




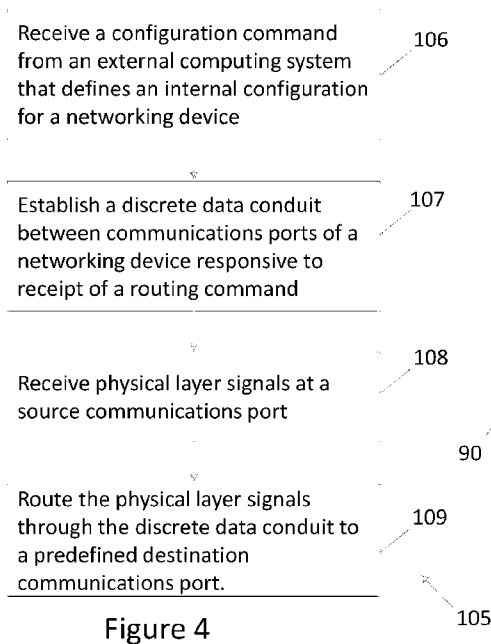
- (51) **International Patent Classification:**
H04L 12/931 (2013.01) G06F 13/38 (2006.01)
H04L 12/00 (2006.01)
- (21) **International Application Number:**
PCT/AU2015/000311
- (22) **International Filing Date:**
26 May 2015 (26.05.2015)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
2014202865 26 May 2014 (26.05.2014) AU
- (71) **Applicant:** ZOMOJO PTY LTD [AU/AU]; Level 6, 76-80 Clarence Street, Sydney, New South Wales 2000 (AU).
- (72) **Inventor:** CHAPMAN, Matthew; Level 6, 76-80 Clarence Street, Sydney, New South Wales 2000 (AU).
- (74) **Agent:** GRIFFITH HACK; Gpo Box 1285, Melbourne, Victoria 3001 (AU).
- (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).

Published:
— with international search report (Art. 21(3))

(54) **Title:** A TRADING SYSTEM



(57) **Abstract:** This specification discloses a protocol agnostic networking apparatus and method of networking. The networking apparatus receives physical layer signal through a plurality of communications ports that interface with external computing systems. A crosspoint switch interconnects the communications ports with discrete reconfigurable data conduits. Each of the data conduits defines a transmission pathway between predetermined communications ports. A management module maintains the data conduits based on configuration commands received from an external computing system. The management module interfaces with the crosspoint switch to make and/or break data conduits responsive to received configuration commands.

WO 2015/179895 A1

A TRADING SYSTEM

Field of the invention

5 The disclosure of this specification relates generally to low latency data communications, and more particularly (but not exclusively) to a reconfigurable networking system and a method for low latency networking.

Background

Communication latency is a measure of delay between releasing communication transmissions
10 and the transmissions being received. Latency in data communication networks is introduced by physical limitations of the networking equipment (including transmission mediums) and transmission signal processing procedures implemented during sending and receipt. The latency of communications may, in some applications, be an important factor in securing a desired outcome. For example, being the first to have a complete electronic trade order received by a
15 stock exchange can establish trade priority. Advantage can then be taken of favourable prices for financial instruments, such as shares, derivatives and futures. Being the first to receive market information may enable a trader to take advantage of favourable market conditions before others. In another example, the outcome of a networked electronic game for an individual player may be determinant on the latency of gaming commands transmitted across the network. Being able to
20 send a low latency instruction to place a wager or a bid at an auction, for example, may increase the probability of securing good odds, a good price, or a successful purchase.

Summary of the invention

In a first aspect, the present invention provides a networking apparatus comprising:

25 a plurality of communications ports that interface with external computing systems to channel physical layer signals,

a crosspoint switch connected to each of the communications ports that routes physical layer signals within the networking apparatus, the crosspoint switch having a plurality of discrete signal channels that are dynamically interconnectable to form data conduits,

30 a management module that controls routing of physical layer signals responsive to commands from an external computing system, the management module being interfaced with

the crosspoint switch to manipulate data conduits by making and/or breaking connections between signal channels, and

5 a processing unit connected in an internal routing loop with the crosspoint switch that decodes internally routed physical layer signals and processes extracted data prior to reversion to a destination communications port, the internal routing loop connecting the processing unit to a plurality of crosspoint switch signal channels.

In an embodiment, the networking apparatus comprises a mounting socket that releasably supports the processing unit, the mounting socket interfacing the processing unit with the internal routing loop.

10 In an embodiment, the networking apparatus comprises a programmable logic device integrated with the processing unit.

In an embodiment, the networking apparatus comprises a communications interface slot that receives a releasable socket module and interfaces individual port sockets within the socket module with signal channels of the crosspoint switch.

15 In an embodiment, the networking apparatus comprises a routing controller integrated with the management module that configures the crosspoint switch for multiplexing operations involving a defined set of source communications ports to minimize multiplexing stages within the processing unit.

20 In an embodiment, the networking apparatus comprises a routing controller that determines a minimum set of logic blocks within the processing unit for multiplexing operations and channels physical layer signals from designated communications ports to input interfaces at the processing unit that correspond to logic blocks within the minimum set to optimize logic block utilization.

25 In an embodiment, the networking apparatus comprises a switching module that facilitates configuration of individual communications ports via a management interface, the management module interconnecting signal channels within the crosspoint switch to link communications ports responsive to configuration commands received via the switching module.

In an embodiment, the configurations commands supported by the switching module via the management interface include timestamping, link quality monitoring, ingress filtering, quality of service rate limiting, packet counting and port-to-port pass-through routing.

In an embodiment, the networking apparatus comprises a switching module that facilitates protocol classification for individual communication ports routed via the processing unit, the management module using the protocol classification for a communications port to initialize the processing unit to decode signals received via corresponding communications ports and/or
5 encoding signals designated for corresponding communications ports.

In an embodiment, the networking apparatus comprises a plurality of physical layer conduits that interface the processing unit with individual signal channels of the crosspoint switch to form the internal routing loop, each of the physical layer conduits defining a dedicated unidirectional pathway for transmission of physical layer signals between the processing unit and a
10 corresponding signal channel of the crosspoint switch.

In an embodiment, the networking apparatus comprises a plurality of physical layer conduits that interface the processing unit with individual signal channels of the crosspoint switch to form the internal routing loop, the internal routing loop having a forwarding conduit and a return conduit for each of the communications ports.

15 In an embodiment, the networking apparatus comprises a plurality of physical layer conduits that interface the communications ports with individual signal channels of the crosspoint switch, each of the communications ports having two physical layer conduits that define a loop between the communications port and the crosspoint switch.

In an embodiment, the networking apparatus comprises a plurality of physical layer conduits that
20 interface the processing unit and the communications ports with dedicated signal channels of the crosspoint switch, the management module interconnecting channels within the crosspoint switch to establish a transmission pathway between respective physical layer conduits.

In an embodiment, the networking apparatus comprises an independent management port interfaced with the management module, the management port facilitating reception of
25 configuration commands from an external computing system via a management network.

In an embodiment, the networking apparatus comprises a signal monitoring unit that samples physical layer signals received via designated communications ports at defined offsets from a nominal sample point to evaluate signal quality, and determines whether the physical layer signal can be recovered at the sample points.

In an embodiment, the networking apparatus comprises a plurality of clock and data recovery modules that facilitate signal conditioning, each of the communications ports being associated with a clock and data recovery module.

5 In an embodiment, the networking apparatus comprises a plurality of AC couplings that connect the crosspoint switch to each of the clock and data recovery modules.

In an embodiment, the networking apparatus comprises a reference clock that is interfaced with each of the clock and data recovery modules.

In a second aspect, the present invention provides a networking method comprising:

10 a networking apparatus interconnecting channels of a crosspoint switch to form reconfigurable data conduits responsive to receipt of configuration commands from an external computing system, the data conduits interconnecting a source communications port and a destination communications port in an internal routing loop that includes a processing unit,

15 the networking apparatus receiving physical layer signals at the source communications port and routing the received physical layer signals from the source communications port to the processing unit via a forwarding data conduit,

the networking apparatus causing the processing unit to decode physical layer signals, process data extracted from the physical layer signals and re-encode extracted data using a communications protocol that is compatible with a designated destination communications port, and

20 the networking apparatus routing the re-encoded physical layer signals from the processing unit to a designated destination communications port via a return data conduit.

