 1330 may also include a register renamer to rename program/instruction reference registers to other registers internal to processor 1300. Reorder/retirement unit 1335 includes components, such as the reorder buffers mentioned above, load buffers, and store buffers, to support out-of-order execution and later in-order retirement of instructions executed out-of-order.

[0071] Scheduler and execution unit(s) block 1340, in one embodiment, includes a scheduler unit to schedule instructions/operation on execution units. For example, a floating point instruction is scheduled on a port of an execution unit that has an available floating point execution unit. Register files associated with the execution units are also included to store information instruction processing results. Exemplary execution units include a floating point execution unit, an integer execution unit, a jump execution unit, a load execution unit, a store execution unit, and other known execution units.

[0072] Lower level data cache and data translation buffer (D-TLB) 1350 are coupled to execution unit(s) 1340. The data cache is to store recently used/operated on elements, such as data operands, which are potentially held in memory coherency states. The D-TLB is to store recent virtual/linear to physical address translations. As a specific example, a processor may include a page table structure to break physical memory into a plurality of virtual pages.

[0073] Here, cores 1301 and 1302 share access to higher-level or further-out cache, such as a second level cache associated with on-chip interface 1310. Note that higher-level or further-out refers to cache levels increasing or getting further way from the execution unit(s). In one embodiment, higher-level cache is a last-level data cache—last cache in the memory hierarchy on processor 1300—such as a second or third level data cache. However, higher level cache is not so limited, as it may be associated with or include an instruction cache. A trace cache—a type of instruction cache—instead may be coupled after decoder 1325 to store recently decoded traces. Here, an instruction potentially refers to a macro-instruction (i.e. a general instruction recognized by the decoders), which may decode into a number of micro-instructions (micro-operations).

[0074] In the depicted configuration, processor 1300 also includes on-chip interface module 1310. Historically, a memory controller, which is described in more detail below, has been included in a computing system external to processor 1300. In this scenario, on-chip interface 1310 is to communicate with devices external to processor 1300, such as system memory 1375, a chipset (often including a memory controller hub to connect to memory 1375 and an I/O controller hub to connect peripheral devices), a memory controller hub, a northbridge, or other integrated circuit. And in this scenario, bus 1305 may include any known interconnect, such as multi-drop bus, a point-to-point interconnect, a serial interconnect, a parallel bus, a coherent (e.g. cache coherent) bus, a layered protocol architecture, a differential bus, and a GTL bus.

[0075] Memory 1375 may be dedicated to processor 1300 or shared with other devices in a system. Common examples of types of memory 1375 include DRAM, SRAM, non-volatile memory (NV memory), and other known storage devices. Note that device 1380 may include a graphic accelerator, processor or card coupled to a memory controller hub, data storage coupled to an I/O controller hub, a wireless transceiver, a flash device, an audio controller, a network controller, or other known device.

[0076] Recently however, as more logic and devices are being integrated on a single die, such as SOC, each of these devices may be incorporated on processor 1300. For example in one embodiment, a memory controller hub is on the same package and/or die with processor 1300. Here, a portion of the core (an on-core portion) 1310 includes one or more controller(s) for interfacing with other devices such as memory 1375 or a graphics device 1380. The configuration including an interconnect and controllers for interfacing with such devices is often referred to as an on-core (or un-core configuration). As an example, on-chip interface 1310 includes a ring interconnect for on-chip communication and a high-speed serial point-to-point link 1305 for off-chip communication. Yet, in the SOC environment, even more devices, such as the network interface, co-processors, memory 1375, graphics processor 1380, and any other known computer devices/interface may be integrated on a single die or integrated circuit to provide small form factor with high functionality and low power consumption.

[0077] In one embodiment, processor 1300 is capable of executing a compiler, optimization, and/or translator code 1377 to compile, translate, and/or optimize application code 1376 to support the apparatus and methods described herein or to interface therewith. A compiler often includes a program or set of programs to translate source text/code into target text/code. Usually, compilation of program/application code with a compiler is done in multiple phases and passes to transform hi-level programming language code into low-level machine or assembly language code. Yet, single pass compilers may still be utilized for simple compilation. A compiler may utilize any known compilation techniques and perform any known compiler operations, such as lexical analysis, preprocessing, parsing, semantic analysis, code generation, code transformation, and code optimization.

