



(51) International Patent Classification:

G06F 11/07 (2006.01) G06F 13/40 (2006.01)  
G06F 11/26 (2006.01) G06F 13/38 (2006.01)  
G06F 11/30 (2006.01)

(21) International Application Number:

PCT/US2018/029496

(22) International Filing Date:

26 April 2018 (26.04.2018)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

62/491,006 27 April 2017 (27.04.2017) US

(71) Applicant: LIQID INC. [US/US]; 329 Interlocken Parkway, Suite 200, Broomfield, Colorado 80021 (US).

(72) Inventors: LONG, Christopher R.; c/o Lqid Inc., 329 Interlocken Parkway, Suite 200, Broomfield, Colorado 80021 (US). HEYD, Andrew Rudolph; c/o Lqid Inc., 329 Interlocken Parkway, Suite 200, Broomfield, Colorado 80021 (US). CANNATA, James Scott; c/o Lqid Inc., 329 Interlocken Parkway, Suite 200, Broomfield, Colorado 80021 (US). PURI, Sumit; c/o Lqid Inc., 329 Interlocken Parkway, Suite 200, Broomfield, Colorado 80021 (US). SCHRAMM, Bryan; c/o Lqid Inc., 329 Interlocken Parkway, Suite 200, Broomfield, Colorado 80021 (US).

(74) Agent: GEHMAN, Leslie P. et al.; Setter Roche LLP, 14694 Orchard Parkway, Building A, Suite 200, Westminster, Colorado 80023 (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, DJ, DK, DM, DO,

(54) Title: PCIe FABRIC CONNECTIVITY EXPANSION CARD

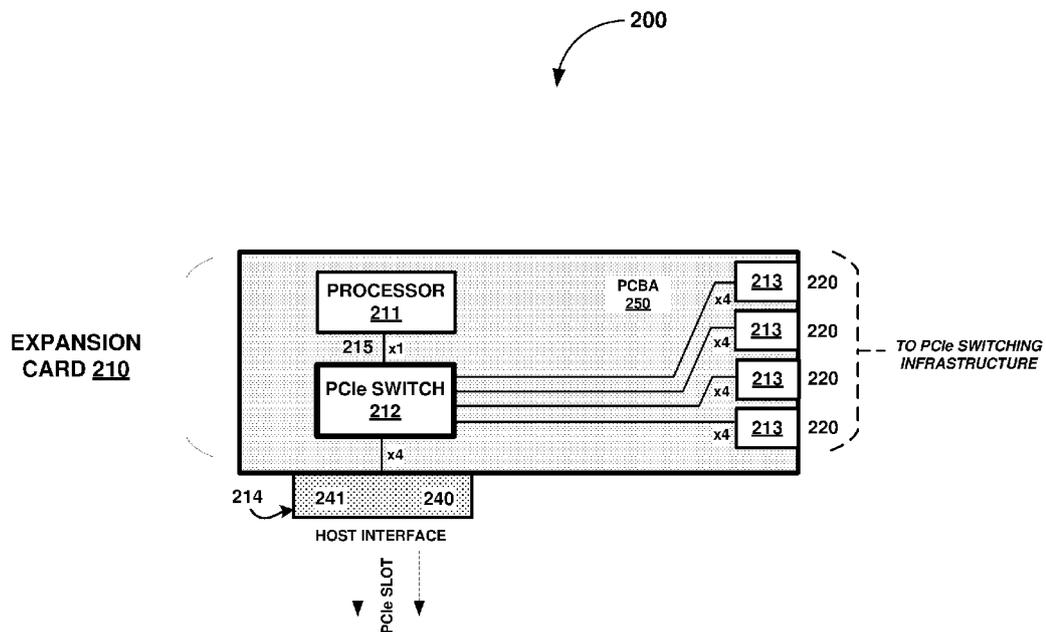


FIGURE 2

(57) Abstract: A connectivity card insertable into a connector of a host system is provided. The connectivity card includes a plurality of Peripheral Component Interconnect Express (PCIe) connectors configured to provide external PCIe ports on the connectivity card, each of the plurality of PCIe connectors capable of carrying PCIe traffic. The connectivity card also includes a PCIe switch circuit configured to communicatively couple the plurality of connectors to a shared connectivity interface carried over a host connector of the connectivity card. The connectivity card further includes control circuitry configured to monitor for connectivity issues that arise with regard to the plurality of PCIe connectors, and responsively mitigate the connectivity issues by at least reconfiguring a communication pathway in the PCIe switch circuit for at least a portion of the PCIe traffic affected by the connectivity issues.

DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, 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))*

**Published:**

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

## PCIe FABRIC CONNECTIVITY EXPANSION CARD

### RELATED APPLICATIONS

[0001] This application hereby claims the benefit of and priority to U.S. Provisional  
5 Patent Application Number 62/491,006, entitled “PCIe FABRIC CONNECTIVITY  
EXPANSION CARD,” filed on April 27, 2017 and which is hereby incorporated by  
reference in its entirety.

### BACKGROUND

[0002] Networked storage and computing systems have been introduced which store and  
10 process large amounts of data in enterprise-class storage environments. These networked  
storage systems typically provide access to bulk data storage over one or more network  
interfaces to end users or other external systems. In addition to storage of data, remote  
computing systems include various processing systems that can provide remote computing  
resources to end users. These networked storage systems and remote computing systems can  
15 be included in high-density installations, such as rack-mounted environments. Some  
computing devices employ Peripheral Component Interconnect Express (PCIe) interfaces to  
connect to peripherals and storage devices. However, typical PCIe implementations employ  
point-to-point host-device architectures.

### 20 OVERVIEW

[0003] In an embodiment, a connectivity card insertable into a connector of a host system  
is provided. The connectivity card includes a plurality of Peripheral Component Interconnect  
Express (PCIe) connectors configured to provide external PCIe ports on the connectivity  
card, each of the plurality of PCIe connectors capable of carrying PCIe traffic.

25 [0004] The connectivity card also includes a PCIe switch circuit configured to  
communicatively couple the plurality of connectors to a shared connectivity interface carried  
over a host connector of the connectivity card.

[0005] The connectivity card further includes control circuitry configured to monitor for  
connectivity issues that arise with regard to the plurality of PCIe connectors, and responsively  
30 mitigate the connectivity issues by at least reconfiguring a communication pathway in the  
PCIe switch circuit for at least a portion of the PCIe traffic affected by the connectivity  
issues.

[0006] In another embodiment, a method of operating a connectivity card insertable into a connector of a host system is provided. The method includes monitoring a plurality of external Peripheral Component Interconnect Express (PCIe) ports for connectivity issues, each of the plurality of PCIe ports capable of carrying PCIe traffic, and each PCIe port  
5 having a corresponding PCIe connector on the connectivity card, and detecting connectivity issues in one of the plurality of PCIe ports.

[0007] The method also includes mitigating the connectivity issues by reconfiguring a communication pathway in a PCIe switch circuit on the connectivity card for at least a portion of the PCIe traffic affected by the connectivity issues, the PCIe switch circuit  
10 configured to communicatively couple the plurality of PCIe ports to a shared connectivity interface carried over a host connector of the connectivity card.

[0008] In a further embodiment, one or more non-transitory computer-readable media having stored thereon program instructions to facilitate operating a connectivity card insertable into a connector of a host system is provided. The program instructions, when  
15 executed by a computing system, direct the computing system to at least monitor a plurality of external Peripheral Component Interconnect Express (PCIe) ports for connectivity issues, each of the plurality of PCIe ports capable of carrying PCIe traffic, and each PCIe port having a corresponding PCIe connector on the connectivity card, and to detect connectivity issues in one of the plurality of PCIe ports.

[0009] The program instructions further direct the computing system to mitigate the connectivity issues by reconfiguring a communication pathway in a PCIe switch circuit on the connectivity card for at least a portion of the PCIe traffic affected by the connectivity issues, the PCIe switch circuit configured to communicatively couple the plurality of PCIe  
20 ports to a shared connectivity interface carried over a host connector of the connectivity card.

## **BRIEF DESCRIPTION OF THE DRAWINGS**

[0010] Many aspects of the disclosure can be better understood with reference to the following drawings. The components in the drawings are not necessarily to scale, emphasis instead being placed upon clearly illustrating the principles of the present disclosure.

Moreover, in the drawings, like reference numerals designate corresponding parts throughout the several views. While several embodiments are described in connection with these  
30 drawings, the disclosure is not limited to the embodiments disclosed herein. On the contrary, the intent is to cover all alternatives, modifications, and equivalents.

[0011] Figure 1 is a system diagram illustrating a computing environment.

[0012] Figure 2 is a block diagram illustrating a connectivity expansion card.

[0013] Figure 3 is a system diagram illustrating a connectivity expansion card environment.

[0014] Figure 4 is a block diagram illustrating a processing system.

5 [0015] Figure 5 is a flow diagram illustrating a method of operating a connectivity expansion card.

[0016] Figure 6A is an oblique view illustrating a connectivity expansion card.

