



(19) **United States**
(12) **Patent Application Publication**
DIAB

(10) **Pub. No.: US 2012/0263176 A1**
(43) **Pub. Date: Oct. 18, 2012**

(54) **UNIFIED BUS ARCHITECTURE FOR POE COMMUNICATION AND CONTROL**

Publication Classification

(51) **Int. Cl.**
H04L 12/50 (2006.01)
(52) **U.S. Cl.** **370/359**
(57) **ABSTRACT**

(75) **Inventor: Wael William DIAB, San Francisco, CA (US)**

(73) **Assignee: Broadcom Corporation, Irvine, CA (US)**


(21) **Appl. No.: 13/538,424**


(22) **Filed: Jun. 29, 2012**

Embodiments of a unified communication and control bus architecture for Ethernet and/or PoE systems are provided. Embodiments enable a unified communication and control bus architecture that significantly simplifies communication and control in Ethernet and/or PoE systems. Embodiments enable significant savings both in terms of cost and complexity as the number of communication and control buses is reduced down to one. Embodiments can be used in various Ethernet and/or PoE implementations, including, for example, single PCB-single PoE, single PCB-multiple PoE, chassis-based switch, and stackable-based switch configurations. Further, embodiments can be implemented using standard Ethernet as well as proprietary implementations.


Related U.S. Application Data

(63) Continuation of application No. 12/179,476, filed on Jul. 24, 2008, now Pat. No. 8,245,056.





(Conventional)
FIG. 1



(Conventional)
FIG. 2




FIG. 3





FIG. 4



(Conventional)
FIG. 5




FIG. 6




FIG. 7




FIG. 8




FIG. 9




FIG. 10

UNIFIED BUS ARCHITECTURE FOR POE COMMUNICATION AND CONTROL

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] The present application is a continuation of U.S. patent application Ser. No. 12/179,476, filed Jul. 24, 2008, now allowed, which is incorporated herein by reference in its entirety.

BACKGROUND OF THE INVENTION

[0002] 1. Field of the Invention

[0003] The present invention relates generally to Power over Ethernet (PoE), and more particularly to a unified bus architecture for PoE communication and control.

[0004] 2. Background Art

[0005] Ethernet communications provide high speed communications between data terminals.

[0006] Power over Ethernet (PoE) systems enable power transmission over the same transmission lines that carry data in an Ethernet link. Generally, power is generated at a Power Source Equipment (PSE) side of the PoE system and is carried over an Ethernet cable to a Powered Device (PD) side of the PoE system.

[0007] As a result of enabling simultaneous power and data transmission, PoE systems have an isolated side and a non-isolated side, where the isolated side includes data subsystems and the non-isolated side includes power supply subsystems.

[0008] The isolated and non-isolated sides of a PoE system have different power requirements. Therefore, communication with the isolated side and communication with the non-isolated side are performed via separate communication buses. This implementation, however, is both expensive and complex, particularly when the number of PoE systems being controlled increases.

[0009] There is a need therefore for simplified and more efficient architectures for communication and control in PoE systems.

BRIEF SUMMARY OF THE INVENTION

[0010] Embodiments of a unified communication and control bus architecture for Ethernet and/or PoE systems are provided.

[0011] Embodiments enable a unified communication and control bus architecture that significantly simplifies communication and control in Ethernet and/or PoE systems. In particular, embodiments enable a communication and control architecture with a single bus for controlling all subsystems in an Ethernet and/or PoE system, notwithstanding differences between the subsystems. For example, embodiments allow for the use of a single bus to communicate with both PHY/Data subsystems and PoE subsystems in a PoE system, notwithstanding the different power requirements between the two types of subsystems.

[0012] Embodiments enable significant savings both in terms of cost and complexity as the number of communication and control buses is reduced down to one. For example, with a single communication and control bus according to embodiments, a single control structure (i.e., software and hardware control paths, addressing scheme, etc.) is needed to communicate between a system controller and the subsystems, and communication from the system controller's

perspective is reduced to the single task of addressing frames to their destination. Further, significant savings can be achieved in terms of the number of pins and ports at the system controller and/or at the subsystems that are needed for communication and control.

[0013] Embodiments can be used in various Ethernet and/or PoE implementations, including, for example, single PCB-single PoE, single PCB-multiple PoE, chassis-based switch, and stackable-based switch configurations. Further, embodiments can be implemented using standard Ethernet as well as proprietary implementations.

[0014] Further embodiments, features, and advantages of the present invention, as well as the structure and operation of the various embodiments of the present invention, are described in detail below with reference to the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS/FIGURES

[0015] The accompanying drawings, which are incorporated herein and form a part of the specification, illustrate the present invention and, together with the description, further serve to explain the principles of the invention and to enable a person skilled in the pertinent art to make and use the invention.