[0078] Larger compilers often include multiple phases, but most often these phases are included within two general phases: (1) a front-end, i.e. generally where syntactic processing, semantic processing, and some transformation/optimization may take place, and

(2) a back-end, i.e. generally where analysis, transformations, optimizations, and code generation takes place. Some compilers refer to a middle, which illustrates the blurring of delineation between a front-end and back end of a compiler. As a result, reference to insertion, association, generation, or other operation of a compiler may take place in any of the aforementioned phases or passes, as well as any other known phases or passes of a compiler. As an illustrative example, a compiler potentially inserts operations, calls, functions, etc. in one or more phases of compilation, such as insertion of calls/operations in a front-end phase of compilation and then transformation of the calls/operations into lower-level code during a transformation phase. Note that during dynamic compilation, compiler code or dynamic optimization code may insert such operations/calls, as well as optimize the code for execution during runtime. As a specific illustrative example, binary code (already compiled code) may be dynamically optimized during runtime. Here, the program code may include the dynamic optimization code, the binary code, or a combination thereof.

[0079] Similar to a compiler, a translator, such as a binary translator, translates code either statically or dynamically to optimize and/or translate code. Therefore, reference to execution of code, application code, program code, or other software environment may refer to: (1) execution of a compiler program(s), optimization code optimizer, or translator either dynamically or statically, to compile program code, to maintain software structures, to perform other operations, to optimize code, or to translate code; (2) execution of main program code including operations/calls, such as application code that has been optimized/compiled; (3) execution of other program code, such as libraries, associated with the main program code to maintain software structures, to perform other software related operations, or to optimize code; or (4) a combination thereof.

[0080] Referring now to FIG. 14, shown is a block diagram of a second system 1400 in accordance with an embodiment of the present invention. As shown in FIG. 14, multiprocessor system 1400 is a point-to-point interconnect system, and includes a first processor 1470 and a second processor 1480 coupled via a point-to-point interconnect 1450. Each of processors 1470 and 1480 may be some version of a processor. In one embodiment, 1452 and 1454 are part of a serial, point-to-point coherent interconnect fabric, such as Intel's Quick Path Interconnect (QPI) architecture. As a result, the invention may be implemented within the QPI architecture.

[0081] While shown with only two processors 1470, 1480, 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.

[0082] Processors 1470 and 1480 are shown including integrated memory controller units 1472 and 1482, respectively. Processor 1470 also includes as part of its bus controller units point-to-point (P-P) interfaces 1476 and 1478; similarly, second processor 1480 includes P-P interfaces 1486 and 1488. Processors 1470, 1480 may exchange information via a point-to-point (P-P) interface 1450 using P-P interface circuits 1478, 1488. As shown in FIG. 14, IMCs 1472 and 1482 couple the processors to respective memories, namely a memory 1432 and a memory 1434, which may be portions of main memory locally attached to the respective processors.

[0083] Processors 1470, 1480 each exchange information with a chipset 1490 via individual P-P interfaces 1452, 1454 using point to point interface circuits 1476, 1494, 1486, 1498. Chipset 1490 also exchanges information with a high-performance graphics circuit 1438 via an interface circuit 1492 along a high-performance graphics interconnect 1439.

[0084] 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.

[0085] Chipset 1490 may be coupled to a first bus 1416 via an interface 1496. In one embodiment, first bus 1416 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.

[0086] As shown in FIG. 14, various I/O devices 1414 are coupled to first bus 1416, along with a bus bridge 1418 which couples first bus 1416 to a second bus 1420. In one embodiment, second bus 1420 includes a low pin count (LPC) bus. Various devices are coupled to second bus 1420 including, for example, a keyboard and/or mouse 1422, communication devices 1427 and a storage unit 1428 such as a disk drive or other mass storage device which often includes instructions/code and data 1430, in one embodiment. Further, an audio I/O 1424 is shown coupled to second bus 1420. 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. 14, a system may implement a multi-drop bus or other such architecture.