[0017] Figure 6B is a front view illustrating a connectivity expansion card.

[0018] Figure 7A is an oblique view illustrating a connectivity expansion card.

10 [0019] Figure 7B is a front view illustrating a connectivity expansion card.

### DETAILED DESCRIPTION

[0020] Figure 1 is a system diagram illustrating a computing environment 100. In this example, computing environment 100 includes server 110, cluster A 120, and cluster B 130.

15 Server 110 is coupled to cluster A 120 via PCIe links 122 and 126. Server 110 is coupled to cluster B 130 via PCIe links 132 and 136.

[0021] PCIe link 122 is connected to cluster A 120 through PCIe connector 124, and to server 110 through PCIe connector 117. PCIe link 126 is connected to cluster A 120 through PCIe connector 128, and to server 110 through PCIe connector 117. PCIe link 132 is

20 connected to cluster B 130 through PCIe connector 134, and to server 110 through PCIe connector 117. PCIe link 136 is connected to cluster B 130 through PCIe connector 138, and to server 110 through PCIe connector 117.

[0022] Server 110 includes motherboard 112, PCIe slot 114, and connectivity card 116. Connectivity card 116 includes four PCIe connectors 117. In a prior art example, if there are connectivity issues with PCIe link 122 (perhaps it has been disconnected from PCIe connector 124), the connectivity issues may cause a cascading failure problem for server 110, referred to as a host panic condition.

[0023] However, in an example embodiment of the present invention, connectivity card 116, which has been monitoring its four PCIe connectors 117 for connectivity issues, automatically switches traffic from PCIe link 122 to PCIe link 126 upon detecting the connectivity issues. Connectivity card 116 performs this redirection of traffic automatically and in such a way that server 110 is not disrupted by the connectivity issues.

30 [0024] Figures 2-3 are blocks diagram illustrating connectivity expansion card 210. Figure 2 illustrates an example physical configuration of connectivity expansion card 210 as

shown for printed circuit board assembly 250. Figure 3 illustrates an example schematic configuration of connectivity expansion card 210. Connectivity expansion card 210 includes one or more external Peripheral Component Interconnect Express (PCIe) compatible ports 213, PCIe switch circuit 212, processor 211, and host connector 214. Processor 211 and  
5 PCIe switch circuit 212 communicate over one or more interfaces 215, which can comprise a x1 PCIe interface. PCIe switch circuit 212 and host connector 214 communicate over one or more interfaces comprising one or more PCIe lanes, such as the x4 PCIe interface shown in Figure 2. Further communication links can be included for communication between the various elements of connectivity expansion card 210.

10 **[0025]** Connectivity expansion card 210 can comprise a HHHL (half-height half-length) PCIe peripheral or expansion card. Other card sizes can be employed, such as a FHHL (full-height half-length) or FHFL (full-height full-length), or HHFL (half-height full length), among others. PCIe edge connector 214 is insertable into a mating PCIe socket of a host system, such as a motherboard or daughterboard of a computer or server system.

15 Connectivity expansion card 210 can receive power over one or more power links provided by the host system over connector 214.

**[0026]** PCIe switch circuit 212 comprises one or more PCIe switch elements and distributes PCIe communications and traffic to one or more of ports 213. PCIe switch circuit 212 is also communicatively coupled to an on-card processor 211 or control system that  
20 handles traffic monitoring, traffic statistics communication, power monitoring, and status monitoring, among other operations.

**[0027]** PCIe-compatible connectors 213 can be used for cluster interconnect and can comprise mini-Serial Attached SCSI (SAS) HD connectors which are employed to carry PCIe signaling over mini-SAS cabling. Other example connectors include zQSFP+, microQSFP,  
25 and OCuLink-2 interconnect. In further examples, MiniSAS HD cables are employed that drive 12Gb/s versus 6Gb/s of standard SAS cables, since 12Gb/s can support PCIe Gen 3. PCIe-compatible connectors 213 can comprise mini-SAS connectors that comprise mini-SAS jacks. Associated cabling for PCIe links 220 can comprise SAS cabling which can include associated shielding, wiring, sheathing, and termination connectors.

30 **[0028]** PCIe-compatible connectors 213 are grouped into a set of four (4) ports in the examples herein to allow connections to other switches, devices, or servers with up to a x16 bus, as well as supporting x1, x4, and x8 bus widths, among others. Each PCIe-compatible connector 213 can also have included indicators (such as light-emitting diodes (LEDs)) that

shine through associated PCIe port connectors to help guide customers in connecting cables, troubleshooting, and indicating connectivity and errors.

**[0029]** PCIe switch circuit 212 communicates with a host system or host module (not pictured) over PCIe link 240. PCIe link 240 can comprise a PCIe link with multiple lanes, 5 such as a “x4” PCIe link, although a different number of PCIe lanes can be employed.

Additionally, more than one PCIe link 240 can be employed for load balancing, redundancy, and failover protection for connectivity expansion card 210. PCIe switch circuit 212 also communicates with four external PCIe-compatible connectors 213 over associated x4 PCIe links, although a different number of links and connectors can be employed. PCIe can 10 support multiple bus widths, such as x1, x4, x8, and x16, among others, with each multiple of bus width comprising an additional “lane” for data transfer. An example embodiment of the present invention also supports transfer of sideband signaling, such as System Management Bus (SMBus) interfaces, as well as associated clocks, power, and bootstrapping, among other signaling.

**[0030]** PCIe switch circuit 212 comprises one or more PCIe crosspoint switches, which logically interconnect various ones of the associated PCIe links based at least on the traffic carried by each PCIe link. PCIe switch circuit 212 establishes switched connections between any PCIe interfaces handled by PCIe switch circuit 212. Each PCIe switch port can comprise a non-transparent (NT) or transparent port or Downstream Port Containment (DCP/eDPC). 15

An NT port can allow some logical isolation between endpoints, while a transparent port does not allow logical isolation, and has the effect of connecting endpoints in a purely switched configuration. Access over an NT port or ports can include additional handshaking between the PCIe switch and the initiating endpoint to select a particular NT port or to allow visibility through the NT port. In other examples, a domain-based PCIe signaling distribution can be 20 included which allows segregation of PCIe ports of a PCIe switch according to user-defined groups. In some examples, PCIe switch circuit 212 comprises a PLX Technology PEX8725 10-port, 24 lane PCIe switch chip. In some examples, PCIe switch circuit 212 comprises a PLX Technology PEX8796 24-port, 96 lane PCIe switch chip.

**[0031]** Although host PCIe link 240 is used in Figures 2-3, it should be understood that 30 additional or different communication links or busses can be employed, such as Ethernet, Serial Attached SCSI (SAS), FibreChannel, Thunderbolt, Serial Attached ATA Express (SATA Express), among other interconnect, network, and link interfaces. Any of the links in Figures 2 and 3 can each use various communication media, such as air, space, metal, optical fiber, or some other signal propagation path, including combinations thereof. Any of the

PCIe links in Figures 2-3 can include any number of PCIe links or lane configurations. Any of the links in Figures 2-3 can each be a direct link or might include various equipment, intermediate components, systems, and networks. Any of the links in Figures 2-3 can each be a common link, shared link, aggregated link, or may be comprised of discrete, separate links.

5 **[0032]** Processor 211 can optionally communicate over at least PCIe link 215 and optional sideband links. Sideband links can include Universal Serial Bus (USB), SMBus, Inter-Integrated Circuit (I2C), controller area network bus (CAN), or any other communication interface, and in some examples are provided over portions of PCIe link 215. In Figure 3, processor 211 includes I2C interface 311 and USB interface 312 for communication over  
10 sideband links 241. I2C interface 311 and USB interface 312 can be included in separate circuitry or included in similar elements as processor 211.

**[0033]** Processor 211 comprises one or more microprocessors, processing devices, multi-core processors, processing circuitry, or other processing system. Processor 211 can include one or more non-transitory memory devices, such as RAM, solid state storage, or other  
15 memory to store instructions that are executable by processor 211 to operate as discussed herein. In some examples, processor 211 comprises an ARM microcontroller, ARM microprocessor, field-programmable gate array (FPGA), application specific integrated circuit (ASIC), application specific processor, or other microprocessor or processing elements. In some examples, processor 211 comprises an ARM-compatible microprocessor  
20 or microcontroller, although other circuitry can be employed. Processor 211 can comprise any processing elements discussed below for control processor 400 of Figure 4. Processor 211 can monitor connectivity status, usage statistics, traffic status, or other usage information through link 215 or other sideband links. PCIe switch circuit 212 can monitor and store this usage information during normal operation and data transfer, and processor 211 can retrieve  
25 this usage information as needed over link 215.

**[0034]** Although processor 211 and PCIe switch circuit 212 are shown as separate elements in Figures 2-3, it should be understood that processor 211 and PCIe switch circuit 212 can be included in the same circuitry.