[0016] FIG. 1 is a block diagram of a conventional Power over Ethernet (PoE) system.

[0017] FIG. 2 illustrates a more detailed illustration of a conventional PoE system.

[0018] FIG. 3 illustrates an example PoE linecard.

[0019] FIG. 4 illustrates an example printed circuit board (PCB) having a PoE system.

[0020] FIG. 5 illustrates a conventional communication and control bus architecture in a chassis having a centralized controller and multiple PoE linecards.

[0021] FIG. 6 illustrates a unified communication and control bus architecture, according to an embodiment of the present invention.

[0022] FIG. 7 illustrates a unified communication and control bus architecture, according to an embodiment of the present invention.

[0023] FIG. 8 illustrates a unified communication and control bus architecture, according to an embodiment of the present invention.

[0024] FIG. 9 illustrates a unified communication and control bus architecture, according to an embodiment of the present invention.

[0025] FIG. 10 illustrates an example bus controller subsystem according to an embodiment of the present invention.

[0026] The present invention will be described with reference to the accompanying drawings. Generally, the drawing in which an element first appears is typically indicated by the leftmost digit(s) in the corresponding reference number.

DETAILED DESCRIPTION OF EMBODIMENT(S)

Introduction

[0027] FIG. 1 illustrates a high level diagram of a conventional Power over Ethernet (PoE) system **100** that provides both DC power and data communications over a common data communications medium. Referring to FIG. 1, power source equipment **102** provides DC power over conductors **104, 110** to a powered device (PD) **106** having a representa-

tive electrical load **108**. The PSE **102** and PD **106** also include data transceivers that operate according to a known communications standard, such as the IEEE Ethernet standard. More specifically, the PSE **102** includes a physical layer device on the PSE side that transmits and receives high speed data with a corresponding physical layer device in the PD **106**, as will be discussed further below. Accordingly, the power transfer between the PSE **102** and the PD **106** occurs simultaneously with the exchange of high speed data over the conductors **104**, **110**. In one example, the PSE **102** is a data switch having multiple ports that is communication with one or more PD devices, such as Internet phones, or a wireless access point.

[0028] The conductor pairs **104** and **110** can carry high speed differential data communications. In one example, the conductor pairs **104** and **110** each include one or more twisted wire pairs, or any other type of cable or communications media capable of carrying the data transmissions and DC power transmissions between the PSE and PD. In Ethernet communications, the conductor pairs **104** and **110** can include multiple twisted pairs, for example four twisted pairs for 10 Gigabit Ethernet. In 10/100 Ethernet, only two of the four pairs carry data communications, and the other two pairs of conductors are unused. Herein, conductor pairs may be referred to as Ethernet cables or communication links for ease of discussion.

[0029] FIG. 2 provides a more detailed circuit diagram of the PoE system **100**, where PSE **102** provides DC power to PD **106** over conductor pairs **104** and **110**. PSE **102** includes a transceiver physical layer device (or PHY) **202** having full duplex transmit and receive capability through differential transmit port **204** and differential receive port **206**. (Herein, transceivers may be referred to as PHYs) A first transformer **208** couples high speed data between the transmit port **204** and the first conductor pair **104**. Likewise, a second transformer **212** couples high speed data between the receive port **206** and the second conductor pair **110**. The respective transformers **208** and **212** pass the high speed data to and from the transceiver **202**, but isolate any low frequency or DC voltage from the transceiver ports, which may be sensitive large voltage values.

[0030] The first transformer **208** includes primary and secondary windings, where the secondary winding (on the conductor side) includes a center tap **210**. Likewise, the second transformer **212** includes primary and secondary windings, where the secondary winding (on the conductor side) includes a center tap **214**. The DC voltage supply **216** generates an output voltage that is applied across the respective center taps of the transformers **208** and **210** on the conductor side of the transformers. The center tap **210** is connected to a first output of a DC voltage supply **216**, and the center tap **214** is connected to a second output of the DC voltage supply **216**. As such, the transformers **208** and **212** isolate the DC voltage from the DC supply **216** from the sensitive data ports **204**, **206** of the transceiver **202**. An example DC output voltage is 48 volts, but other voltages could be used depending on the voltage/power requirements of the PD **106**.

[0031] The PSE **102** further includes a PSE controller **218** that controls the DC voltage supply **216** based on the dynamic needs of the PD **106**. More specifically, the PSE controller **218** measures the voltage, current, and temperature of the outgoing and incoming DC supply lines so as to characterize the power requirements of the PD **106**.