[0087] Computing systems can include various combinations of components. These components may be implemented as ICs, portions thereof, discrete electronic devices, or other modules, logic, hardware, software, firmware, or a combination thereof adapted in a computer system, or as components otherwise incorporated within a chassis of the computer system. However, it is to be understood that some of the components shown may be omitted, additional components may be present, and different arrangement of the components shown may occur in other implementations. As a result, the invention described above may be implemented in any portion of one or more of the interconnects illustrated or described herein.

[0088] A processor, in one embodiment, includes a microprocessor, multi-core processor, multithreaded processor, an ultra low voltage processor, an embedded processor, or other known processing element. In the illustrated implementation, processor acts as a main processing unit and central hub for communication with many of the various components of the system. As one example, processor is implemented as a system on a chip (SoC). As a specific illustrative example, processor includes an Intel® Architecture Core™-based processor such as an i3, i5, i7 or another such processor available from Intel Corporation, Santa Clara, CA. However, understand that other low power processors such as available from Advanced Micro Devices, Inc. (AMD) of Sunnyvale, CA, a MIPS-based design from MIPS Technologies, Inc. of Sunnyvale, CA, an ARM-based design licensed from ARM Holdings, Ltd. or customer thereof, or their licensees or adopters may instead be present in other embodiments such as an Apple A5/A6 processor, a Qualcomm Snapdragon processor, or TI OMAP processor. Note that many of the customer versions of such processors are modified and varied; however, they may support or recognize a specific instructions set that performs defined algorithms as set forth by the processor licensor. Here, the microarchitectural implementation may vary, but the architectural function of the processor is usually consistent. Certain details regarding the architecture and operation of processor in one implementation will be discussed further below to provide an illustrative example.

[0089] Processor, in one embodiment, communicates with a system memory. As an illustrative example, which in an embodiment can be implemented via multiple memory

devices to provide for a given amount of system memory. As examples, the memory can be in accordance with a Joint Electron Devices Engineering Council (JEDEC) low power double data rate (LPDDR)-based design such as the current LPDDR2 standard according to JEDEC JESD 209-2E (published April 2009), or a next generation LPDDR standard to be referred to as LPDDR3 or LPDDR4 that will offer extensions to LPDDR2 to increase bandwidth. In various implementations the individual memory devices may be of different package types such as single die package (SDP), dual die package (DDP) or quad die package (13P). These devices, in some embodiments, are directly soldered onto a motherboard to provide a lower profile solution, while in other embodiments the devices are configured as one or more memory modules that in turn couple to the motherboard by a given connector. And of course, other memory implementations are possible such as other types of memory modules, e.g., dual inline memory modules (DIMMs) of different varieties including but not limited to microDIMMs, MiniDIMMs. In a particular illustrative embodiment, memory is sized between 2GB and 16GB, and may be configured as a DDR3LM package or an LPDDR2 or LPDDR3 memory that is soldered onto a motherboard via a ball grid array (BGA).

[0090] To provide for persistent storage of information such as data, applications, one or more operating systems and so forth, a mass storage may also couple to processor. In various embodiments, to enable a thinner and lighter system design as well as to improve system responsiveness, this mass storage may be implemented via a SSD. However in other embodiments, the mass storage may primarily be implemented using a hard disk drive (HDD) with a smaller amount of SSD storage to act as a SSD cache to enable non-volatile storage of context state and other such information during power down events so that a fast power up can occur on re-initiation of system activities. A flash device may be coupled to processor, e.g., via a serial peripheral interface (SPI). This flash device may provide for non-volatile storage of system software, including a basic input/output software (BIOS) as well as other firmware of the system.

[0091] In various embodiments, mass storage of the system is implemented by a SSD alone or as a disk, optical or other drive with an SSD cache. In some embodiments, the mass storage is implemented as a SSD or as a HDD along with a restore (RST) cache module. In various implementations, the HDD provides for storage of between 320GB-4 terabytes (TB) and upward while the RST cache is implemented with a SSD having a capacity of 24GB-256GB. Note that such SSD cache may be configured as a single level cache (SLC) or multi-

level cache (MLC) option to provide an appropriate level of responsiveness. In a SSD-only option, the module may be accommodated in various locations such as in a mSATA or NGFF slot. As an example, an SSD has a capacity ranging from 120GB-1TB.

[0092] 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.

[0093] 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.

[0094] 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.

[0095] Use of the phrase ‘to’ or ‘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.