**[0035]** Holdup circuitry can optionally be included in some examples on connectivity  
30 expansion card 210 to provide power to the connectivity card when input power has been lost or removed for the connectivity card. In some examples, the connectivity card is removed from an associated mating connector and input power is lost due to the removal. In other examples, power is lost to a host system into which connectivity expansion card 210 is connected, such as during a facility power outage or when an associated power supply fails.



The various holdup circuitry can also be accompanied by a power controller circuit to selectively provide power to the elements of connectivity expansion card 210.

**[0036]** Figure 3 also includes system elements for example context of connectivity expansion card 210. In Figure 3, connectivity expansion card 210 is coupled to host system 5 340, such as inserted into a PCIe slot of host system 340. Connectivity expansion card 210 provides PCIe ports that are external to a chassis, case, or enclosure of host system 340. These external ports carry PCIe links 220 and can be coupled over associated cabling to one or more rackmount switch elements, such as rackmount PCIe switches 320-321 in Figure 3. In this manner, host system 340 can be coupled into a cluster or rackmount PCIe fabric using 10 connectivity expansion card 210. Other external PCIe devices can also be coupled over links 220.

**[0037]** Figure 4 is a block diagram illustrating control processor 400. Control processor 400 illustrates an example of any of the control modules or processors discussed herein, such as processor 211 of Figures 2 and 3. In addition, control processor 400 can be illustrative of 15 any processing system a connectivity card discussed herein.

**[0038]** Control processor 400 includes communication interface 401 and processing system 410. Processing system 410 includes processing circuitry 411, random access memory (RAM) 412, and storage 413, although further elements can be included. Example contents of RAM 412 are further detailed in RAM space 420, and example contents of 20 storage 413 are further detailed in storage system 460.

**[0039]** Processing circuitry 411 can be implemented within a single processing device but can also be distributed across multiple processing devices or sub-systems that cooperate in executing program instructions. Examples of processing circuitry 411 include general purpose central processing units, microprocessors, application specific processors, and logic 25 devices, as well as any other type of processing device. In some examples, processing circuitry 411 includes physically distributed processing devices, such as cloud computing systems.

**[0040]** Communication interface 401 includes one or more communication and network interfaces for communicating over communication links, networks, such as packet networks, 30 the Internet, and the like. The communication interfaces can include PCIe interfaces, serial links, such as SPI links, I2C links, USB links, UART links, or one or more local or wide area network communication interfaces which can communicate over Ethernet or Internet protocol (IP) links. Communication interface 401 can include network interfaces configured to communicate using one or more network addresses, which can be associated with different

network links. Examples of communication interface 401 include network interface card equipment, transceivers, modems, and other communication circuitry.

**[0041]** RAM 412 and storage 413 together can comprise a non-transitory data storage system, although variations are possible. RAM 412 and storage 413 can each comprise any storage media readable by processing circuitry 411 and capable of storing software. RAM 412 can include volatile and nonvolatile, removable and non-removable media implemented in any method or technology for storage of information, such as computer readable instructions, data structures, program modules, or other data. Storage 413 can include non-volatile storage media, such as solid-state storage media, flash memory, phase change memory, or magnetic memory, including combinations thereof. RAM 412 and storage 413 can each be implemented as a single storage device but can also be implemented across multiple storage devices or sub-systems. RAM 412 and storage 413 can each comprise additional elements, such as controllers, capable of communicating with processing circuitry 311.

**[0042]** Software stored on or in RAM 412 or storage 413 can comprise computer program instructions, firmware, or some other form of machine-readable processing instructions having processes that when executed a processing system direct control processor 400 to operate as described herein. For example, software can drive processor 400 to monitor operating statistics and status for various PCIe traffic and other modules, monitor for connectivity issues that arise with regard to the plurality of PCIe connectors, and responsively mitigate the connectivity issues by at least reconfiguring a communication pathway in the PCIe switch circuit for at least a portion of the PCIe traffic affected by the connectivity issues, among other operations. The software can also include user software applications, application programming interfaces (APIs), or user interfaces. The software can be implemented as a single application or as multiple applications. In general, the software can, when loaded into a processing system and executed, transform the processing system from a general-purpose device into a special-purpose device customized as described herein.

**[0043]** RAM space 420 illustrates a detailed view of an example configuration of RAM 412. It should be understood that different configurations are possible. In an example embodiment of the present invention, applications 421 and custom kernel 422 can reside in flash memory and run in an execute in place (XIP) configuration. This enables quick boots while providing sufficient memory to execute telemetry and PCIe monitoring functions. Software applications 423-424 each comprise executable instructions which can be executed by processor 400 for operating a connectivity card or other circuitry according to the

operations discussed herein. Specifically, PCIe monitor 423 monitors for connectivity issues that arise with regard to the plurality of PCIe connectors. PCIe redirect control 424 mitigates connectivity issues by at least reconfiguring a communication pathway in the PCIe switch circuit for at least a portion of the PCIe traffic affected by the connectivity issues.

5 [0044] Applications 421 and custom kernel 422 can reside in RAM space 420 during execution and operation of control processor 400, and can reside in storage system 460 during a powered-off state, among other locations and states. Applications 421 and custom kernel 422 can be loaded into RAM space 420 during a startup or boot procedure as described for computer operating systems and applications.

10 [0045] Storage system 460 illustrates a detailed view of an example configuration of storage 413. Storage system 460 can comprise flash memory such as NAND flash or NOR flash memory, phase change memory, magnetic memory, among other solid-state storage technologies. As shown in Figure 4, storage system 460 includes system software 461.

[0046] Control processor 400 is generally intended to represent a computing system with  
15 which at least software 461 and 421-424 are deployed and executed in order to render or otherwise implement the operations described herein. However, control processor 400 can also represent any computing system on which at least software 461 and 421-424 can be staged and from where software 461 and 421-424 can be distributed, transported, downloaded, or otherwise provided to yet another computing system for deployment and  
20 execution, or yet additional distribution.

[0047] Figure 5 is a flow diagram illustrating a method of operating a connectivity card. Operations of Figure 5 can be performed by any of processor 211, PCIe switch 212, or control processor 400, including combinations thereof. Processor 211 will be discussed below as an exemplary processing system.

25 [0048] In Figure 5, control processor 400 monitors ports 213 and connector 214 for connectivity issues, (operation 500). Control processor 400 then detects connectivity issues among any of ports 213 and connector 214, (operation 502). These connectivity issues can include cable connectivity, such as when an associated cable connected to ports 213 is disconnected or re-connected. The connectivity issues can include failed internal links, such  
30 as when traces or portions of the links or ports internal to card 210 fail or are intermittent. Normally, a connectivity issue - such as connectivity loss - would cause a cascading failure problem for the host, referred to as a host panic condition. Control processor 400 can prevent host panic conditions by isolating the connectivity issue to within PCIe switch circuit 212, and control processor 400 instructs PCIe switch circuit 212 to redirect the affected PCIe

traffic to a redundant link, by reconfiguring a communication pathway within PCIe switch circuit 212, (operation 504).

5 [0049] For example, in Figure 3, expansion card 210 is coupled to two separate rackmount PCIe switches 320-321. If a connectivity issue arises with a cable or other portion associated with a first PCIe link between card 210 and switch 320, then processor 211 can detect this connectivity issue and redirect traffic that was formerly carried over the failed link. This redirection can occur within PCIe switch circuit 212 to redirect the traffic to another link to switch 320, or to another switch 321. The nature of the failure might preclude failure redirection to another link on the same switch 320. In this manner, connectivity failures -  
10 such as failed cables, links, or inadvertent cable disconnection, can be isolated from a host system by processor 211 and PCIe switch circuit 212. Processor 211 or control processor 400 can responsively mitigate PCIe connectivity issues by at least reconfiguring a communication pathway in PCIe switch circuit 212 for at least a portion of the PCIe traffic affected by the connectivity issues.

15 [0050] Figure 6A is an oblique view and Figure 6B is a front view illustrating a connectivity expansion card 600. In this example embodiment of the present invention, connectivity expansion card 600 comprises control processor (or control circuitry) 610, PCIe switch circuit 620, PCIe connectors 630, and host connector 640.

20 [0051] Figure 7A is an oblique view and Figure 7B is a front view illustrating a connectivity expansion card 700. In this example embodiment of the present invention, connectivity expansion card 700 comprises control processor (or control circuitry) integrated with PCIe switch circuit 710, PCIe connectors 720, and host connector 730.

[0052] The included descriptions and figures depict specific embodiments to teach those skilled in the art how to make and use the best mode. For the purpose of teaching inventive principles, some conventional aspects have been simplified or omitted. Those skilled in the art will appreciate variations from these embodiments that fall within the scope of the invention. Those skilled in the art will also appreciate that the features described above can be combined in various ways to form multiple embodiments. As a result, the invention is not limited to the specific embodiments described above, but only by the claims and their  
25  
30 equivalents.