[0032] Further, the PSE controller **218** detects and validates a compatible PD, determines a power classification signature

for the validated PD, supplies power to the PD, monitors the power, and reduces or removes the power from the PD when the power is no longer requested or required. During detection, if the PSE finds the PD to be non-compatible, the PSE can prevent the application of power to that PD device, protecting the PD from possible damage. IEEE has imposed standards on the detection, power classification, and monitoring of a PD by a PSE in the IEEE 802.3 standard, which is incorporated herein by reference.

[0033] Still referring to FIG. 2, the contents and functionality of the PD **106** will now be discussed. The PD **106** includes a transceiver physical layer device **219** having full duplex transmit and receive capability through differential transmit port **236** and differential receive port **234**. A third transformer **220** couples high speed data between the first conductor pair **104** and the receive port **234**. Likewise, a fourth transformer **224** couples high speed data between the transmit port **236** and the second conductor pair **110**. The respective transformers **220** and **224** pass the high speed data to and from the transceiver **219**, but isolate any low frequency or DC voltage from the sensitive transceiver data ports.

[0034] The third transformer **220** includes primary and secondary windings, where the secondary winding (on the conductor side) includes a center tap **222**. Likewise, the fourth transformer **224** includes primary and secondary windings, where the secondary winding (on the conductor side) includes a center tap **226**. The center taps **222** and **226** supply the DC power carried over conductors **104** and **106** to the representative load **108** of the PD **106**, where the load **108** represents the dynamic power draw needed to operate PD **106**. A DC-DC converter **230** may be optionally inserted before the load **108** to step down the voltage as necessary to meet the voltage requirements of the PD **106**. Further, multiple DC-DC converters **230** may be arrayed in parallel to output multiple different voltages (3 volts, 5 volts, 12 volts) to supply different loads **108** of the PD **106**.

[0035] The PD **106** further includes a PD controller **228** that monitors the voltage and current on the PD side of the PoE configuration. The PD controller **228** further provides the necessary impedance signatures on the return conductor **110** during initialization, so that the PSE controller **218** will recognize the PD as a valid PoE device, and be able to classify its power requirements.

[0036] During ideal operation, a direct current (I_{DC}) **238** flows from the DC power supply **216** through the first center tap **210**, and divides into a first current (I_1) **240** and a second current (I_2) **242** that is carried over conductor pair **104**. The first current (I_1) **240** and the second current (I_2) **242** then recombine at the third center tap **222** to reform the direct current (I_{DC}) **238** so as to power PD **106**. On return, the direct current (I_{DC}) **238** flows from PD **106** through the fourth center tap **226**, and divides for transport over conductor pair **110**. The return DC current recombines at the second center tap **214**, and returns to the DC power supply **216**. As discussed above, data transmission between the PSE **102** and the PD **106** occurs simultaneously with the DC power supply described above. Accordingly, a first communication signal **244** and/or a second communication signal **246** are simultaneously differentially carried via the conductor pairs **104** and **110** between the PSE **102** and the PD **106**. It is important to note that the communication signals **244** and **246** are differential signals that ideally are not effected by the DC power transfer.

PoE Communication and Control

[0037] A detailed description of a PoE system, including a PSE and a PD, has been provided in FIGS. 1 and 2 above. In

particular, the coupling at the PSE side of the PoE system between data ports 204 and 206 of transceiver/PHY 202 and DC supply 216 has been described to identify the required isolation between two subsystems that can be found at the PSE side. Indeed, a PoE system at the PSE side generally includes an isolated side and a wire or non-isolated side, which typically operate at different power requirements. This is illustrated, for example, in FIG. 3, which shows an example PoE system 302. PoE system 302 can be used at the PSE side of a PoE arrangement. As shown in FIG. 3, PoE system 302 includes an isolated side 304 and a non-isolated side 306. Typically, the isolated side 304 includes a PHY/Data subsystem 310, which includes, for example, a transceiver such as transceiver 202 and/or Ethernet data subsystems. The non-isolated side 306 includes a PoE subsystem 312, which includes, for example, PoE power circuitry such as DC Supply 216 and PSE Controller 218.

[0038] Having different power requirements, the isolated side 304 and the non-isolated side 306 of PoE system 302 need to be electrically isolated from one another. For this reason, for example, the coupling between data ports 204 and 206 of transceiver 202 and conductor pairs 104 and 110 occurs through transformers 208 and 212. Further, any communication between the isolated side 304 and the non-isolated side 306 must be done through an isolation device 308. In FIG. 2, for example, communication between transceiver 202 and PSE Controller 218 occurs via an opto-isolator 248, which enables a serial communications interface between the two. Other types of isolation devices may also be used to enable this serial communications interface between transceiver 202 and PSE Controller 218. Further detail can be found in commonly owned U.S. patent application Ser. No. 12/168,577, titled "High Speed Isolation Interface for Use in Power Source Equipment (PSE) System," filed Jul. 7, 2008.