[0096] Furthermore, use of the phrases ‘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.

[0097] 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.

[0098] 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.

[0099] 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.

[00100] 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).

[00101] The following examples pertain to embodiments in accordance with this Specification. There is disclosed in an example 1, an apparatus comprising: a first pair of conductors to carry differential signals; at least one ground conductor adjacent to the first pair of conductors, the ground conductor to be connected to a ground plane; and at least one particular conductor to carry sideband signals, where the particular conductor is to be connected to a ground plane via a resonance mitigation circuit, and the resonance mitigation circuit comprises a resistor. Layout of the conductors may be according to a Peripheral Component Interconnect (PCI) Express (PCIe)-based specification.

[00102] In example 2, the apparatus of example 1 can optionally include the resonance mitigation circuit which further includes a capacitor in series with the resistor.

[00103] In example 3, the apparatus of any of examples 1 and 2, may optionally include the resonance mitigation circuit, which further includes a tee-connection, and the resistor and capacitor are in a first of at least two branches of the tee-connection.

[00104] In example 4, the resonance mitigation circuit of the apparatus of example 1 can optionally mitigate resonance manifesting in the particular conductor when the particular conductor has at least one floating connection.

[00105] In example 5, the floating connection of the apparatus of any of examples 1 and 4 can optionally be one of an open connection and a connection with a mismatched impedance

[00106] In example 6, the differential signals of the apparatus of example 1 can optionally include high speed signals with frequencies higher than frequencies of the sideband signals.

[00107] In example 7, the differential signals of the apparatus of any of examples 1 and 6 can optionally include signals with frequencies greater than 100 MHz.

[00108] In example 8, the resonance mitigation circuit of the apparatus of any of examples 1, 6, and 7 can optionally damp high frequency resonant energy while allowing operation of the sideband signals.

[00109] In example 9, the apparatus of example 1 can optionally include the first pair of conductors, ground conductor, and particular conductor.

[00110] In example 10, the connector of the apparatus of any of examples 1 and 9 can optionally include card edge contacts of an add-in card.

[00111] In example 11, the add-in card of any of examples 1, 9, and 10 can optionally include an add-in card compliant with a PCIe-based protocol.

[00112] In example 12, the connector of the apparatus of any of examples 1 and 9 can optionally include a card edge connector to accept an add-in card.

[00113] In example 13, the connector of the apparatus of any of examples 1 and 9 can optionally include the connector comprises card edge contacts of a baseboard card.

[00114] In example 14, the baseboard card of any of examples 1, 9, and 13 can optionally include a baseboard card compliant with a PCIe-based protocol.

[00115] There is disclosed in example 15, a method that includes sending a differential signal on a pair of conductors of a connector comprising a plurality of conductors, wherein the differential signal radiates energy to another one of the plurality of conductors, damping the energy manifesting on the other conductor from the differential signal using a resonance mitigation circuit, and sending a sideband signal on the other conductor. The other conductor can be coupled to a ground plane via the resonance mitigation circuit, and the resonance mitigation circuit can include a resistor and a capacitor in series. The layout of the plurality of conductors of the connector can be according to a Peripheral Component Interconnect (PCI) Express (PCIe)-based specification.

[00116] There is disclosed in example 16, a system including a board comprising a card edge connector and an add-in card connected to the card edge connector. The card edge connector can include a first pair of conductors to carry differential signals, at least one ground conductor adjacent to the first pair of conductors, the ground conductor to be connected to a ground plane, and at least one particular conductor to carry sideband signals, wherein the particular conductor is to be connected to a ground plane via a resonance mitigation circuit, and the resonance mitigation circuit comprises a resistor in series with a capacitor.

[00117] In example 17, the add-in card of the system of example 16 can optionally include a pair of conductors to correspond to the first pair of conductors and carry the differential signals, a ground conductor to correspond to the ground conductor of the card edge connector, and a sideband conductor to correspond to the particular conductor.

[00118] In example 18, the sideband conductor of the system of any one of examples 16 and 17 can optionally be connected to a ground plane via a respective resonance mitigation circuit, and the respective resonance mitigation circuit can include a resistor in series with a capacitor.

[00119] In example 19, the sideband conductor of the system of any one of examples 16, 17, and 18 can optionally be unterminated at the add-in card.