**CLAIMS:**

What is claimed is:

1. A connectivity card, comprising:
  - a plurality of Peripheral Component Interconnect Express (PCIe) connectors
  - 5 configured to provide external PCIe ports on the connectivity card, each of the plurality of PCIe connectors capable of carrying PCIe traffic;
  - a PCIe switch circuit configured to communicatively couple the plurality of connectors to a shared connectivity interface carried over a host connector of the connectivity card; and
  - 10 control circuitry configured to monitor for connectivity issues that arise with regard to the plurality of PCIe connectors, and responsively mitigate the connectivity issues by at least reconfiguring a communication pathway in the PCIe switch circuit for at least a portion of the PCIe traffic affected by the connectivity issues.
- 15 2. The connectivity card of claim 1, wherein the control circuitry is further configured to mitigate the connectivity issues by determining an identity of an external PCIe port affected by the connectivity issues, and redirecting traffic from the affected external PCIe port to a different external PCIe port.
- 20 3. The connectivity card of claim 1, wherein the control circuitry is further configured to determine which external PCIe ports are affected by the connectivity issues and which external PCIe ports are unaffected by the connectivity issues, and redirecting PCIe traffic from the affected external PCIe ports to the unaffected external PCIe ports.
- 25 4. The connectivity card of claim 1, wherein the plurality of PCIe connectors comprises at least four PCIe connectors.
5. The connectivity card of claim 1, wherein the connectivity interface and the host connector comprise a PCIe link with multiple lanes.
- 30 6. The connectivity card of claim 5, wherein the PCIe link has four lanes.
7. The connectivity card of claim 1, wherein the connectivity interface and the host connector comprise at least two PCIe links.

8. A method of operating a connectivity card, the method comprising:  
monitoring a plurality of external Peripheral Component Interconnect Express (PCIe) ports for connectivity issues, each of the plurality of PCIe ports capable of carrying  
5 PCIe traffic, and each PCIe port having a corresponding PCIe connector on the connectivity card;  
detecting connectivity issues in one of the plurality of PCIe ports; and  
mitigating the connectivity issues by reconfiguring a communication pathway in a  
PCIe switch circuit on the connectivity card for at least a portion of the PCIe traffic affected  
10 by the connectivity issues, the PCIe switch circuit configured to communicatively couple the plurality of PCIe ports to a shared connectivity interface carried over a host connector of the connectivity card.
9. The method of claim 8, wherein mitigating the connectivity issues comprises:  
15 determining an identity of an external PCIe port affected by the connectivity issues; and  
redirecting traffic from the affected external PCIe port to a different external PCIe port.
- 20 10. The method of claim 8, wherein mitigating the connectivity issues comprises:  
determining which external PCIe ports are affected by the connectivity issues and which external PCIe ports are unaffected by the connectivity issues; and  
redirecting PCIe traffic from the affected external PCIe ports to the unaffected  
external PCIe ports.  
25
11. The method of claim 8, wherein the plurality of PCIe connectors comprises at least four PCIe connectors.
12. The method of claim 8, wherein the connectivity interface and the host connector  
30 comprise a PCIe link with multiple lanes.
13. The method of claim 12, wherein the PCIe link has four lanes.

14. The method of claim 8, wherein the connectivity interface and the host connector comprise at least two PCIe links.

15. One or more non-transitory computer-readable media having stored thereon  
5 program instructions to facilitate operating a connectivity card, wherein the program instructions, when executed by a computing system, direct the computing system to at least:  
monitor a plurality of external Peripheral Component Interconnect Express (PCIe) ports for connectivity issues, each of the plurality of PCIe ports capable of carrying PCIe traffic, and each PCIe port having a corresponding PCIe connector on the connectivity card;  
10 detect connectivity issues in one of the plurality of PCIe ports; and  
mitigate the connectivity issues by reconfiguring a communication pathway in a PCIe switch circuit on the connectivity card for at least a portion of the PCIe traffic affected by the connectivity issues, the PCIe switch circuit configured to communicatively couple the plurality of PCIe ports to a shared connectivity interface carried over a host connector of the  
15 connectivity card.

16. The one or more non-transitory computer-readable media of claim 15, wherein the program instructions direct the computing system to mitigate the connectivity issues by:  
determining an identity of an external PCIe port affected by the connectivity  
20 issues; and  
redirecting traffic from the affected external PCIe port to a different external PCIe port.

17. The one or more non-transitory computer-readable media of claim 15, wherein the  
25 program instructions direct the computing system to mitigate the connectivity issues by:  
determining which external PCIe ports are affected by the connectivity issues and which external PCIe ports are unaffected by the connectivity issues; and  
redirecting PCIe traffic from the affected external PCIe ports to the unaffected  
external PCIe ports.

30

18. The one or more non-transitory computer-readable media of claim 15, wherein the plurality of PCIe connectors comprises at least four PCIe connectors.

19. The one or more non-transitory computer-readable media of claim 15, wherein the connectivity interface and the host connector comprise a PCIe link with multiple lanes.