[0039] In addition to communication between the isolated side 304 and the non-isolated side 306 in a PoE system (which generally includes configuration and status polling by the isolated side), generally both sides need to communicate with a CPU/Controller, which provides them with command and control information.

[0040] The CPU/Controller may be located on the same printed circuit board (PCB) as the PoE system. For example, as shown in FIG. 4, PCB 402 includes PoE system 302 and a CPU/Controller 404. PCB 402 may represent a linecard with a single PCB, which is generally used as part of a chassis-based switch configuration, or a stackable unit with a single PCB, which is generally used as part of a stackable-based switch configuration (commonly known as pizza box). Note that a linecard or a stackable unit may include one or more PCBs such as PCB 402, for example.

[0041] Alternatively, the CPU/Controller 404 may be located externally to the PoE system 302. For example, in a chassis-based implementation, a centralized CPU/Controller is located in a supervisory slot of the chassis and controls one or more Ethernet and/or PoE linecards located on different slots of the chassis.

[0042] Regardless of the location of the CPU/Controller relative to the PoE system, at least two distinct communication buses are needed in conventional systems for the CPU/Controller to communicate with the isolated side and the non-isolated side of the PoE system. This communication and control architecture, again, is necessitated by the different power requirements between the isolated and non-isolated sides of the PoE system.

[0043] For example, in the single PCB-single PoE implementation 400 of FIG. 4, two separate communication buses 408 and 410 connect CPU/Controller 404 respectively to PHY/Data subsystem 310 and PoE subsystem 312 of PoE linecard 302. Communication bus 408 may be an out-of-band communication and management channel, separate from a data channel that is generally available between CPU/Controller 404 and PHY/Data subsystem 310. Similarly, communication bus 410 is a communication and management channel for PoE subsystem 312. Note that communication bus 410 may need to cross an isolation barrier to reach PoE subsystem 312. As described above, this can be done via an isolation device 406.

[0044] Similarly, in a multiple PoE implementation, at least two separate communication buses will be needed to couple the CPU/Controller to the PHY/Data subsystems and the PoE subsystems of the multiple PoE linecards/stackables. In an embodiment, two communication buses are used and implemented as shared buses, with one bus serving the PHY/Data subsystems and the other serving the PoE subsystems. However, the two communication buses cannot be combined due to the different power requirements between the PHY/Data and the PoE subsystems. In an alternative embodiment, two dedicated buses for each linecard/stackable unit are used, thereby resulting in 2xN buses in an N linecards/stackable-based system. In either embodiment, the CPU/Controller will need to run at least two different addressing schemes and enable at least two different software and hardware control paths to communicate with the PHY/Data and PoE subsystems as the CPU/Controller will handle two distinct buses.

[0045] The above described communication and control architectures can also be extended to a chassis-based or a stackable-based switch implementation, in which a centralized CPU/Controller located in a supervisory slot of the chassis/stackable-based switch controls one or more Ethernet and/or PoE linecards/stackable units located on different slots of the chassis/stackable-based switch. A chassis-based implementation 500 is illustrated in FIG. 5. As shown in FIG. 5, chassis-based switch 502 includes a supervisory slot 504a and a plurality of linecard slots 504b-n. Supervisory slot 504a contains a centralized CPU/Controller 512. In another embodiment, chassis-based switch 502 also includes a backup supervisory slot containing a backup CPU/Controller. Line card slots 504b-n each contains a respective Ethernet or PoE linecard. For example, in embodiment 500, linecard slots 504-n respectively contain PoE linecards 302b-n. As above, PoE linecard 302 includes a PHY/Data subsystem 310 and a PoE subsystem 312. Further, linecard 302 may include an isolation device 510, needed to cross the isolation barrier to PoE subsystem 312.

[0046] Chassis-based switch implementation 500 implements a dedicated bus communication and control architecture, in which a dedicated communication bus is used for each PHY/Data or PoE subsystem of each linecard in the chassis-based switch. Indeed, as shown in FIG. 5, communication buses 506b-n are each dedicated to serve a respective one of PHY/Data subsystems 310b-n of PoE linecards 302b-n. Similarly, communication buses 508b-n are each dedicated to serve a respective one of PoE subsystems 312b-n of PoE linecards 302b-n.