In an embodiment, the method comprises:

the networking apparatus routing physical layer signals from a plurality of source communications ports to the processing unit via distinct forwarding data conduits,

25 the processing unit extracting data from the physical layer signals and multiplexing extracted data received via each of the distinct forwarding data conduits into a multiplexed signal stream, and

the networking apparatus routing the multiplexed signal stream to a destination communications port via a return data conduit.

In an embodiment, the method comprises the networking apparatus causing the processing unit to decode physical layer signals received via a plurality of source communications ports, extract data frames from the decoded signals, interleave frames received via each of the communications ports in a multiplexed data stream and encode the multiplexed data stream using a
5 communications protocol that is compatible with a destination communications port.

In an embodiment, the method comprises:

the networking apparatus determining a routing configuration for the crosspoint switch that minimizes multiplexer stages within the processing unit for a defined set of source communications ports, and

10 the networking apparatus interconnecting distinct channels within the crosspoint switch to form forwarding data conduits that conform to the determined routing configuration.

In an embodiment, the method comprises:

the networking apparatus determining a minimum set of logic blocks within the processing unit for multiplexing operations, and

15 the networking apparatus interconnecting distinct channels within the crosspoint switch to channel physical layer signals from designated communications ports to input interfaces at the processing unit that correspond to logic blocks within the minimum set to optimize logic block utilization.

In an embodiment, the method comprises the networking apparatus causing the processing unit
20 to timestamp the data extracted from each of the physical layer signals prior to multiplexing the extracted data.

In an embodiment, the method comprises:

the networking apparatus managing protocol classification for communications ports routed via the processing unit, including receiving protocol classifications for individual
25 communications ports, and

the networking apparatus using the received protocol classifications to initialize the processing unit for compatibility with a corresponding communications ports to facilitate decoding received signals and/or encoding data for transmission.

In an embodiment, the method comprises the networking apparatus facilitating configuration of individual communications ports, including receiving configuration commands from an external computing system, and reconfiguring the crosspoint switch responsive to the received configuration commands.

- 5 In an embodiment, the method comprises the networking apparatus breaking data conduits within the crosspoint switch responsive to received configuration commands to prevent routing of physical layer signals from a source communications port.

In an embodiment, the method comprises the networking apparatus autonomously routing physical layer signals received via a source communications port to the processing unit
10 responsive to receiving a configuration command that specifies multiplexing, timestamping, link quality monitoring, ingress filtering, quality of service rate limiting or packet counting operations.

In an embodiment, the method comprises the networking apparatus receiving a configuration command from an external computing system and establishing a routing loop between
15 communications ports of the networking apparatus responsive to the received configuration command.

In an embodiment, the method comprises the networking apparatus conditioning physical layer signals before routing the signals to a destination communications port.

In an embodiment, the method comprises the networking apparatus performing a clock recovery
20 and signal regeneration process to condition the physical layer signals.

In an embodiment, the method comprises the networking apparatus shaping degraded physical layer signals.

In an embodiment, the method comprises the networking apparatus replicating received physical layer signals and routing the replicated signals to a plurality of destination communications
25 ports.

Any features of the above disclosed systems, methods, network, and module may be combined, if possible.

Brief description of the figures

Embodiments of the invention will now be described by way of example only with reference to the accompanying figures, in which:

Figure 1 is a schematic diagram of an embodiment of a networking apparatus illustrating the establishment of a data conduit between communications ports;

5 **Figure 2** is a schematic representation of a rack mountable enclosure in which the networking apparatus of figure 1 may be housed;

Figure 3 is a schematic diagram of another embodiment of a networking apparatus that incorporates a plurality of dedicated clock and data recovery (CDR) modules;

10 **Figure 4** is a flow diagram of a networking method that the networking apparatus may implement;

Figure 5 is a schematic representation of a data replication system;

Figure 6 is a schematic representation of a failsafe computing network configuration; and

Figure 7 is a schematic representation of a system for determining latency.

15 **Figure 8a** is a block diagram representing the unstructured signal multiplexing within the processing unit.

Figure 8b is a block diagram representing optimised signal multiplexing within the processing unit.

Detailed description

20 Embodiments of a networking apparatus are disclosed in this specification. The disclosed apparatus facilitates communication between network capable computing systems that are connected to the networking apparatus. A collection of communications ports interface the networking apparatus with external computing systems. The communications ports channel physical layer signal from the corresponding computing systems to a crosspoint switch within the networking apparatus. The networking apparatus forms reconfigurable data conduits within
25 the crosspoint switch to interconnect designated communications ports and facilitate the exchange of data between external computing systems. The data conduits define transmission pathways within the networking apparatus for physical layer signals. The networking apparatus is capable of channeling data signals between external computing systems via the data conduits at layer 1 of the Open Systems Interconnection (OSI) model.

The disclosed networking apparatus has an integrated processing unit that is capable of extracting data from physical layer signals received via the communications ports. Physical layer signals received from designated computing systems are channeled to the processing unit for intermediate processing before transmission to a destination computing system. The processing unit decodes internally routed physical layer signals, processes data extracted from the signals and re-encodes the data for a designated communications port. Typical processing operations implemented by the processing unit can include multiplexing data from a plurality of computing systems for transmission to a selected destination, timestamping, link quality monitoring, ingress filtering, quality of service rate limiting and transmission analysis (such as data packet counts).

The networking apparatus routes all physical layer signals received via the communications ports with the crosspoint switch. The crosspoint switch has a series of discrete signal channels that facilitate transmission of physical layer signals. The networking apparatus interconnects designated signal channels to form data conduits. The data conduits channel physical layer signals between the respective input/output interfaces of the crosspoint switch and the components connected to the respective signal channels. The crosspoint switch may use semiconductor switches, optical splitters and/or other components to channel signals.

The networking apparatus controls the crosspoint switch to channel physical layer signals directly between external computing systems and route signals internally via the processing unit. The crosspoint switch creates suitable connections (such as electrical or optical pathways) to facilitate this functionality. The routing configuration of the crosspoint switch is set by configuration commands that are independent of the physical layer signals channeled by the networking apparatus. Channeling physical layer signals via preconfigured data conduits eliminates the decoding process implemented in conventional network switches to extract routing information (typically contained in packet headers). This can produce significant latency reductions (especially for latency sensitive applications such as high speed trading and network gaming). The disclosed networking apparatus is also capable of performing higher level functions by routing physical layer signals via the internal processing unit. An integrated management module facilitates configuration of the networking apparatus (including configuration of the crosspoint switch and optional processing functions) from an external computing system.

The processing unit is connected in an internal routing loop with the crosspoint switch. The routing loop comprises a series of dedicated physical layer conduits that connect individual

crosspoint switch interfaces with corresponding I/O interfaces at the processing unit. The networking apparatus routes physical layer signals via the processing unit by forming a reconfigurable data conduit between a source communications port and a corresponding input interface. The processing unit decodes physical layer signals received via the source
5 communications port and processes data extracted from the signals. The crosspoint switch controls routing between communications ports and the processing unit. Physical layer signals from a designated source communications ports are routed to the processing unit via preconfigured forwarding conduits. Re-encoded signals are subsequently routed from the processing unit to a destination communications port via a return conduit.

10 The networking apparatus can bypass the processing unit for latency sensitive signals that are not designated for internal processing. Pass-through routing between external systems is executed at layer 1 of the Open Systems Interconnection (OSI) model. The networking apparatus facilitates layer 1 routing by establishing a direct data conduit within the crosspoint switch that channels physical layer signals from a source communications port directly to a destination
15 communications port.

The configuration of the signal channels within the crosspoint switch determines the routing path for physical layer signals received by the networking apparatus. This enables the networking apparatus to route signals without processing signal metadata (such as packet headers). The crosspoint switch eliminates routing overheads associated with higher layer signal transmission
20 protocols (such as TCP/IP and UDP) by channeling signals via predefined pathways (instead of determining signal destinations from associated metadata after the signal is received and decoded).

The crosspoint switch is 'data agnostic'. Physical layer signals are channeled within the networking apparatus via predefined data conduits that are not influenced by higher level
25 transmission protocols. This enables the networking apparatus to channel data signals that adhere to different transmission protocols without prejudice. The transmission protocols accompanying internally routed data signals may be used by the processing unit to facilitate internal processing operations (such as packet counting). The processing unit is typically preconfigured for compatibility with a designated protocol to facilitate this functionality.