20. The one or more non-transitory computer-readable media of claim 19, wherein the  
5 PCIe link has four lanes.



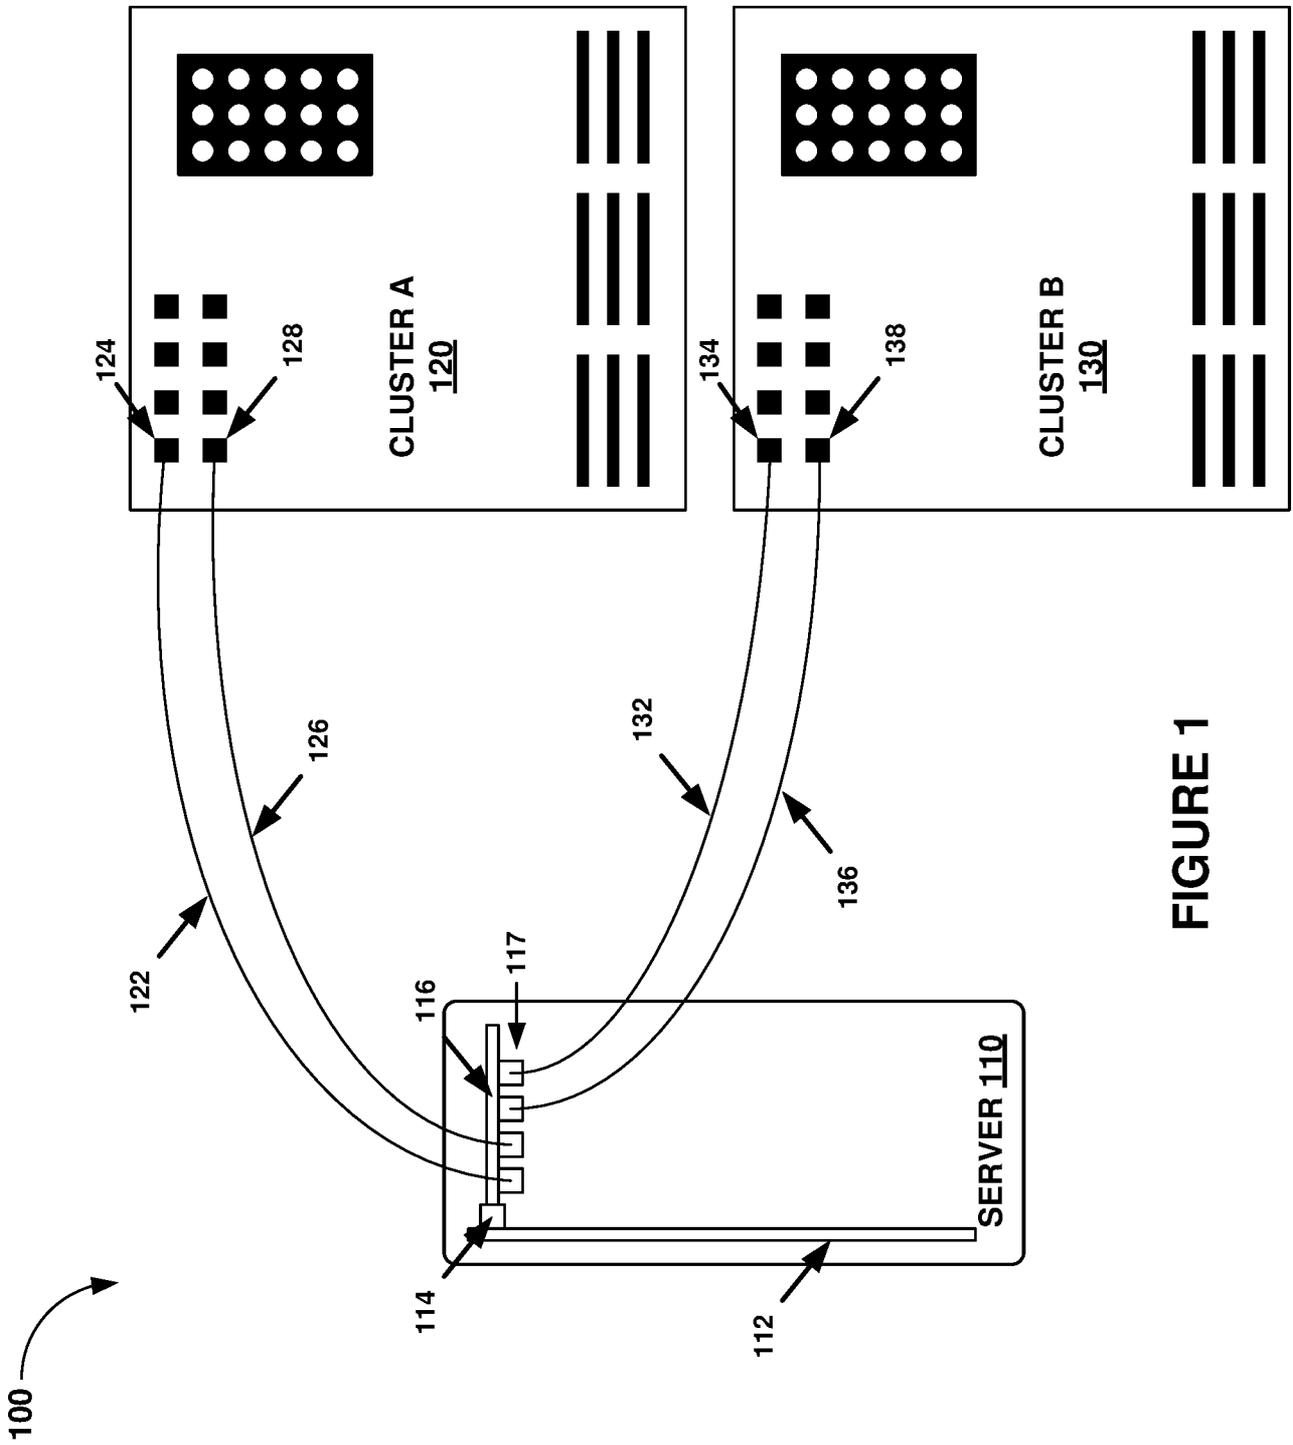


FIGURE 1