[0047] Accordingly, the number of buses needed in implementation 500 is at least two times the number of linecards in chassis 502. Clearly, this is both an expensive and complex implementation to maintain as the number of linecards

increases. While simplification in the number of buses can be achieved by using shared buses, one for PHY/Data subsystems and one for PoE subsystems, as described above, the complexity will remain. For example, referring to FIG. 5, CPU/Controller 512 would still need to access a distinct port to communicate over either shared bus and to address a particular subsystem. This also requires reserving at least two ports on supervisory slot 504a for communication and control. Further, as linecards are added and/or removed from chassis-based switch 502, re-configuration both at CPU/Controller 512 and at the added linecard will be needed, to ensure that CPU/Controller 512 can properly address the added linecard and that the added linecard is properly configured to communicate over the shared bus.

[0048] Embodiments of the present invention enable a unified communication and control bus architecture that significantly simplifies communication and control in Ethernet and/or PoE systems. In particular, embodiments enable a unified communication and control architecture with a single bus for controlling all subsystems in an Ethernet and/or PoE system, notwithstanding differences between the subsystems. For example, embodiments allow for the use of a single bus to communicate with both PHY/Data subsystems and PoE subsystems in a PoE system, notwithstanding the different power requirements between the two types of subsystems.

[0049] According to embodiments of the present invention, the unified communication and control architecture employs a single bus that implements an Ethernet or an Ethernet-like standard. Several implementations of the unified communication and control architecture according to embodiments of the present invention will be presented below. These implementations are provided for the purpose of illustration only and are not limiting of the scope of embodiments of the present invention. Further, although the implementations will be described with reference to chassis-based switch implementations, embodiments of the present invention are not limited as such and can be readily implemented in stackable-based switch implementations, for example. As will be further described below, generally, the implementations differ with respect to the number of bus tap points per linecard as well as the isolation implemented between the supervisory slot and the linecard slots. For example, FIGS. 6 and 7 illustrate implementations according to embodiments of the present invention with full isolation (according to the Ethernet standard) between the supervisory slot and the linecard slots. In contrast, FIGS. 8 and 9 illustrate implementations according to embodiments of the present invention with no isolation between the supervisory slot and the linecard slots.

[0050] FIG. 6 illustrates a unified communication and control bus architecture 600 according to an embodiment of the present invention. Embodiment 600 will be described with reference to a chassis-based switch implementation, but can be readily implemented in a stackable-based switch implementation (i.e., pizza box implementation).

[0051] As shown in FIG. 6, chassis-based switch 602 includes a supervisory slot 604a and a plurality of linecard slots 604b-n. As in implementation 500, supervisory slot 604a includes a centralized CPU/Controller 512, and linecard slots 604b-n each includes a respective Ethernet and/or PoE linecard 302. In contrast to implementation 500, however, architecture 600 uses a single bus 606 for communication and control between centralized CPU/Controller 512 and PHY/Data subsystems 310b-n and PoE subsystems 312b-n on linecards 302b-n.

[0052] According to embodiment 600, full isolation is implemented between centralized CPU/Controller 512 and linecards 302b-n. In an embodiment, isolation is implemented according to the Ethernet standard requirements, which impose a high degree of isolation between anything attached to the Ethernet cable (i.e., linecards 302b-n) and any circuitry sending and receiving transmissions over that cable (i.e., CPU/Controller 512). In an embodiment, isolation is achieved by a way of an isolation device 610, which may include an isolation transformer, for example.

[0053] With full isolation between CPU/Controller 512 and linecards 302b-n, CPU/Controller 512 will have a floating output as seen from linecards 302b-n. As such, logic re-referencing will be needed between the CPU/Controller side of the interface and the PHY/Data or PoE side of the interface. In an embodiment, this can be done using re-referencing circuitry 612 for re-referencing the output of CPU/Controller 512 from a first logic environment to a second logic environment. The second logic environment may be either that of PHY/Data subsystem 310 or that of PoE subsystem 312 (note that PHY/Data subsystem 310 and PoE subsystem 312 have different logic environments).

[0054] In embodiment 600, re-referencing circuitry 612 re-references the output of CPU/Controller 512 to the logic environment of PHY/Data subsystem 310. As such, no further isolation is needed beyond re-referencing circuitry 612 in the path to PHY/Data subsystem 310. However, isolation will be needed in the path to PoE subsystem 312 and is provided by isolation device 510. Alternatively, re-referencing circuitry 612 can be configured to re-reference the output of CPU/Controller 512 to the logic environment of PoE subsystem 312. As such, no further isolation would be needed beyond re-referencing circuitry 612 in the path to PoE subsystem 312, but isolation will be needed in the path to PHY/Data subsystem 310 (i.e., isolation device 510 would be moved to the PHY/Data subsystem path, instead of the PoE subsystem path).