30 The disclosed networking apparatus has a plurality of communications ports that facilitate connection with external computing systems. Each of the communications ports incorporate a transceiver module (such as SFP or SFP+) that channels physical layer signals from external computing systems to the crosspoint switch. The networking apparatus may support releasable

transceiver modules that can be interchanged to reflect networking requirements (such as swapping electrical transceivers for optical transceivers).

5 A series of physical layer conduits connect the communications ports with individual channels of the crosspoint switch. The physical layer conduits may interface directly with the transceiver module of a corresponding port or a block interface (such as a socket that receives releasable transceiver modules) that connects to a set of ports.

10 The illustrated physical layer conduits comprise conductive pathways (such as tracks, signal traces, strip lines and/or micro strip lines and wires) between individual port transceivers and signal channel interfaces at the crosspoint switch. Each of the communications ports is typically linked with two crosspoint switch channels (a 'send' channel and a 'receive' channel) via dedicated physical layer conduits. This facilitates bidirectional signal transmission between the ports and the crosspoint switch.

15 The crosspoint switch interconnects designated communications ports with discrete reconfigurable data conduits. Each of the data conduits define a transmission pathway between predefined communications ports for physical layer signals. Physical layer signals can be routed directly between communications ports (a 'pass-through' configuration) or channeled via the processing unit. Both configurations enable physical layer signals to be channeled from one computing system to another computing system with low transmission latency relative to conventional network switches that derive routing information from signal metadata.

20 A management module controls signal routing within the networking apparatus. The management module is interfaced with the crosspoint switch and manipulates data conduits by making and/or breaking connections between signal channels. The routing path for physical layer signals received by the networking apparatus is defined by the configuration of the crosspoint switch (i.e. the data conduits formed by the signal channels).

25 The management module preconfigures the crosspoint switch and processing unit responsive to configuration commands received from an external computing system. The configuration commands are independent of the physical layer signals being channeled by the networking device. The management module makes and/or breaks data conduits within the crosspoint switch responsive to received configuration commands.

30 A switching module facilitates configuration of individual communications ports via a management interface. The switching module may be integrated with the management module or implemented as an independent component. The management interface is typically accessible

via an independent management network (commonly accessible via the management module). The management module interconnects signal channels within the crosspoint switch to link communications ports responsive to configuration commands received via the management interface.

- 5 The processing unit is connected in an internal routing loop with the crosspoint switch. A plurality of physical layer conduits (such as conductive tracks, signal traces, strip lines and/or micro strip lines and wires) convey signals from the crosspoint switch to the processing unit and back from the processing unit to crosspoint switch. Each of the physical layer conduits links an individual channel of the crosspoint switch with a dedicated I/O interface at the processing unit.
- 10 The crosspoint switch establishes data conduits between designated communications ports and the I/O interfaces to channel physical layer signals via the internal routing loop.

Figure 1 shows a schematic diagram of an embodiment of a networking device for use in a data network 12. The networking device 10 can be implemented in a variety of networks, including wide area networks such as the internet, an Ethernet network or Infiniband (TM) network. The
15 illustrated networking device 10 has a plurality of communication ports, generally identified by reference numerals 14 to 30. The communications ports 14 to 30 convey data signals between the networking device 10 and external components (such as external computing systems).

The networking device incorporates a crosspoint switch 32 that channels physical layer signals between the communications ports. Each of the communications ports 14 to 30 is connected to a
20 discrete signal channel within the crosspoint switch. A plurality of dedicated physical layer conduits interface the communications ports 14 to 30 with the respective signal channels. The crosspoint switch 32 establishes connections between individual signal channels to form discrete reconfigurable data conduits. The data conduits interconnect the respective communication ports 14 to 30 to facilitate signal transmissions.

25 **Pass-through routing between communications ports**

The networking device 10 is capable of routing physical layer signals directly between communications ports via the crosspoint switch 32 in a 'pass-through' routing configuration. The networking device 10 facilitates pass-through routing by linking designated communications ports with discrete data conduits formed between corresponding signal channels within the
30 crosspoint switch 32. The networking device 10 is capable of operating at layer 1 of the Open Systems Interconnection (OSI) model (true physical-layer networking) during pass-through routing.

Pass-through routing is protocol independent (“protocol agnostic”). Physical layer signals are channeled through the crosspoint switch 32 without decoding or interpreting signal metadata. Existing ‘physical layer’ devices often incorporate higher level signal interpretation (such as retrieving header information from transmitted signals) despite claiming ‘physical layer’ or
5 ‘layer 1’ operation. Higher level signal interpretation can require logical operations that introduce latency.

The data conduits formed within the crosspoint switch 32 define a pathway through the networking device 10 for physical layer signals. Physical layer signals are routed from a source communications port to a destination communications port via a corresponding data conduit
10 without processing. The crosspoint switch 32 shown in Figure 1 is configured to connect ports 14 and 16 in a ‘one-to-one’ routing configuration with a dedicated data conduit 34. The data conduits formed by the networking device 10 are readily reconfigurable.

The networking device 10 can facilitate both unidirectional and bidirectional data channeling in ‘one-to-one’ routing configurations. The crosspoint switch 32 may also establish ‘one-to-many’
15 port configurations with conduits that interconnect more than two ports. This is facilitated by connecting a ‘source’ signal channel to multiple ‘destination’ signal channels within the crosspoint switch. A ‘one-to-many’ routing configuration is also depicted in Figure 1. Ports 18, 20 and 22 are interconnected by conduit 36.

‘One-to-many’ routing configurations are typically used where data distribution equality is
20 important. Physical layer signals from a source communications port are simultaneously channeled to each destination communications port with negligible discrimination. The networking device 10 is restricted to unidirectional data channeling in pass-through ‘one-to-many’ routing configurations

Embodiments of the layer 1 networking device 10 disclosed in this specification are capable of
25 operating over a wide range of bit rates. For example, the networking device 10 may be compatible with any one of 10BASE5 Ethernet bit rate, 10BASET Ethernet bit rate, Fast Ethernet, Gigabit Ethernet, 10 Gigabit Ethernet and 100 Gigabit Ethernet.

The networking device 10 illustrated in Figure 1 incorporates a management module 50 that facilitates configuration of the crosspoint switch 32. The management module 50 receives
30 configuration commands from an external computing system. The configuration commands define the routing ‘state’ of the crosspoint switch. The management module 50 makes and/or

breaks connections between signal channels within the crosspoint switch responsive to received configuration commands.

Commands that affect routing changes within the networking device 10 can be restricted to an isolated management network. Typical configuration commands define signal channel

5 connections within the crosspoint switch 32 to link designated communications ports and initiate internal routing via the processing unit.

The illustrated management module 50 comprises a control system (such as a microprocessor or dedicated control circuitry) and a communications interface (such as a serial and/or Ethernet interface 52). The control system regulates operation of the networking device 10, including

10 configuration of the crosspoint switch 32. The communications interface connects the management module to an external computing system or network to facilitate service communications (such as diagnostic exchanges and reception of configuration commands).

The management module 50 makes and/or breaks data conduits within the crosspoint switch 32 responsive to configuration commands received from an external computing system. This

15 enables the networking device 10 to be configured without physical interaction. The management module 50 may also control active indicia 54 (such as LED status lights) mounted to the exterior of the enclosure for low level fault detection (such as faults effecting the management module communication interface 52).

A flow diagram representation of a pass-through networking process is illustrated in Figure 4.

20 The flow diagram 105 represents a protocol agnostic process for channeling physical layer signals between communications ports. The method illustrated comprises four operations 106, 107, 108, 109. The operations documented in the flow diagram 105 are:

- Receiving a configuration command from an external computing system (operation 106). The command is transmitted independently of signals being channeled through the device

25 and may be received by a dedicated communications interface (such as the management module interface 52). The configuration command defines an internal configuration for the networking device.

- Establishing a discrete data conduit between communications ports of the networking

30 device responsive to receipt of a configuration command (operation 107). The configuration command may define an internal configuration for the networking device (such as the configuration command identified in operation 106) or cause the networking

device 10, 60 to revert to a predefined default configuration (such as an initialization state or a failsafe configuration).