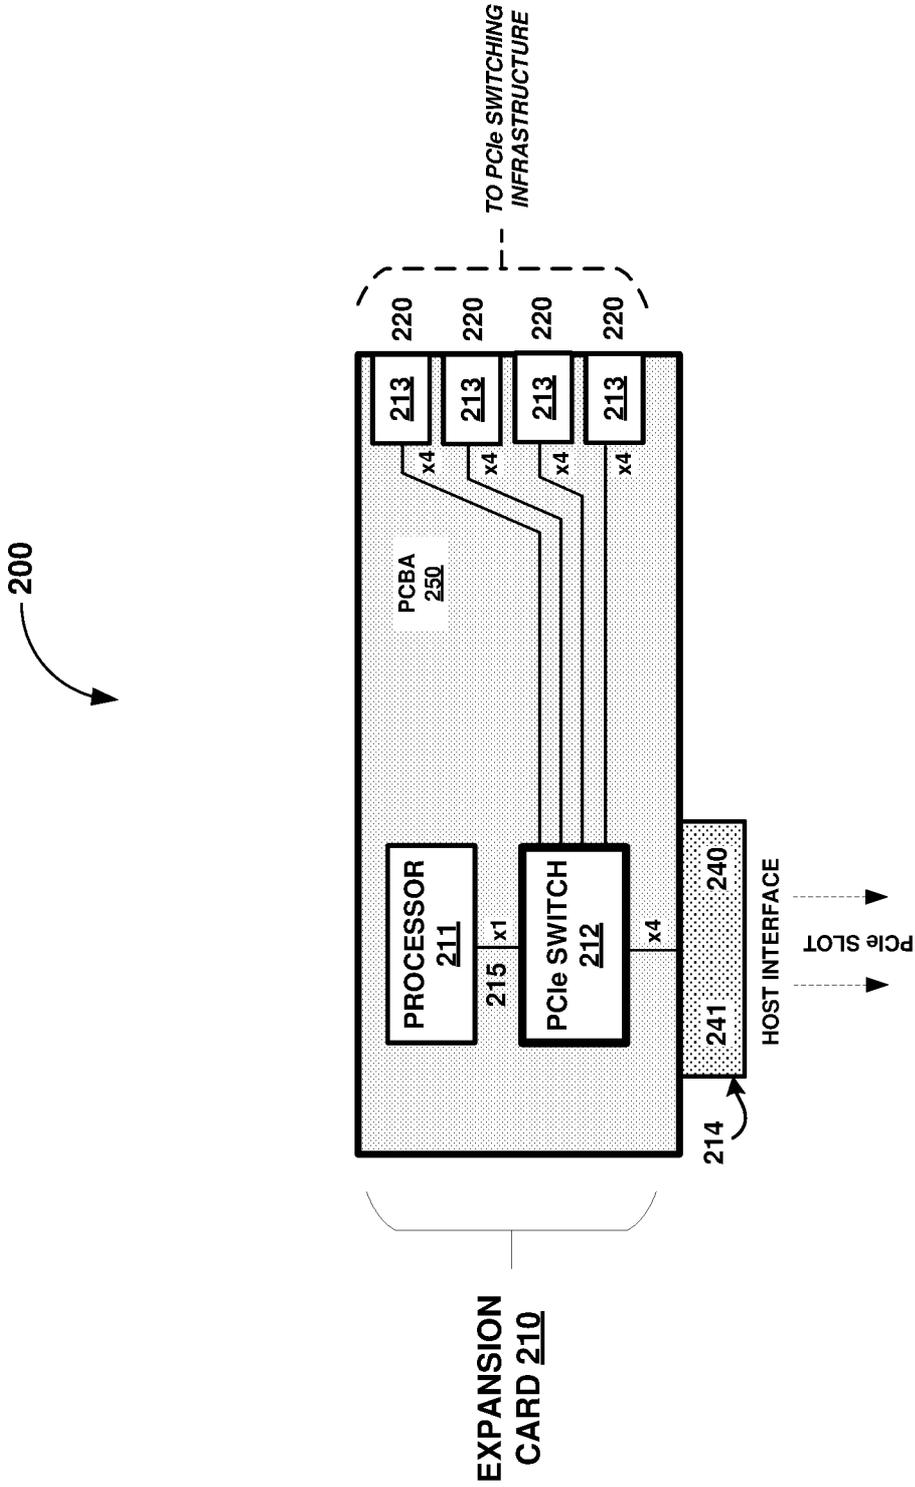


FIGURE 2

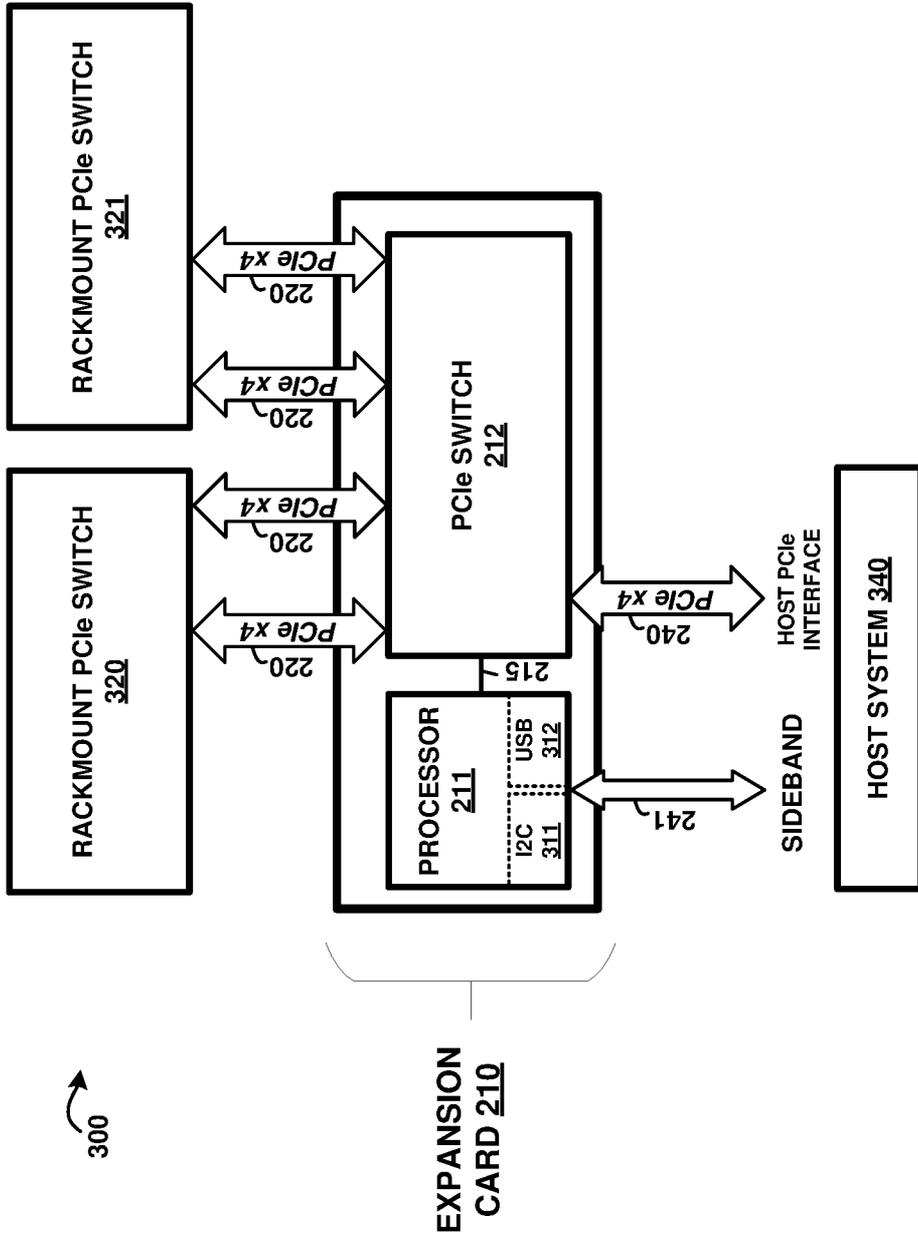


FIGURE 3

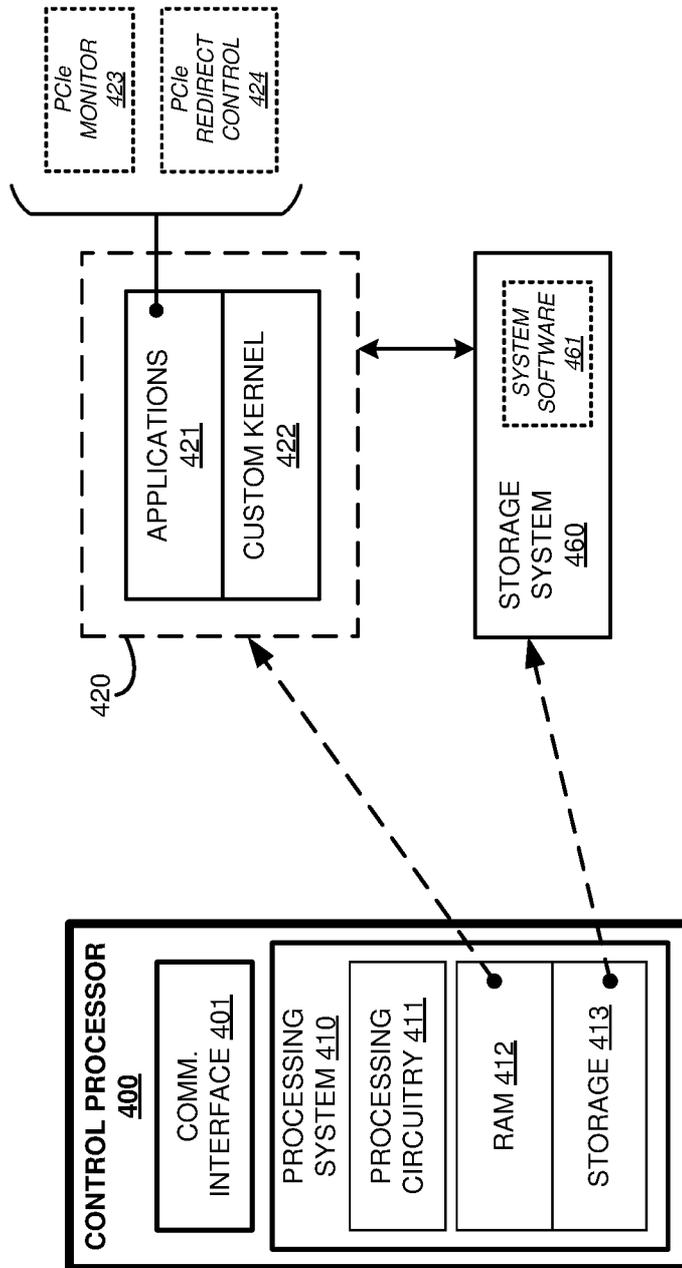
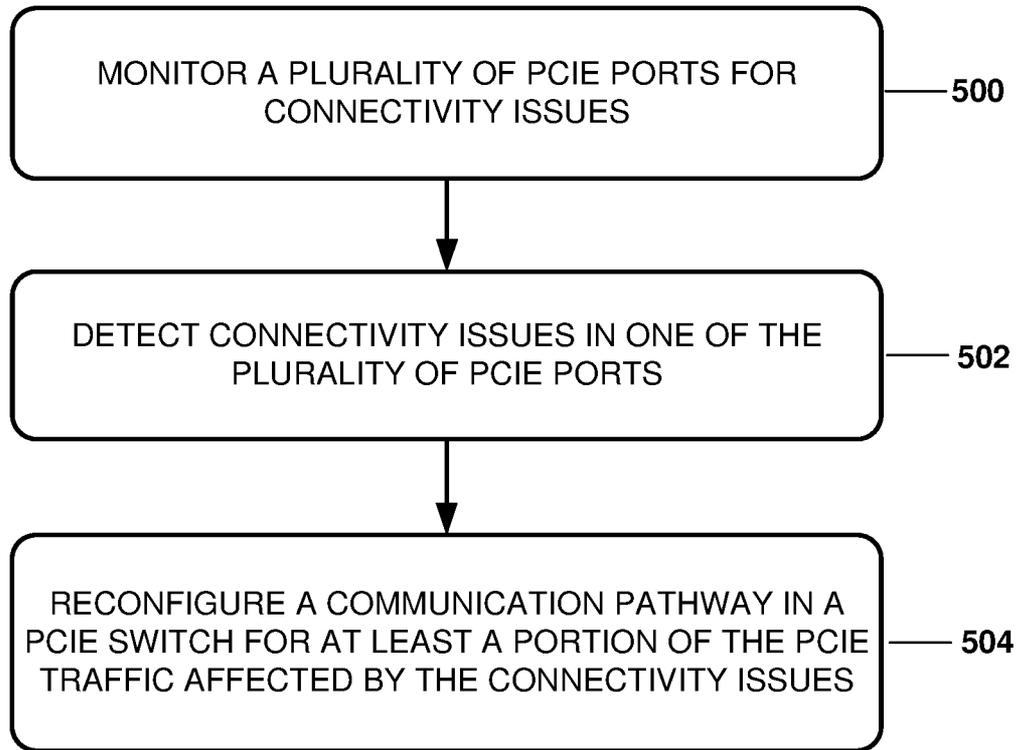