[0055] Embodiment 600 represents a single bus tap point per linecard implementation, with a single bus tap point 608 used for each linecard slot 604. As such, a single bus controller subsystem 614 is needed per linecard slot. In an embodiment, bus controller subsystems 614b-n enable bus access arbitration between CPU/Controller 512, PHY/Data subsystems 310b-n, and PoE subsystems 312b-n. Further, each bus controller subsystem 614b-n acts to examine frames communicated over bus 606 and to forward to its respective subsystem(s) frames addressed thereto. In an embodiment, bus controller subsystem 614 includes multiple outputs in the direction of the linecard with embedded circuitry that switches a received frame onto the appropriate output according to whether the frame is intended for the PHY/Data or the PoE subsystem of the linecard. In an embodiment, bus controller subsystems 614b-n each includes an Ethernet bus controller subsystem.

[0056] An example bus controller subsystem according to an embodiment of the present invention is illustrated in FIG. 10. As shown in FIG. 10, bus controller subsystem 1000 includes a physical layer (PHY) module 1002, a medium access layer (MAC) module 1004, a host controller subsystem 1006, and a host interface module 1006. In an embodiment, host controller subsystem 1006 includes programmable logic, which may be implemented using field programmable gate array (FPGA) logic, for example. In an

embodiment, host interface module **1006** includes a PCI (Peripheral Component Interconnect) interface.

[0057] In an embodiment, bus **606** is a shared bus that unifies all the communication and control buses that are conventionally needed to connect supervisory slot **604a** to linecard slots **604b-n** into a single bus at the backplane of chassis-based switch **602**. In an embodiment, bus **606** implements a shared access protocol, such as CSMA/CD (Carrier Sense Multiple Access/Collision Detection), for example. Accordingly, bus **606** is implemented as a multi-drop interface as shown in embodiment **600**, with CPU/Controller **512**, PHY/Data subsystems **310b-n**, and PoE subsystems **312b-n** attached to bus **606** through respective tap points **608a**, **608b-n**, and **608b-n**. Alternatively, other shared access protocols may also be used, including, for example, token-based reservation protocols. Further, other shared bus topologies may also be used, including, for example, ring topologies.

[0058] As would be understood by a person skilled in the art based on the teachings herein, architecture **600** is not limited to the embodiment shown in FIG. 6. Indeed, architecture **600** may also be used in single PCB-single PoE, single PCB-multiple PoE, and stackable-based switch implementations.

[0059] In an embodiment, bus **606** enables a local area network (LAN) that implements IEEE 802.3 (the Ethernet standard) and operates according to the half-duplex mode of operation of the Ethernet standard. Thus, bus **606** also implements CSMA/CD as a shared access protocol. In an embodiment, bus **606** implements one of 1000BASE-KX, 10 GBASE-KX4, and 10 GBASE-KR. Further, bus **606** may use any of the shielded twisted pair cabling standards.

[0060] FIG. 7 illustrates another unified communication and control bus architecture **700** according to an embodiment of the present invention. Embodiment **700** will be described with reference to a chassis-based switch implementation, but can be readily implemented in a stackable-based switch implementation (i.e., pizza box implementation).

[0061] Embodiment **700** is similar to embodiment **600** described above in that it also uses full isolation between CPU/Controller **512** and linecards **706b-n** of chassis-based switch **702**. However, as shown in FIG. 7, embodiment **700** uses two bus tap points **710** and **712** per linecard slot **704**. As such, in an embodiment, two re-referencing circuits **714** and **716** are used for each linecard slot **704**, in order to re-reference the output of CPU/Controller **512** from the logic environment of CPU/Controller **512** to logic environments compatible with PHY/Data subsystem **310** and PoE subsystem **312**, respectively. Similarly, two bus controller subsystems **718** and **720** are used per linecard slot **704**, with each bus controller subsystem dedicated to either PHY/Data subsystem **310** or PoE subsystem **312**.

[0062] In an embodiment, bus **708** enables a local area network (LAN) that implements IEEE 802.3 (the Ethernet standard) and operates according to the half-duplex mode of operation of the Ethernet standard. In an embodiment, bus **708** implements one of 1000BASE-KX, 10 GBASE-KX4, and 10 GBASE-KR.

[0063] In an embodiment, bus controller subsystems **718b-n** and **720b-n** enable bus access arbitration between CPU/Controller **512**, PHY/Data subsystems **310b-n**, and PoE subsystems **312b-n**. Further, each bus controller subsystem **718b-n** or **720b-n** acts to examine frames communicated over bus **708** and to forward to its respective subsystem(s) frames

addressed thereto. In an embodiment, bus controller subsystems **718** and **720** include Ethernet bus controller subsystems.

[0064] FIG. 8 illustrates another unified communication and control bus architecture **800** according to an embodiment of the present invention. Embodiment **800** will be described with reference to a chassis-based switch implementation, but can be readily implemented in a stackable-based switch implementation (i.e., pizza box implementation).