- 5 • Receiving physical layer signals at a source communications port (operation 108). The networking device 10, 60 is capable of processing signals that adhere to different protocols as the physical layer signals are channeled without extracting transmission information (such as signal headers).
- 10 • Routing the physical layer signals through the discrete data conduit to a predefined destination communications port. The data conduit linking the source communications port to the destination communications port is established in advance of signal reception as the signal transmission pathway is configured independently of the physical layer signal being channeled.

15 Operation 106 of the flow diagram (receipt of configuration commands) may be bypassed when the operating configuration of the networking device 10, 60 satisfies the routing configuration needed for scheduled signal transmissions (i.e. no routing changes are required to channel scheduled transmissions). Configuration commands may be transmitted to the networking device 10, 60 to initiate reconfiguration of data conduits. They may also be used to initialize the device (although the networking devices 10, 60 may store preconfigured initialization states).

20 The networking device 10, 60 may implement other operations that supplement the general networking functionality summarized in Figure 4. These operations include:

- breaking data conduits responsive to received configuration commands to prevent transmission of physical layer signals from a source communications port to a destination communications port.
- 25 • adding a destination communications port to an established data conduit responsive to a configuration command to expand the distribution of physical layer signals from a source communications port.
- distributing a communication received by a service port to a plurality of client ports; and
- conditioning a received physical layer signal before routing to a destination.
- 30 communications port by performing a clock recovery and signal regeneration process to condition to received physical layer signal and/or shaping the degraded physical layer signal.

Internal signal routing and processing

The networking device 10 depicted in Figure 1 is also capable of internally processing received signals prior to transmission. A dedicated processing unit 66 facilitates internal processing operations. Typical internal processing operations include multiplexing, timestamping, link
5 quality monitoring, ingress filtering, quality of service rate limiting and transmission analysis (such as data packet counts). The processing unit 66 decodes internally routed physical layer signals and processes data extracted from the signal.

The processing unit 66 is connected in an internal routing loop 68 with the crosspoint switch 32. The illustrated processing unit 66 comprises a processor 66a (commonly an integrated circuit or
10 programmable logic device such as Field Programmable Gate Array) and a complimentary mounting socket 66b that supports the processor 66a. The mounting socket 66b releasably engages the processor 66a, providing a physical layer interface with the crosspoint switch 32 (via the internal routing loop). The illustrated routing loop 66 connects the processing unit 66 in a closed circuit with the crosspoint switch 32 (i.e. the crosspoint switch 32 exclusively interfaces
15 the processing unit 66 with the communications ports). The mounting socket 66b may also interface the processor 66a with the management module 50 and other components (such as a power supply). The processor 66a implements the processes functionality of the processing unit 66 (such as processing data extracted from physical layer signals received from the crosspoint switch 32). The configuration of the illustrated processing unit 66 enables the processor 66a to
20 be removed for diagnostic testing and/or replacement (including general upgrades).

The crosspoint switch 32 routes physical layer signals from designated communications ports to the processing unit 66. Processed signals are routed back from the processing unit 66 to a destination port via the crosspoint switch. The crosspoint switch 32 channels physical layer
25 signals via the internal routing loop 68 by forming data conduits between designated communications ports and selected I/O interfaces at the processing unit 66 (defined by the mounting socket 66b).

A switching module facilitates configuration of individual communications ports via a management interface. The switching module is integrated with the management module 50 depicted in Figure 1. The management interface is accessible via the management module
30 network interface 52. The management module 50 configures the processing unit 66 for compatibility with designated communications ports responsive to configuration commands received via the management interface.

The switching module facilitates protocol classification for individual communication ports routed via the processing unit 66. The management module 50 uses the protocol classification for each communications port to configure the processing unit 66 for compatibility with the designated communications ports (including initializing the processing unit 66 to decode signals received via corresponding communications ports and/or encode signals designated for corresponding communications ports). The protocol classification may define a networking protocol (such as TCP/IP or UDP) or a transmission protocol (Fast Ethernet, Gigabit Ethernet, 10 Gigabit Ethernet or 100 Gigabit Ethernet). The networking device 10 is capable of exchanging data between non-compatible network nodes through protocol conversion. The processing unit 66 facilitates protocol conversions by extracting data from physical layer signals received from a source network and re-encoding the data with a communications protocol that is compatible with a designated destination network. The networking device 10 may implement this process for incompatible networking protocols and/or transmission protocols.

A series of physical layer conduits interface the crosspoint switch 32 with the processing unit 66. Each of the physical layer conduits connects an individual signal channel within the crosspoint switch to a dedicated I/O interface at the processing unit. A set of 'forwarding' physical layer conduits define transmission pathways between the crosspoint 32 switch and individual input interfaces at the processing unit 66. A set of 'return' physical layer conduits define transmission pathways between output interfaces at the processing unit 66 and individual signal channels within the crosspoint switch 32. The internal routing loop 68 ideally has a forwarding conduit and a return conduit for each of the communications ports to optimize bandwidth (although other configurations are possible).

Individual communications ports are designated for internal processing via commands from an external computing system. The management module 50 configures the crosspoint switch 32 and processing unit 66 responsive to received commands. The commands can define the 'state' of the crosspoint switch 32, prescribe internal processing operations (implemented by the processing unit 66), identify the communications protocol used by the signals (to facilitate defined processing operations) and supply other configuration parameters (such as prioritization ratings for latency sensitive signals). The networking device 10 facilitates independent routing and processing configuration for each of the communications ports. The processing unit 66 decodes internally routed physical layer signals, processes extracted data and encodes the data for transmission to a designated destination communications port. Typical processing operations include multiplexing, timestamping, link quality monitoring, ingress filtering, quality of service rate limiting and packet counting.

The management module 50 is capable of determining an internal routing configuration for the networking device 10 that facilitates the internal processing designated by received configuration commands. This enables the networking apparatus 10 to optimize the interconnection of signal channels (within the crosspoint switch 32) with I/O interfaces (at the processing unit 66) for designated processing operations.

A dedicated routing controller (not shown) integrated with the management module 50 implements a routing optimization algorithm to determine the routing configuration of the crosspoint switch 32 for internal processing. The optimization algorithm is typically derived from the processing capacity of the processor 66a (such as available I/O interfaces and internal processing capabilities) and the interface bandwidth provided by the internal routing loop. The routing controller determines a routing configuration for the crosspoint switch 32 that optimizes utilization of logic blocks within the processing unit for internal processing operations (such as multiplexing). The management module 50 uses the routing configurations determined by the routing controller to channel physical layer signals from designated communications ports to input interfaces at the processing.

A block representation of an unstructured multiplexing operation 150 with four input signals 161, 164, 166, 169 is depicted in Figures 8a. The depicted operation is implemented in three distinct multiplexing stages (represented by mux blocks 151, 152, 153). Each stage equates to a distinct logic block within the processor 66a. The input channels available for each multiplexing stage are limited by the capacity of logic blocks within the programmable logic device. The illustrated logic blocks (configured as multiplexers 151, 152, 153) each have four signal channel inputs. Service inputs (such as a selector channel) are not depicted.

The programmable logic device represented in Figure 8a has ten input interface 160 to 169. Each of the input interfaces 160 to 169 is connected to a crosspoint switch by a dedicated physical layer conduit. The crosspoint switch routes physical layer signals to the processing unit via the physical layer conduits. The set of input interfaces used for processing operations is determined by the configuration of the crosspoint switch. The management module 50 routes physical layer signals from individual communications ports to dedicated input interfaces at the processing unit by interconnecting the respective signal channels within the crosspoint switch. The crosspoint switch is configured to route signals to four input interfaces 161, 164, 166, 169 in Figure 8a.

The multiplexing operation 150 depicted in Figure 8a produces a single multiplexed output 159 from four independent input interfaces 161, 164, 166, 169. Physical layer signals from the

respective input interfaces 161, 164, 166, 169 are ‘combined’ over three multiplexing stages (represented by mux blocks 151, 152, 153). The first stage mux block 151 receives physical layer signals from a single input interface 161. The output 151a from the first multiplexing stage (mux block 151) is combined with physical layer signals from two independent input interfaces 164, 166 at the second stage mux block 152. The third stage mux block 153 receives the output 152a of the second multiplexing stage (mux block 152) and a single input interface 169.