[00120] In example 20, the resonance mitigation circuit of the system of example 16 can optionally consist of the resistor connected in series to the capacitor.

[00121] 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.

[00122] 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.

CLAIMS:

1. An apparatus comprising:
a first pair of conductors to carry differential signals;
at least one ground conductor neighboring the first pair of conductors, the ground
conductor to be connected to a ground plane; and
5 at least one particular conductor to carry sideband signals, wherein the particular
conductor is to be connected to a ground plane via a resonance mitigation circuit, and the
resonance mitigation circuit comprises a resistor.
2. The apparatus of Claim 1, wherein the resonance mitigation circuit further comprises
10 a capacitor in series with the resistor.
3. The apparatus of any one of Claims 1-2, wherein the resonance mitigation circuit
further comprises a tee-connection, and the resistor and capacitor are in a first of at least two
branches of the tee-connection.
15
4. The apparatus of any one of Claims 1-3, wherein the resonance mitigation circuit is to
mitigate resonance manifesting in the particular conductor when the particular conductor has
at least one floating connection.
- 20 5. The apparatus of Claim 4, wherein the floating connection comprises one of an open
connection and a connection with a mismatched impedance.
6. The apparatus of any one of Claims 1-5, wherein the differential signals comprise
high speed signals with frequencies higher than frequencies of the sideband signals.
25
7. The apparatus of any one of Claims 1-6, wherein the differential signals comprise
signals with frequencies greater than 100 MHz.
8. The apparatus of any one of Claims 1-7, wherein the resonance mitigation circuit is
30 to damp high frequency resonant energy while allowing operation of the sideband signals.

9. The apparatus of any one of Claims 1-8, wherein the apparatus comprises a connector comprising the first pair of conductors, ground conductor, and particular conductor.

10. The apparatus of Claim 9, wherein the connector comprises card edge contacts of an add-in card.

11. The apparatus of Claim 10, wherein the add-in card comprises an add-in card compliant with a PCIe-based protocol.

12. The apparatus of Claim 9, wherein the connector comprises a card edge connector to accept an add-in card.

13. The apparatus of Claim 9, wherein the connector comprises card edge contacts of a baseboard card.

14. The apparatus of Claim 13, wherein the baseboard card comprises a baseboard card compliant with a PCIe-based protocol.

15. A method comprising:

sending a differential signal on a pair of conductors of a connector comprising a plurality of conductors, wherein the differential signal radiates energy to another one of the plurality of conductors;

damping the energy manifesting on the other conductor from the differential signal using a resonance mitigation circuit, wherein the other conductor is coupled to a ground plane via the resonance mitigation circuit, and the resonance mitigation circuit comprises a resistor and a capacitor in series; and

sending a sideband signal on the other conductor.

16. The method of Claim 15, wherein the resonance mitigation circuit further comprises a tee-connection, and the resistor and capacitor are in a first of at least two branches of the tee-connection.

17. The method of any one of Claims 15-16, wherein the other conductor has at least one floating connection.

18. The method of Claim 17, wherein the floating connection comprises one of an open connection and a connection with a mismatched impedance.

19. The method of any one of Claims 15-18, wherein the differential signals comprise high speed signals with frequencies higher than frequencies of the sideband signals.

20. The method of any one of Claims 15-19, wherein the differential signals comprise signals with frequencies greater than 100 MHz.

21. The method of any one of Claims 15-20, wherein the connector further comprises one or more ground conductors.

22. The method of any one of Claims 15-21, wherein the connector comprises card edge contacts of an add-in card.

23. The method of Claim 22, wherein the add-in card comprises an add-in card compliant with a PCIe-based protocol.

24. The method of any one of Claims 15-21, wherein the connector comprises a card edge connector to accept an add-in card.

25. The method of any one of Claims 15-21, wherein the connector comprises card edge contacts of a baseboard card.

26. The method of Claim 25, wherein the baseboard card comprises a baseboard card compliant with a PCIe-based protocol.

27. A system comprising means to perform the method of any one of Claims 15-26.

28. The system of Claim 27, wherein the means comprise code executable to perform steps of at least some of the methods of Claims 15-26.