FIGURE 4



**FIGURE 5**

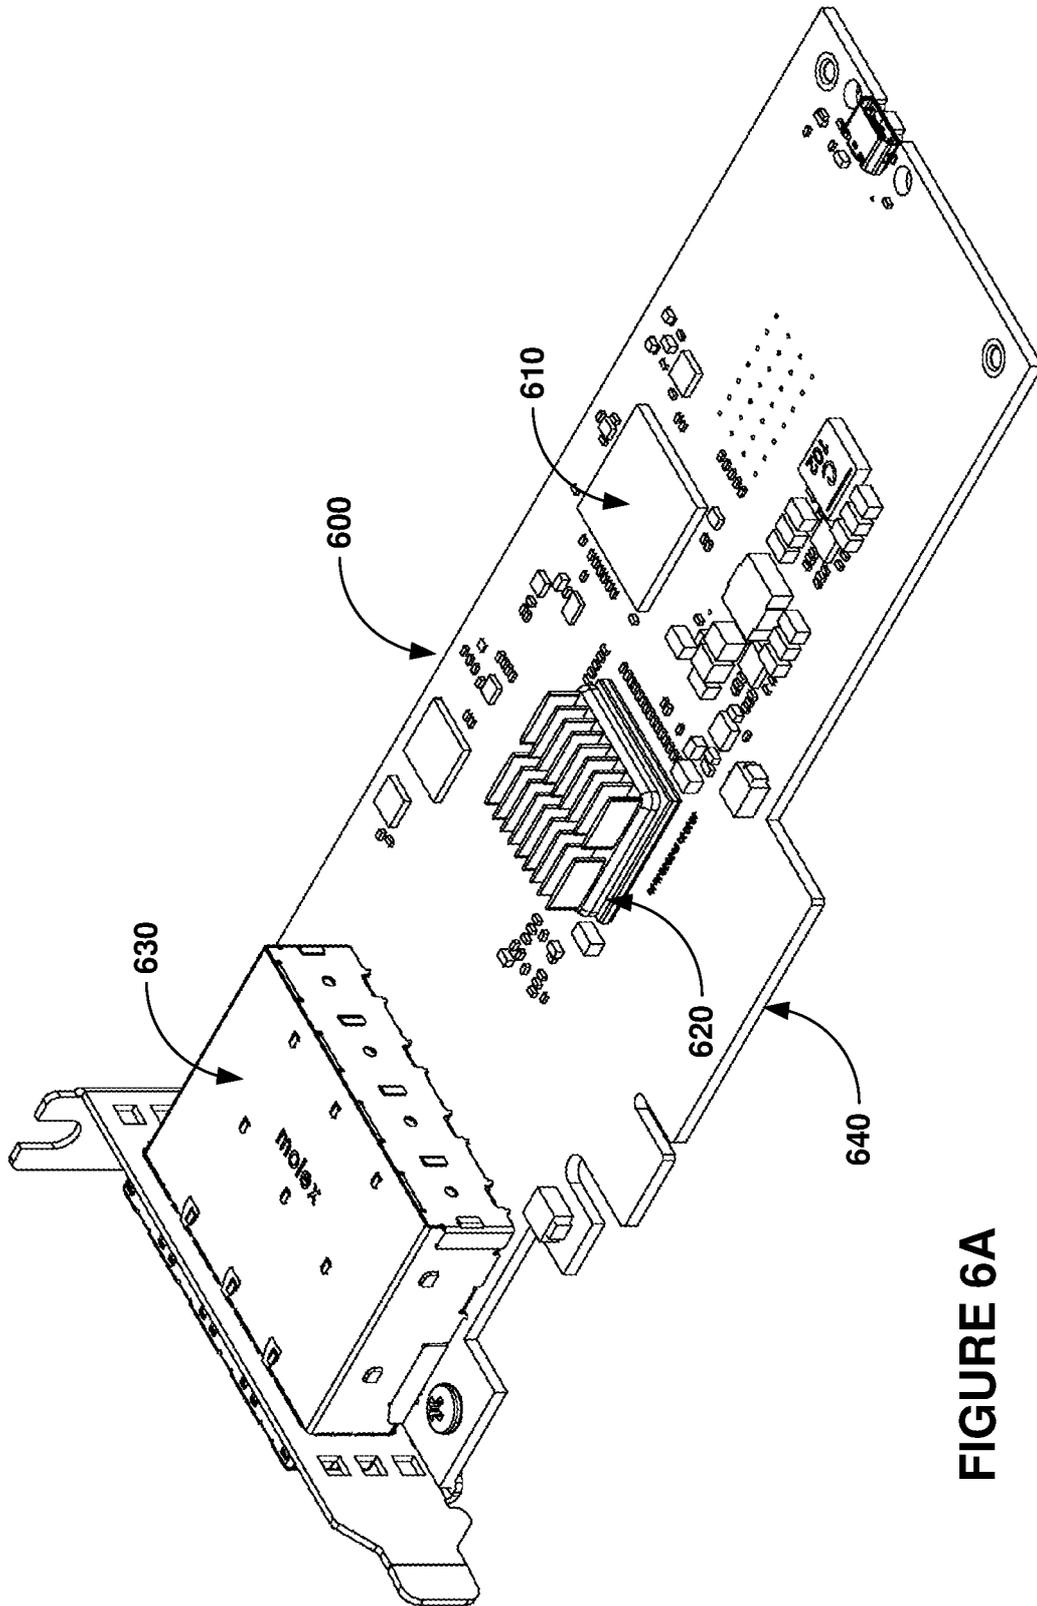


FIGURE 6A

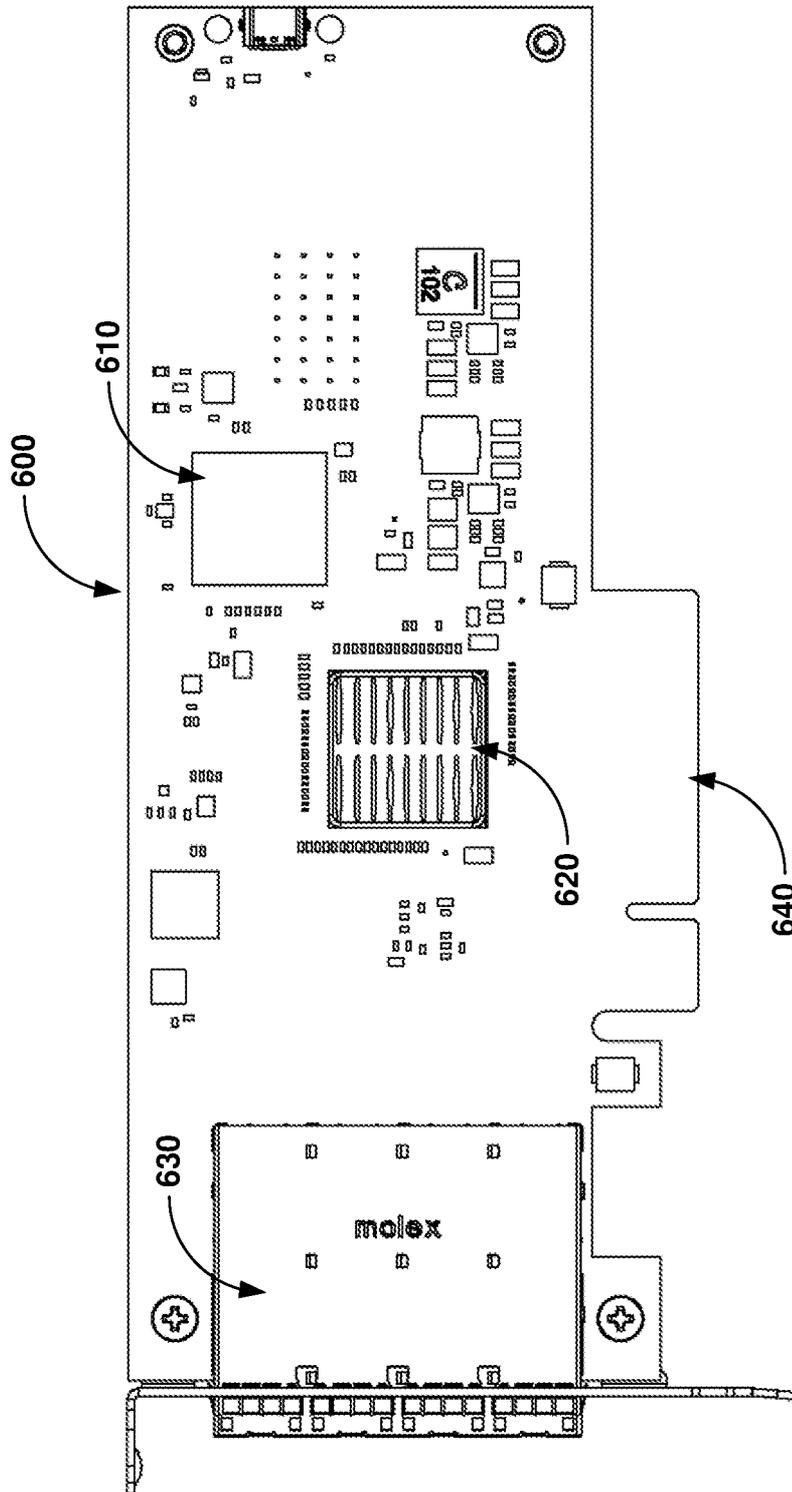


FIGURE 6B

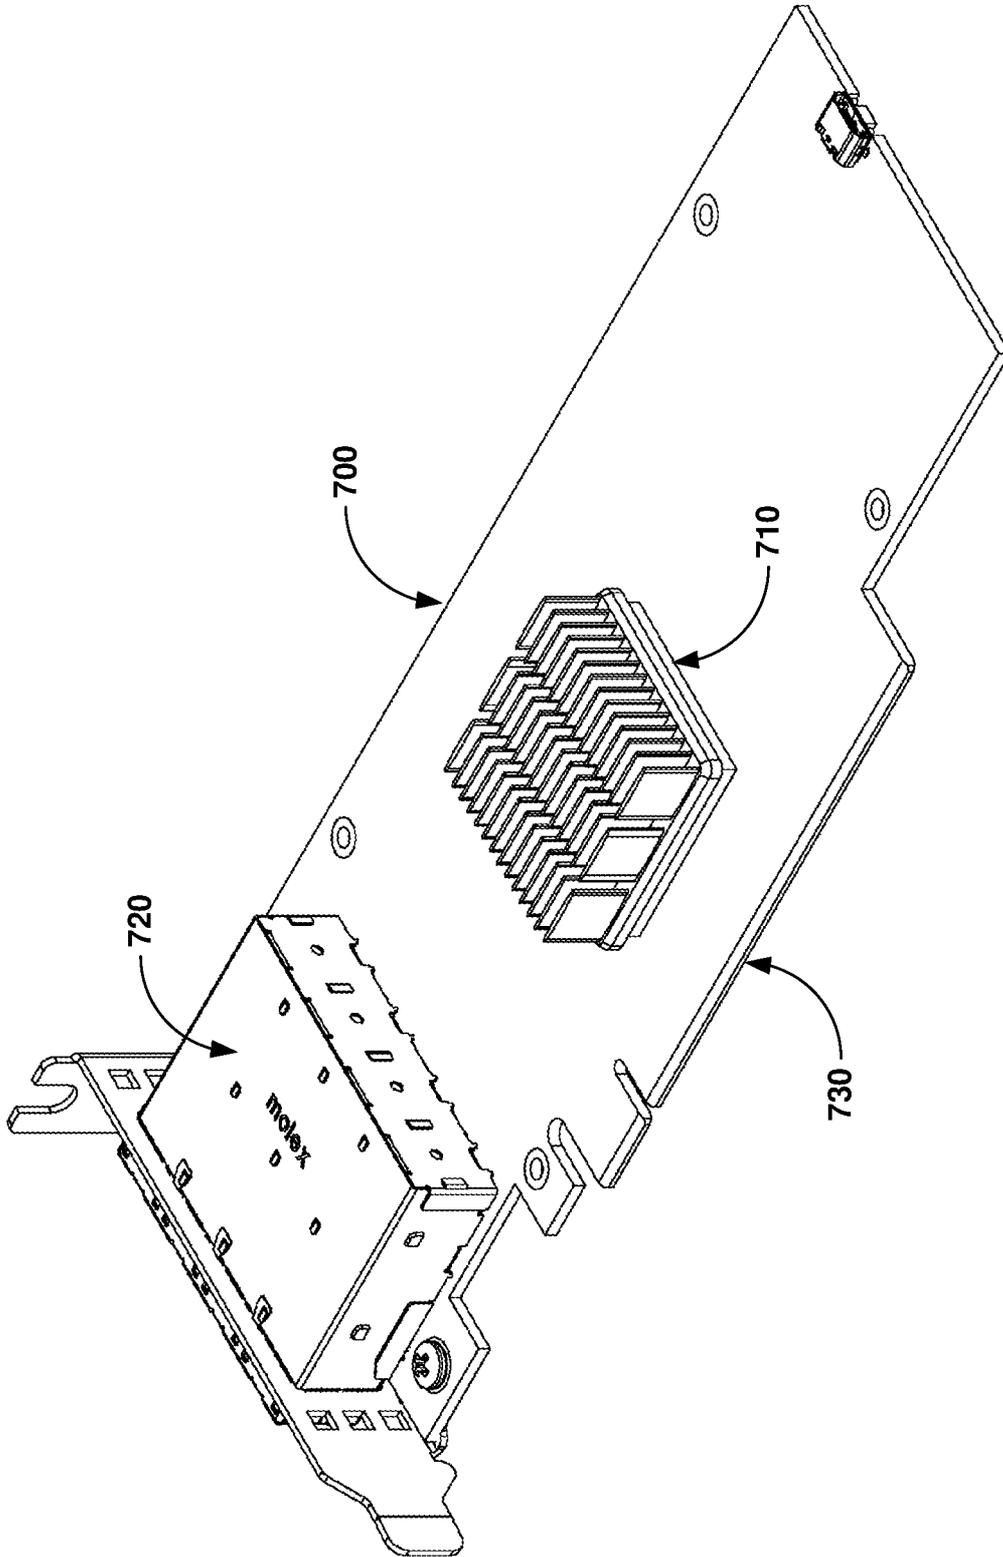


FIGURE 7A



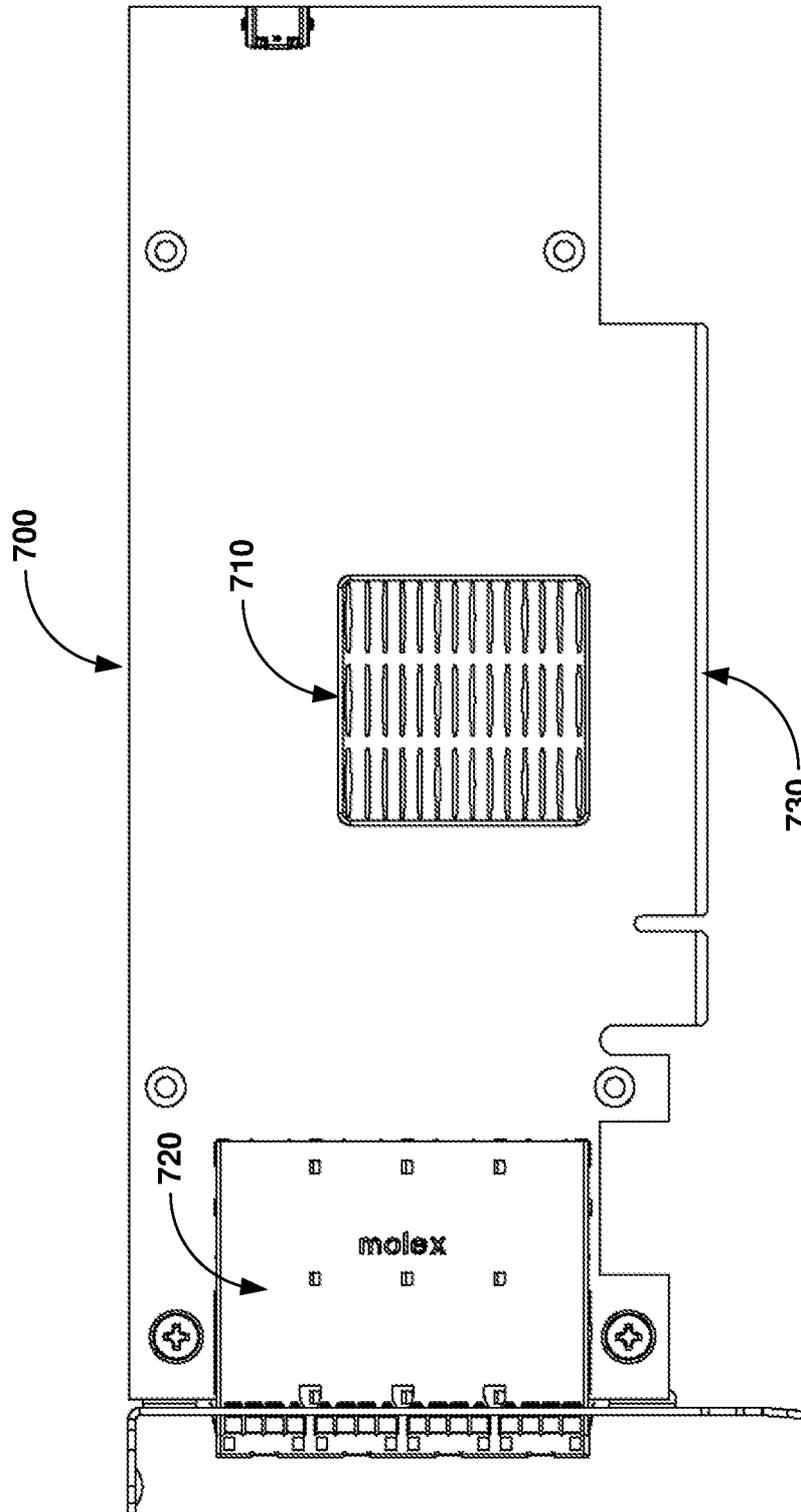


FIGURE 7B

## INTERNATIONAL SEARCH REPORT

International application No.

PCT/US18/29496

## A. CLASSIFICATION OF SUBJECT MATTER

IPC - G06F 11/07, 11/26, 11/30, 13/40, 13/38 (2018.01)

CPC -

G06F 11/1423, 11/142, 11/0793, 13/4022, 13/4027, 13/4063, 13/4068, 13/38

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)

See Search History document

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

See Search History document

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

See Search History document

## C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2015/0347345 A1 (CIRRASCALE CORPORATION) December 3, 2015; figures 4, 9, paragraphs [0015], [0023], [0044], [0045], [0051], claims 1, 8, 12	1-20
Y	US 2016/0147628 A1 (INTERNATIONAL BUSINESS MACHINES CORPORATION) May 26, 2016; figure 1, paragraphs [0007], [0030], [0035], [0063], [0065], [0068], [0069], claim 1	1-20
Y	US 2016/0072640 A1 (YANG, J et al.) March 10, 2016; paragraphs [0052], [0089], claims 1, 5	2, 3, 9, 10, 16, 17
Y	US 2016/0070661 A1 (QUANTA COMPUTER INC.) March 10, 2016; figure 2, paragraphs [0015], [0020]	5, 6, 12, 13, 19, 20
Y	US 8,756,360 B1 (RICHARD, J) June 17, 2014; column 7, lines 15-22, claim 3	7, 14

Further documents are listed in the continuation of Box C.

See patent family annex.

* Special categories of cited documents:	"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
"A" document defining the general state of the art which is not considered to be of particular relevance	"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
"E" earlier application or patent but published on or after the international filing date	"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
"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)	"&" document member of the same patent family
"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	

Date of the actual completion of the international search

21 June 2018 (21.06.2018)

Date of mailing of the international search report

10 JUL 2018

Name and mailing address of the ISA/

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents  
P.O. Box 1450, Alexandria, Virginia 22313-1450  
Facsimile No. 571-273-8300

Authorized officer

Shane Thomas

PCT Helpdesk: 571-272-4300  
PCT OSP: 571-272-7774