The latency performance of some processing operations can be impacted by the selection of input interfaces used to route physical signals to the processing unit. Each of the multiplexing stages 151, 152, 153 depicted in Figure 8a incurs a latency penalty. The latency of each stage is generally attributable to the inherent processing overhead of the corresponding logic blocks and the transmission of physical layer signals between blocks within the programmable logic device. The overall latency of the illustrated multiplexing operation 150 is a cumulative total of individual latencies at each multiplexing stage 151, 152, 153.

The management module 50 is capable of optimizing the routing configuration of the crosspoint switch to reduce processing latency. An optimized multiplexing operation 170 is depicted in Figure 8b. The illustrated multiplexing operation 170 produces a single multiplexed output 179 from four independent input interfaces 160, 161, 162, 163. The output 179 is derived directly from the first mux block 151.

The crosspoint switch is configured to route physical layer signals from four independent communications ports to the processing unit. The management module 50 establishes a ‘forwarding’ data conduit within the crosspoint switch for each of the designated communications ports. The forwarding conduits define a transmission pathway from the respective communications ports to a dedicated input interface 160, 161, 162, 163 at the processing unit.

The designated communications ports are interfaced with the first multiplexer stage (mux block 151) via corresponding input interfaces 160, 161, 162, 163 in Figure 8b. This routing configuration eliminates the subsequent multiplexing stages illustrated in Figure 8a (the second 152 and third 153 stage mux blocks are not utilized in the multiplexing operation 170 depicted in Figure 8b) and the corresponding latency attributable to the subsequent stages.

The management module 50 also establishes a ‘return’ data conduit from the processing unit to a destination communications port for the multiplexed output 179. The return conduit defines a

transmission pathway from an output interface at the processing unit to a designated communications port.

The management module 50 uses the routing optimizations generated by the routing controller to organize data conduits within the crosspoint switch and configure the processing unit. The management module 50 may also modify the configuration of the programmable logic device to reduce latency for defined processing operations.

Structural elements of the illustrated networking device

The illustrated embodiment of the networking device 10 has a multilayer printed circuit board 38. Components are mounted to the circuit board 38 and generally interconnected to each other by physical layer conduits 40 (typically conductive pathways such as tracks, signal traces, strip lines and/or micro strip lines and wires). The illustrated printed circuit board 38 is coupled to various surface mounted and/or through hole components (including the crosspoint switch 32 and mounting socket 66b).

Figure 2 shows a rack mountable enclosure 44 in which the printed circuit board 38 can be housed (although other housings are possible). The multilayered circuit board 38 illustrated in Figure 1 is arranged for horizontal disposition within the rack mountable enclosure, which is in turn configured to be horizontally mounted within a rack. This arrangement reduces the amount of rack space used by the networking device.

The illustrated enclosure occupies a single rack slot. The rack mountable enclosure is ideally mounted in either the top position or the bottom position of a rack. Generally, communications cables approach a rack in a data center from one of the top and bottom ends of the rack. Having the enclosure mounted at the end of the rack approached by the communication cables may result in a relatively lower latency.

A power supply 46 is mounted to the printed circuit board 38 in the embodiment illustrated in Figure 1. The power supply 46 produces a DC operating voltage for the networking device components (typically a relatively low voltage, such as 12, 24 or 48 volts) from an AC source (usually a relatively high voltage associated, such as 110V or 240V mains supply from the electricity grid). A DC regulator 48 (such as a switched mode power supply module) is also mounted to the printed circuit board 38 in the illustrated embodiment. The DC regulator 48 receives the voltage output from the power supply 46 and powers a set of active conductive rails that are integrated with the circuit board 38. The power supply and DC regulator may also be

mounted to the enclosure separate from the printed circuit board 38. The networking device 10 may also be powered by the rails of the rack where the device is mounted.

5 A fan 16 is ideally mounted to the circuit board 38 or the enclosure 44 to regulate the temperature of the networking device components. The fan provides airflow across the multilayer printed circuit board to extract waste heat.

Each of the communication ports 14 to 30 incorporates a network interface, such as a transceiver or dedicated transmitter and receiver modules. A transceiver 15 is illustrated with port 14 in Figure 1. Suitable transceiver modules include GBIC, XFP, XAUI and Small Form Factor Pluggable Plus (SFP+) transceivers. The transceiver modules may be integrated with the
10 networking device 10 or form part of a interchangeable communications interface (such as releasable optical and/or electrical transceiver blocks). The transceivers are arranged to interface with physical layer conduits that connect the networking device 10 with external computing devices. Typical physical layer conduits include optical fiber network cables and/or electrical network cables (such as copper cabling). The transceivers interface with the physical layer
15 conduits to send and receive electromagnetic communications (such as optical and/or an electrical signals).

The illustrated transceivers are configured to engage two fibre optic LC connectors. This enables the networking device 10 to receive electromagnetic communications through one optical fibre and transmit electromagnetic communications through the other (i.e. the networking
20 device 10 uses each optic fibre for unidirectional communications). The transceivers generate electrical signals from received optical signals and communicate the generated electrical signals to electrical conductors integrated with the printed circuit board 38. The transceivers may support various physical layer standards, such as Ethernet protocol standards SONET, Fibre Channel, or other suitable communications standards.

25 The illustrated transceivers are grouped and housed in SFP cage enclosures 17 that are fixed to the printed circuit board 38. The enclosures 17, and consequently the communication ports, are disposed adjacent the periphery 58 of the multilayered printed circuit board 38. The cage structure provides an electrical connection between electrical contacts on the transceivers and conductive tracks 40 in (usually stripline and/or micro stripline tracks) formed on or within the
30 circuit board 38. The enclosures 17 may also act as Faraday cages to reduce electromagnetic interference.

The crosspoint switch 32 is interfaced with the communications ports to facilitate exchange of physical layer signals with the corresponding transceivers. Physical layer signals received by the communications ports are transmitted directly to the crosspoint switch 32. The crosspoint switch 32 channels the physical layer signals between communications ports using discrete data conduits that are preconfigured using an independent interface (the management module 50 facilitates configuration of the crosspoint 56 switch in the illustrated embodiment).

The crosspoint switch 56 ideally has some or all of the following attributes:

- non-blocking switching capabilities,
- programmable input equalization, and
- output pre-emphasis settings.

The data conduits established by the crosspoint switch 32 are ideally unregistered and asynchronous so that there are no phase, frequency or signal pattern restriction imposed by the crosspoint switch 32. An example of a suitable crosspoint switch is the VITESSE (TM) VSC3172 (other crosspoint switches may be used).

The crosspoint switch 32 is capable of channeling physical layer signals to the processing unit 66 prior to transmission to a recipient communications port. The illustrated crosspoint switch 32 and mounting socket 66b are connected by a series of conductive pathways (such as tracks, signal traces, strip lines and/or micro strip lines and wires). The conductive pathways transfer physical layer signals between the respective components. The mounting socket 66b releasably supports a processor 66a, such as a Virtex 5 FPGA from Xilinx.

The networking device 10 establishes discrete data conduits within the crosspoint switch responsive to configuration commands from an external computing device. The configuration commands are independent of the physical layer signals being channeled through the networking device. The data conduits interconnect components within the networking device 10 to channel physical layer signals between external computing systems. The networking device 10 is capable of routing physical layer signals directly between external systems without higher layer processing (such as determining the intended receipt of a data transmission from signal metadata). The signals may be simultaneously routed to a plurality of destinations ports in 'one-to-many' pass-through routing configurations.

A service port 22 is designated for dedicated broadcast applications. The service port 22 receives broadcast communications from a broadcast server. The remaining communication

ports 14 to 20 and 24 to 30 are designated for client computing systems that receive communications from the broadcasts server. This configuration is often used to facilitate nondiscriminatory data dissemination from the server to the client computing systems. The crosspoint switch 32 is typically configured in a 'one-to-many' pass-through routing configuration for server broadcasts, so that communications from the server are simultaneously channeled from the service port 22 to each of the 'client' ports. In financial applications, the broadcast server may disseminate price ranges for financial instruments (such as shares, derivatives and futures) and the quantity of instruments sold in a defined time period.