29. A system comprising:

5 a board comprising a card edge connector;

an add-in card connected to the card edge connector,

wherein the card edge connector comprises:

a first pair of conductors to carry differential signals;

10 at least one ground conductor approximate to the first pair of conductors, the ground conductor to be connected to a ground plane; and

at least one particular conductor to carry sideband signals, wherein the particular conductor is to be connected to a ground plane via a resonance mitigation circuit, and the resonance mitigation circuit comprises a resistor in series with a capacitor.

15 30. The system of Claim 29, wherein the add-in card comprises:

a pair of conductors to correspond to the first pair of conductors and carry the differential signals;


a ground conductor to correspond to the ground conductor of the card edge connector; and

20 a sideband conductor to correspond to the particular conductor.

31. The system of any one of Claims 29-30, wherein the sideband conductor is to be connected to a ground plane via a respective resonance mitigation circuit, and the respective resonance mitigation circuit comprises a resistor in series with a capacitor.

25 32. The system of any one of Claims 29-30, wherein the sideband conductor is unterminated at the add-in card.

33. The system of any one of Claims 29-32, wherein the resonance mitigation circuit
30 consists of the resistor connected in series to the capacitor.

**FIG. 1**

Layered Protocol Stack 200




FIG. 2




FIG. 3




FIG. 4





FIG. 5

5/19

600



630 632 620 622 640 642 610 612 650 652

FIG. 6





FIG. 7A



Capacitor-Resistor combination absorbs high frequency coupled energy while blocking DC and low frequency signals




FIG. 7C




FIG. 7D




FIG. 7E




FIG. 7F





FIG. 7G






FIG. 9





FIG. 10



Baseboard connector vias

Resistor-only sideband termination

1100 ↗

FIG. 11




FIG. 12

1200




FIG. 13




FIG. 14

A. CLASSIFICATION OF SUBJECT MATTER**H04L 25/20(2006.01)I, H04L 25/03(2006.01)I, H04L 12/861(2013.01)I, H04L 29/10(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)

H04L 25/20; H01R 13/646; H01R 43/00; G06F 1/26; H01P 7/08; H01R 12/71; H02J 3/00; H04L 25/03; H04L 12/861; H04L 29/10

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: sideband, resonance, mitigation, ground

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2014-0357105 A1 (TIMOTHY D. WIG) 04 December 2014 See paragraphs [0003], [0035], [0037], [0045], [0050], [0051], [0071], [0081] , claims 1-10 and figure 4.	1-33
Y	US 2005-0013145 A1 (NORBERT HUBER) 20 January 2005 See paragraphs [0025], [0028], claims 1, 8, 9 and figure 1.	1-33
A	US 2014-0141654 A1 (TIMOTHY WIG et al.) 22 May 2014 See paragraphs [0055]-[0060], claims 1-9 and figures 4-10.	1-33
A	US 2014-0049343 A1 (JUN SAKAI) 20 February 2014 See abstract, paragraphs [0007]-[0013] and figures 13, 14, 16, 17.	1-33
A	US 2013-0132746 A1 (VIA TECHNOLOGIES, INC.) 23 May 2013 See abstract, claims 1-12 and figures 1, 2.	1-33



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 another citation or other special reason (as specified)

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

08 August 2016 (08.08.2016)

Date of mailing of the international search report

08 August 2016 (08.08.2016)

Name and mailing address of the ISA/KR

International Application Division

Korean Intellectual Property Office

189 Cheongsa-ro, Seo-gu, Daejeon, 35208, Republic of Korea

Facsimile No. +82-42-481-8578

Authorized officer

LEE, EUN KYU

Telephone No. +82-42-481-3580



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2016/033690

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2014-0357105 A1	04/12/2014	None	
US 2005-0013145 A1	20/01/2005	US 7019990 B2	28/03/2006
US 2014-0141654 A1	22/05/2014	None	
US 2014-0049343 A1	20/02/2014	JP 5931851 B2	08/06/2016
		US 8994470 B2	31/03/2015
		WO 2012-147803 A1	01/11/2012
US 2013-0132746 A1	23/05/2013	CN 102609071 A	25/07/2012
		CN 102609071 B	07/01/2015
		TW 201321957 A	01/06/2013
		US 9195297 B2	24/11/2015