[0065] Embodiment **800** has similar bus topology to embodiment **500** described above. However, embodiment **800** opts for no isolation between CPU/Controller **512** and linecards **302b-n** of chassis-based switch **802**. Note that this implementation may not be fully compliant with Ethernet or Ethernet-like standards, which impose isolation requirements. However, since CPU/Controller **512** and PHY/Data subsystems **310b-n** typically have common power levels, this implementation can be appropriate with isolation needed only in the paths between CPU/Controller **512** and PoE subsystems **312b-n**. In embodiment **800**, this is achieved using isolation devices **510b-n**, coupled respectively between bus controller subsystems **614b-n** and PoE subsystems **312b-n**.

[0066] Embodiment **800** is a single bus tap point per linecard implementation, with each linecard **302b-n** having a single respective bus tap point **808b-n**. Accordingly, a single Ethernet bus controller subsystem **614** is needed for each linecard **302b-n**.

[0067] In an embodiment, bus **806** enables a local area network (LAN) that implements IEEE 802.3 (the Ethernet standard) and operates according to the half-duplex mode of operation of the Ethernet standard. In an embodiment, bus **806** implements one of 1000BASE-KX, 10 GBASE-KX4, and 10 GBASE-KR.

[0068] In an embodiment, bus controller subsystems **614b-n** enable bus access arbitration between CPU/Controller **512**, PHY/Data subsystems **310b-n**, and PoE subsystems **312b-n**. Further, each bus controller subsystem **614b-n** acts to examine frames communicated over bus **806** and to forward to its respective subsystem(s) frames addressed thereto. In an embodiment, bus controller subsystems **614b-n** include Ethernet bus controller subsystems.

[0069] FIG. 9 illustrates another unified communication and control bus architecture **900** according to an embodiment of the present invention. Embodiment **900** will be described with reference to a chassis-based switch implementation, but can be readily implemented in a stackable-based switch implementation (i.e., pizza box implementation).

[0070] As shown in FIG. 9, embodiment **900** implements two bus tap points **908** and **910** off bus **906** for each linecard slot **904** of chassis-based switch **902**. Further, embodiment **900** implements no isolation between CPU/Controller **512** and linecards **706b-n**.

[0071] Since CPU/Controller **512** and Data/PHY subsystems **310b-n** typically have common power levels, there is no need for isolation in the paths that connect them. Indeed, as shown in FIG. 9, it is sufficient to couple a bus controller subsystem **718** between CPU/Controller **512** and PHY/Data subsystems **310b-n**. On the other hand, PoE subsystems **312b-n** have different logic levels than CPU/Controller **512**. As such, isolation devices **716b-n** need to be inserted in the path between CPU/Controller **512** and PoE subsystems **312b-n** as shown in FIG. 9. Separate bus controller subsystems **720b-n** will also be needed for PoE subsystems **312b-n** in this two bus tap points per linecard implementation.

[0072] In an embodiment, bus 906 enables a local area network (LAN) that implements IEEE 802.3 (the Ethernet standard) and operates according to the half-duplex mode of operation of the Ethernet standard. In an embodiment, bus 906 implements one of 1000BASE-KX, 10 GBASE-KX4, and 10 GBASE-KR.

[0073] In an embodiment, bus controller subsystems 718*b-n* and 720*b-n* enable bus access arbitration between CPU/Controller 512, PHY/Data subsystems 310*b-n*, and PoE subsystems 312*b-n*. Further, each bus controller subsystem 718*b-n* or 720*b-n* acts to examine frames communicated over bus 906 and to forward to its respective subsystem(s) frames addressed thereto. In an embodiment, bus controller subsystems 718 and 720 include Ethernet bus controller subsystems.

[0074] As would be understood by a person skilled in the art based on the teachings herein, choosing between the different implementations described above in FIGS. 6-9 for a given PoE system depends in large part on the size, cost, and performance requirements of the system.

[0075] Accordingly, embodiments of the present invention enable a unified communication and control bus architecture that significantly simplifies communication and control in Ethernet and/or PoE systems. In particular, embodiments enable a communication and control architecture with a single bus for controlling all subsystems in an Ethernet and/or PoE system, notwithstanding differences between the subsystems. For example, embodiments of the present invention allow for the use of a single bus to communicate with both PHY/Data subsystems and PoE subsystems in a PoE system, notwithstanding the different power requirements between the two types of subsystems.