The illustrated service port 22 is disposed adjacent the crosspoint switch 32. It is also centrally disposed with respect to the remaining client communications ports. This configuration can reduce the average latency of communications travelling between the service port and the client ports. For example, embodiments having the service port to one side of all of the client ports may have higher average and maximum latency between the client ports and the service ports. Having the service port as close as possible to the crosspoint switch 32 may result in a relative reduction in latency.

The crosspoint switch 32 illustrated in Figure 1 is arranged to condition degraded communications that are received by the networking device. Some degree of signal degradation is generally inevitable during high speed communication transmission. Signal degradation increases the probability of data errors. The crosspoint switch 32 shapes the waveform of degraded communications to conform as close as possible with a desired waveform shape (usually comprising a train of square pulses) by adjusting various frequency components of the physical layer signal. The crosspoint switch may also perform a pre-emphasis routine to compensate for expected losses during transmission. Pre-emphasis generally involves overdriving high frequency components of the signal.

Another embodiment of networking device 60 is illustrated schematically in Figure 3. The networking device 60 retains the fundamental structure and operation of the device 10 illustrated in Figure 1, Similar reference numerals are used in Figure 3 to identify components with similar or identical form and/or function to corresponding components identified in Figure 1.

The networking device 60 illustrated in Figure 3 incorporates a signal monitoring unit 65 that evaluates the quality of received physical layer signals. A centralized signal monitoring unit 65 is depicted in Figure 3. The illustrated signal monitoring unit 65 samples physical layer signals received via designated communications ports at defined offsets from a nominal sample point to evaluate signal quality and determines whether received physical layer signals can be recovered

at the sample points. This function may also be facilitated with distributed signal monitoring units (not shown) that are disposed on the data path between individual communications ports and the crosspoint switch.

The illustrated networking device 60 also incorporates a plurality of dedicated clock and data recovery (CDR) modules 62. Each of the clock and data recovery (CDR) modules 62 is
5 interfaced to a communication port. The illustrated clock and data recovery (CDR) modules 62 are AC coupled to the crosspoint switch. A reference clock 64 is mounted to the circuit board and interfaced with each of the clock and data recovery modules in the illustrated embodiment. Tracks within the circuit board connect the reference clock 64 with the individual clock and data
10 recovery (CDR) modules 62. The clock and data recovery (CDR) modules 62 often incorporate dedicated signal monitoring units.

The reference clock 64 generates an approximate frequency reference that is transmitted to each of the dedicated clock and data recovery (CDR) modules. The clock and data recovery (CDR) modules 62 generate a clock from the approximate frequency reference and phase align
15 transitions in the signal stream using an internal phase locked loop. The sampling phase is then determined with reference to an eye diagram structure. The clock and data recovery (CDR) modules regenerate degraded signals to reduce jitter and produce a more open eye relative to the original signal received by the networking device 60.

The networking device 10 may be used to remotely configure and/or reconfigure a data network
20 12 by sending configuration commands from an external computing system. The networks 12 illustrated in Figures 1 and 3 comprise a broadcast server 72 that is interfaced with the service port 22 of the respective networking devices 10, 60. A plurality of client systems 70, 74 are interfaced with client communications ports 18, 30. The service port 22 and the client ports 18, 30 can be remotely linked (by establishing a data conduit between the respective ports) or
25 isolated (by breaking an existing data conduit) from an external computing system by issuing appropriate configuration commands. This allows the dissemination of physical layer signals from the server 72 to the clients 70, 74 to be manipulated from a remote computing system.

Configuration commands that reconfigure the crosspoint switch 32 are typically received and processed by the management module 50. The illustrated management module 50 receives the
30 configuration commands through a dedicated management port 52. These commands may be issued by a remote computing system, allowing the networking device 10, 60 to remain physically isolated.

The ability to remotely configure the networking device 10, 60 is desirable in some applications. For example, a stock exchange may have a secure data center that hosts financial order server(s) and/or client machines participating in an electronic trading network. Physical access to the secure data center is generally restricted during trading hours. Consequently, it can be
5 problematic to reconfigure electronic trading networks that require manual intervention during trading hours. The networking device 10, 60 disclosed in this specification facilitates reconfiguration of the electronic trading network during trading hours by communicating with computing systems on a segregated management network. The management network can often include computing systems located remote of the data center without breaching trading protocols
10 as the management network is isolated from the trading network. The networking device 10, 60 is also advantageous in this exemplary application because it operates at relatively low latency compared to contemporary networking systems. This is particularly desirable in latency sensitive environments.

The networking devices 10, 60 illustrated in Figures 1 and 3 incorporates a client controller 76
15 that monitors client systems for irregularities. The client controller 76 can isolate client systems that exhibit irregularities by breaking the data conduit(s) that interconnect the client system with the network 12. Client irregularities are typically quantified by application specific parameters. The client controller may monitor designated parameters directly or receive updates from a remote system.

20 The client controller 76 isolates client systems when a monitored client parameter satisfies defined disconnection conditions. Disconnection conditions for financial applications can include pre-order risk conditions, post-order risk conditions and financial ordering system subscription conditions. Breaking data conduits in these situations disconnect the client system from a financial order server. 'On-board' client controllers (such as the client controller
25 illustrated in Figure 1 and 3) may be an integrated sub-module of the management unit or a separate unit interfaced with the crosspoint switch (either directly or indirectly through the management module).

The client controller can also be implemented by an external computing system (i.e. remote of the printed circuit board and the enclosure 44) that is interface with the management network.
30 External client controllers ideally interface with the management module through the dedicated management port 52. The management module may transmit client information, such as recent client trading behavior and the financial condition of the client, to external client controllers to facilitate remote client monitoring.

A data replication system 80 is illustrated in Figure 5. The replication system can be used in various environments to disseminate information. In financial applications, the replication system can be used in conjunction with a financial ordering system. The illustrated data replication system comprises two networking devices 10, 80 (a primary networking device 10 and a secondary networking device 80). The networking devices 10, 80 are both interfaced with the broadcast port 82 of a financial order server 72. They are also interfaced with a plurality of client systems 70, 74.

The secondary networking device 80 illustrated in Figure 5 is internally configured to replicate physical layer signals received from the financial order server 72 for dissemination to the client computing systems 70, 74. A 'one-to-many' pass-through configuration is established within the crosspoint switch of the secondary networking device 80 to channel communications from the financial order server 72 ('one') to the client systems 72, 74 ('many'). The primary networking device 10 is configured for 'one-to-one' bidirectional communications between the financial order system 72 and the client computing systems 72, 74.

The replication system 80 illustrated in Figure 5 is particularly advantageous in latency sensitive applications (such as financial applications where accurate market information is critical).

A failsafe computing network 90 is illustrated in Figure 6. The computing network 90 comprises a plurality of servers 92, 94 that are interfaced to a common network 96. The network 96 interconnects the servers 92, 94 to a plurality of client machines 98,100. Typical networks include the internet, local area networks wide area networks.

Generally, the servers 92, 94 are deterministic computing systems that are maintained with the same (or very similar) internal states. An active server 92 services the client systems 98, 100 during regular network operation. The secondary server 94 functions as a 'live backup'. If the active server 92 fails, the backup server 94 is brought 'online' to service the client systems 98, 100 in place of the failed server 92.

The illustrated computing network 90 incorporates a low latency networking device 102 that functions in accordance with the disclosure of this specification. Each of the servers 92 and 94 are interfaced with the networking device 102 through dedicated communications ports 93, 95. A crosspoint switch (integrated with the networking device 102) establishes data conduits that interconnect the server communication ports 93, 95 to the network communication ports 97. The network port 97 is centrally disposed port in the illustrated embodiment, with the server ports 93 and 95 disposed on opposite side. This port allocation minimizes latency for communication

channeled from the network to the servers 92, 94 (such as stock orders placed by the client computing systems 98, 100).

A fault detector 104 monitors the active server 92 for faults. The illustrated fault detector 104 is interfaced with the active server 92 and the networking device 102. The fault monitor 104 is arranged to detect server faults and initialize failover to the backup server (typically by issuing a failover command to the networking device 102). The networking device 102 typically responds to failover commands by establishing a data conduit between the backup server port 97 with the network port 97. The networking device 102 may also break the data conduit between failed servers and the networking port 97 to isolate the failed server.

10 A latency determination system 110 is illustrated in Figure 7. The system 110 comprises a networking device 112 that operates in accordance with the disclosure of this specification. Communications ports are omitted from the system schematic depicted in Figure 7 to make the data conduits formed within the crosspoint switch 114 easily distinguishable.

The networking device 112 is disposed into a two-way physical link 128, 130 that interconnects two computing systems 132, 134. The crosspoint switch 114 illustrated in Figure 7 is configured to replicate data signals received from the computing systems 132, 134 (the incoming signals are received through communications ports 116, 122) and forward the replicated signal to a time stamp module.

20 Communications sent from the first computing system 132 are replicated by the networking device 112 using a 'one-to-many' routing configuration (represented by data conduit 136). The networking system 112 channels the replicated signals to a time stamp module 138 and the second computing device 134 respectively. The time stamp module generates a time stamp on receipt of the replicated signal.

25 Receipt of communications signal from the first computing system 132 triggers a response communication from the second computing system 134. The response communication transmitted by the second computing system 134 is similarly replicated and distributed by the networking device 112.

30 The response signal transmitted by the second computing system 134 is replicated using a 'one-to-many' routing configuration (represented by data conduit 136). The networking system channels the replicated signals to the time stamp module 138 and the first computing system 132 respectively. The time stamp module generates another time stamp on receipt of the replicated response signal. The two time stamps can then be differenced to determine the network latency.

Variations and/or modifications may be made to the embodiments described within this specification without departing from the spirit or ambit of the invention. For example, while the networks described above may comprise optical and/or electrical Ethernet (for example 10 Mb, 40Mb, 1Gb, 10Gb, 40Gb, 100Gb, 400Gb, 1Tb), it will be understood that other network types and protocols may be used, such as INFINIBAND and WiFi. Alternatively or additionally, one or more of the connections may alternatively be a serial port connection, a USB port connection, a FireWire (TM) port connection, a ThunderBolt (TM) port connection, a PCI or PCIe connection, a SONET (or SDH) connection with or without a sonet demultiplexing device, or generally any suitable type of connection. The present embodiments are, therefore, to be considered in all respects as illustrative and not restrictive.

Prior art, if any, described herein is not to be taken as an admission that the prior art forms part of the common general knowledge in any jurisdiction.

In the claims which follow and in the preceding description of the invention, except where the context requires otherwise due to express language or necessary implication, the word “comprise” or variations such as “comprises” or “comprising” is used in an inclusive sense, that is to specify the presence of the stated features but not to preclude the presence or addition of further features in various embodiments of the invention.

Claims

1. A networking apparatus comprising:

a plurality of communications ports that interface with external computing systems to channel physical layer signals,

5 a crosspoint switch connected to each of the communications ports that routes physical layer signals within the networking apparatus, the crosspoint switch having a plurality of discrete signal channels that are dynamically interconnectable to form data conduits,

a management module that controls routing of physical layer signals responsive to commands from an external computing system, the management module being interfaced with
10 the crosspoint switch to manipulate data conduits by making and/or breaking connections between signal channels, and

a processing unit connected in an internal routing loop with the crosspoint switch that decodes internally routed physical layer signals and processes extracted data prior to reversion to a destination communications port, the internal routing loop connecting the processing unit to a
15 plurality of crosspoint switch signal channels.

2. The networking apparatus of claim 1 comprising a mounting socket that releasably supports the processing unit, the mounting socket interfacing the processing unit with the internal routing loop.

3. The networking apparatus of claim 1 or claim 2 comprising a programmable logic device
20 integrated with the processing unit.

4. The networking apparatus of any one of claims 1 to 3 comprising a communications interface slot that receives a releasable socket module and interfaces individual port sockets within the socket module with signal channels of the crosspoint switch.

5. The networking apparatus of any one of claims 1 to 4 comprising a routing controller
25 integrated with the management module that configures the crosspoint switch for multiplexing operations involving a defined set of source communications ports to minimize multiplexing stages within the processing unit.

6. The networking apparatus of any one of claims 1 to 4 comprising a routing controller that determines a minimum set of logic blocks within the processing unit for multiplexing operations

and channels physical layer signals from designated communications ports to input interfaces at the processing unit that correspond to logic blocks within the minimum set to optimize logic block utilization.

7. The networking apparatus of any one of claims 1 to 6 comprising a switching module that facilitates configuration of individual communications ports via a management interface, the management module interconnecting signal channels within the crosspoint switch to link communications ports responsive to configuration commands received via the switching module.
8. The networking apparatus of claim 7 wherein the configurations commands supported by the switching module via the management interface include timestamping, link quality monitoring, ingress filtering, quality of service rate limiting, packet counting and port-to-port pass-through routing.
9. The networking apparatus of any one of claims 1 to 6 comprising a switching module that facilitates protocol classification for individual communication ports routed via the processing unit, the management module using the protocol classification for a communications port to initialize the processing unit to decode signals received via corresponding communications ports and/or encoding signals designated for corresponding communications ports.
10. The networking apparatus of any one of claims 1 to 9 comprising a plurality of physical layer conduits that interface the processing unit with individual signal channels of the crosspoint switch to form the internal routing loop, each of the physical layer conduits defining a dedicated unidirectional pathway for transmission of physical layer signals between the processing unit and a corresponding signal channel of the crosspoint switch.
11. The networking apparatus of any one of claims 1 to 8 comprising a plurality of physical layer conduits that interface the processing unit with individual signal channels of the crosspoint switch to form the internal routing loop, the internal routing loop having a forwarding conduit and a return conduit for each of the communications ports.
11. The networking apparatus of any one of claims 1 to 10 comprising a plurality of physical layer conduits that interface the communications ports with individual signal channels of the crosspoint switch, each of the communications ports having two physical layer conduits that define a loop between the communications port and the crosspoint switch.
12. The networking apparatus of any one of claims 1 to 8 comprising a plurality of physical layer conduits that interface the processing unit and the communications ports with dedicated signal

channels of the crosspoint switch, the management module interconnecting channels within the crosspoint switch to establish a transmission pathway between respective physical layer conduits.

- 5 13. The networking apparatus of any one of claims 1 to 12 comprising an independent management port interfaced with the management module, the management port facilitating reception of configuration commands from an external computing system via a management network.
- 10 14. The networking apparatus of any one of claims 1 to 13 comprising a signal monitoring unit that samples physical layer signals received via designated communications ports at defined offsets from a nominal sample point to evaluate signal quality, and determines whether the physical layer signal can be recovered at the sample points.
- 15 15. The networking apparatus of any one of claims 1 to 14 comprising a plurality of clock and data recovery modules that facilitate signal conditioning, each of the communications ports being associated with a clock and data recovery module.
- 15 16. The networking apparatus of claim 15 comprising a plurality of AC couplings that connect the crosspoint switch to each of the clock and data recovery modules.
17. The networking apparatus of claim 15 or claim 16 comprising a reference clock that is interfaced with each of the clock and data recovery modules.
18. A networking method comprising:
- 20 a networking apparatus interconnecting channels of a crosspoint switch to form reconfigurable data conduits responsive to receipt of configuration commands from an external computing system, the data conduits interconnecting a source communications port and a destination communications port in an internal routing loop that includes a processing unit,
- 25 the networking apparatus receiving physical layer signals at the source communications port and routing the received physical layer signals from the source communications port to the processing unit via a forwarding data conduit,
- 30 the networking apparatus causing the processing unit to decode physical layer signals, process data extracted from the physical layer signals and re-encode extracted data using a communications protocol that is compatible with a designated destination communications port, and

the networking apparatus routing the re-encoded physical layer signals from the processing unit to a designated destination communications port via a return data conduit.

19. The method of claim 18 comprising:

5 the networking apparatus routing physical layer signals from a plurality of source communications ports to the processing unit via distinct forwarding data conduits,

the processing unit extracting data from the physical layer signals and multiplexing extracted data received via each of the distinct forwarding data conduits into a multiplexed signal stream, and

10 the networking apparatus routing the multiplexed signal stream to a destination communications port via a return data conduit.

20. The method of claim 18 or claim 19 comprising the networking apparatus causing the processing unit to decode physical layer signals received via a plurality of source communications ports, extract data frames from the decoded signals, interleave frames received via each of the communications ports in a multiplexed data stream and encode the multiplexed data stream using a communications protocol that is compatible with a destination communications port.