[0076] As such, significant savings both in terms of cost and complexity can be achieved as the number of communication and control buses is reduced down to one. For example, with a single communication and control bus according to embodiments of the present invention, a single control structure (i.e., software and hardware control paths, addressing scheme, etc.) is needed to communicate between the system controller and the subsystems, and communication from the controller's perspective is reduced to the single task of addressing frames to their destination. Further, significant savings can be achieved in terms of the number of pins and ports at the system controller and/or at the subsystems that are needed for communication and control.

[0077] Embodiments of the present invention can be used in various Ethernet and/or PoE implementations, including, for example, single PCB-single PoE, single PCB-multiple PoE, chassis-based switch, and stackable-based switch configurations. Further, embodiments of the present invention can be implemented using standard Ethernet as well as proprietary implementations.

CONCLUSION

[0078] While various embodiments of the present invention have been described above, it should be understood that they have been presented by way of example only, and not limitation. It will be apparent to persons skilled in the relevant art that various changes in form and detail can be made therein without departing from the spirit and scope of the invention. Thus, the breadth and scope of the present invention should not be limited by any of the above-described exemplary embodiments, but should be defined only in accordance with the following claims and their equivalents.

What is claimed is:

1. A switch, comprising:
 - a supervisory slot, including a controller;
 - a first linecard slot, including a first Power Over Ethernet (PoE) linecard;
 - a second linecard slot, including a second PoE linecard; and
 - a shared communication and control bus configured to couple the supervisory slot to each of the first and second linecard slots.
2. The switch of claim 1, wherein at least one of the first and second PoE linecards includes a transceiver physical layer device (PHY) and a PoE subsystem.
3. The switch of claim 2, wherein the PoE subsystem includes at least one of a DC power supply and a Power Source Equipment (PSE) controller.
4. The switch of claim 1, wherein the supervisory slot further includes an isolation device, the isolation device configured to isolate the controller from at least one of the first and second PoE linecards.
5. The switch of claim 1, wherein at least one of the first and second linecard slots comprises:
 - re-referencing circuitry configured to re-reference an output of the controller from a first logic environment of the controller to a second logic environment.
6. The switch of claim 5, wherein the second logic environment corresponds to a logic environment of a transceiver physical layer device (PHY) of said at least one of the first and second linecard slots.
7. The switch of claim 6, wherein said at least one of the first and second linecard slots further comprises:
 - an isolation device configured to isolate said re-referencing circuitry from a PoE subsystem of said at least one of the first and second linecard slots.
8. The switch of claim 5, wherein the second logic environment corresponds to a logic environment of a PoE subsystem of said at least one of the first and second linecard slots.
9. The switch of claim 8, wherein said at least one of the first and second linecard slots further comprises:
 - an isolation device configured to isolate said re-referencing circuitry from a transceiver physical layer device (PHY) of said at least one of the first and second linecard slots.
10. The switch of claim 5, wherein said at least one of the first and second linecard slots further comprises:
 - a bus controller, coupled between the re-referencing circuitry and a PoE linecard of said at least one of the first and second linecard slots.
11. The switch of claim 1, wherein the shared communication and control bus implements a shared access protocol.
12. The switch of claim 1, wherein the shared communication and control bus implements Carrier Sense Multiple Access/Collision Detection (CSMA/CD).
13. The switch of claim 1, wherein the shared communication and control bus implements a token-based shared access protocol.
14. The switch of claim 1, wherein the shared communication and control bus is a multi-drop bus, and wherein each of the supervisory slot, the first linecard slot, and the second linecard slot is attached to the bus via a respective tap point.
15. A switch, comprising:
 - a supervisory slot, including a controller;
 - a linecard slot, including a Power Over Ethernet (PoE) linecard and a bus controller; and

a communication and control bus configured to couple the supervisory slot to the linecard slot via the bus controller.

16. The switch of claim **15**, wherein the PoE linecard comprises a transceiver physical layer device (PHY) and a PoE subsystem.

17. The switch of claim **16**, wherein the PoE linecard further comprises an isolation device coupled in a path between the PoE subsystem and the bus controller.

18. A switch, comprising:
a supervisory slot, including a controller;
a linecard slot, including a Power Over Ethernet (PoE) linecard; and
a communication and control bus configured to couple the supervisory slot to the linecard slot,
wherein the PoE linecard includes a transceiver physical layer device (PHY) and a PoE subsystem, and

wherein the linecard slot further comprises:

a first bus controller configured to couple the PHY of the PoE linecard to a first tap point of the communication and control bus; and

a second bus controller configured to couple the PoE subsystem of the PoE linecard to a second tap point of the communication and control bus.

19. The switch of claim **18**, wherein the linecard slot further comprises an isolation device coupled between the second bus controller and the second tap point of the communication and control bus.

20. The switch of claim **18**, wherein the PoE subsystem includes at least one of a DC power supply and a Power Source Equipment (PSE) controller.

* * * * *