21. The method of claim 19 or claim 20 comprising:

20 the networking apparatus determining a routing configuration for the crosspoint switch that minimizes multiplexer stages within the processing unit for a defined set of source communications ports, and

the networking apparatus interconnecting distinct channels within the crosspoint switch to form forwarding data conduits that conform to the determined routing configuration.

22. The method of claim 19 or claim 20 comprising:

25 the networking apparatus determining a minimum set of logic blocks within the processing unit for multiplexing operations, and

the networking apparatus interconnecting distinct channels within the crosspoint switch to channel physical layer signals from designated communications ports to input interfaces at the processing unit that correspond to logic blocks within the minimum set to optimize logic block utilization.

23. The method of any one of claims 19 to 21 comprising the networking apparatus causing the processing unit to timestamp the data extracted from each of the physical layer signals prior to multiplexing the extracted data.

24. The method of any one of claims 18 to 23 comprising:

5 the networking apparatus managing protocol classification for communications ports routed via the processing unit, including receiving protocol classifications for individual communications ports, and

 the networking apparatus using the received protocol classifications to initialize the processing unit for compatibility with a corresponding communications ports to facilitate
10 decoding received signals and/or encoding data for transmission.

25. The method of any one of claims 18 to 24 comprising the networking apparatus facilitating configuration of individual communications ports, including receiving configuration commands from an external computing system, and reconfiguring the crosspoint switch responsive to the received configuration commands .

15 26. The method of claim 25 comprising the networking apparatus breaking data conduits within the crosspoint switch responsive to received configuration commands to prevent routing of physical layer signals from a source communications port.

27. The method of claim 25 comprising the networking apparatus autonomously routing physical layer signals received via a source communications port to the processing unit responsive to
20 receiving a configuration command that specifies multiplexing, timestamping, link quality monitoring, ingress filtering, quality of service rate limiting or packet counting operations.

28. The method of any one of claims 18 to 27 comprising the networking apparatus receiving a configuration command from an external computing system and establishing a routing loop between communications ports of the networking apparatus responsive to the received
25 configuration command.

29. The method of anyone of claims 18 to 28 comprising the networking apparatus conditioning physical layer signals before routing the signals to a destination communications port.

30. The method of claim 29 comprising the networking apparatus performing a clock recovery and signal regeneration process to condition the physical layer signals.

31. The method of claim 29 or claim 30 comprising the networking apparatus shaping degraded physical layer signals.

32. The method of any one of claims 18 to 31 comprising the networking apparatus replicating received physical layer signals and routing the replicated signals to a plurality of destination communications ports.

33. A networking apparatus comprising a crosspoint switch that channels physical layer signals between communications ports and a processing unit connected in an internal routing loop with the crosspoint switch that receives physical layer signals from the crosspoint switch intermediate source and destination communications ports.

10 34. A networking switch comprising a routing matrix that channels communications signals without decoding protocol data encoded with the communications signals and a processing unit connected to the routing matrix.

35. The networking switch of claim 35 comprising an internal routing loop that connects the processing unit

15




Figure 1





Figure 2

Figure 5

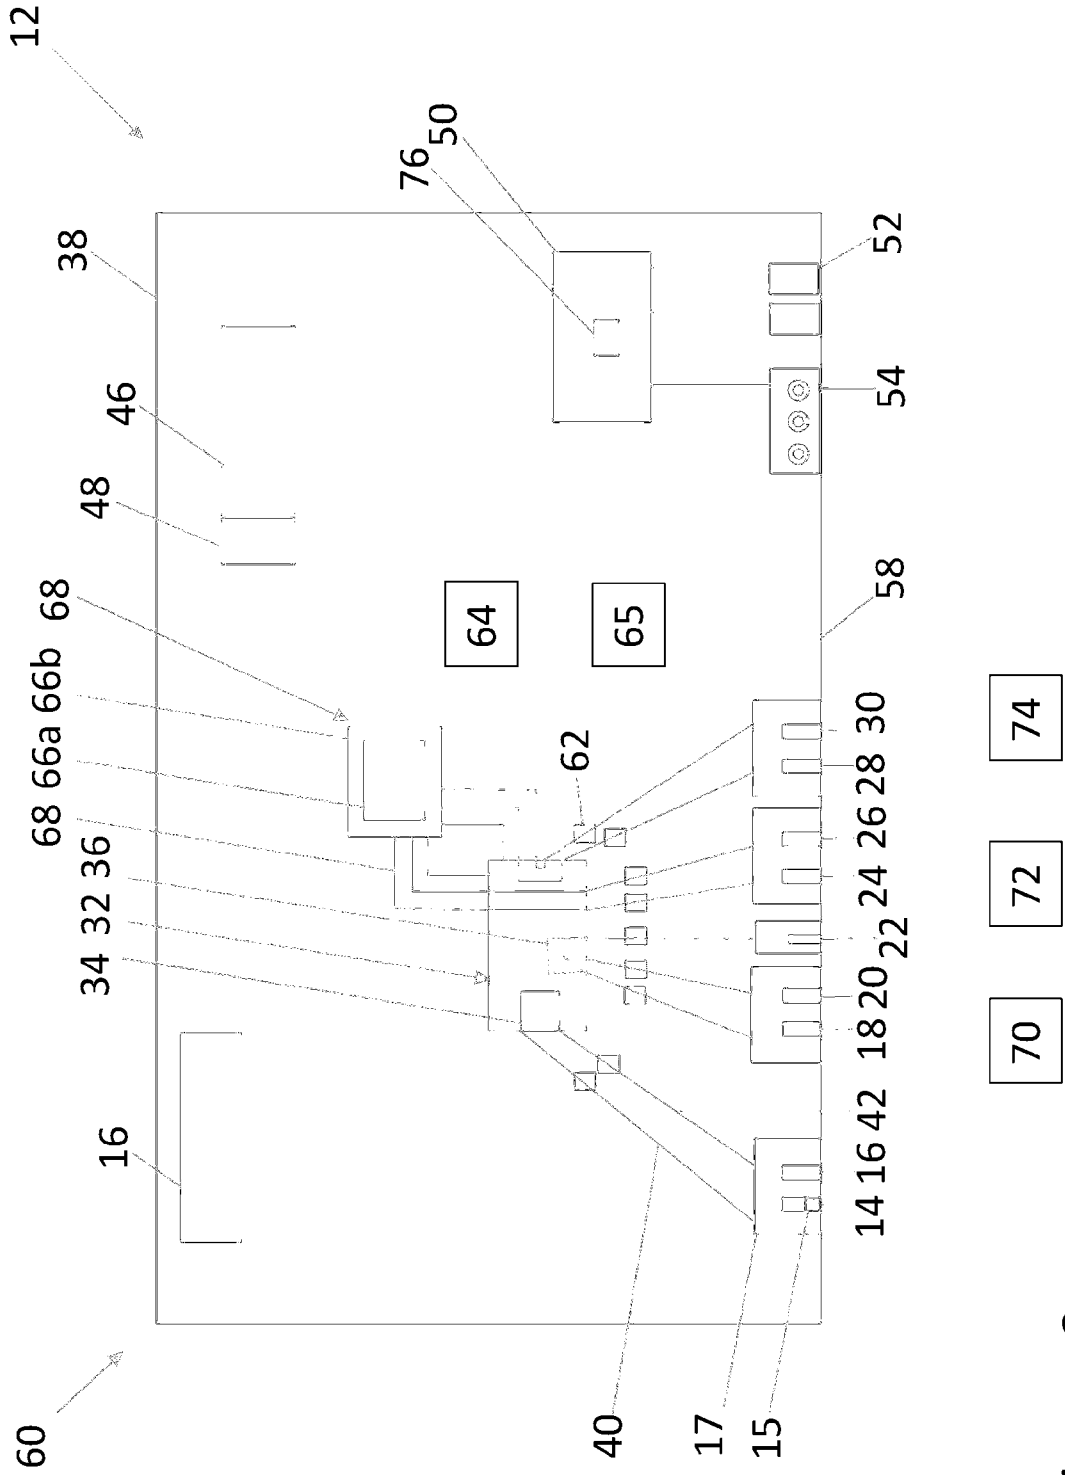
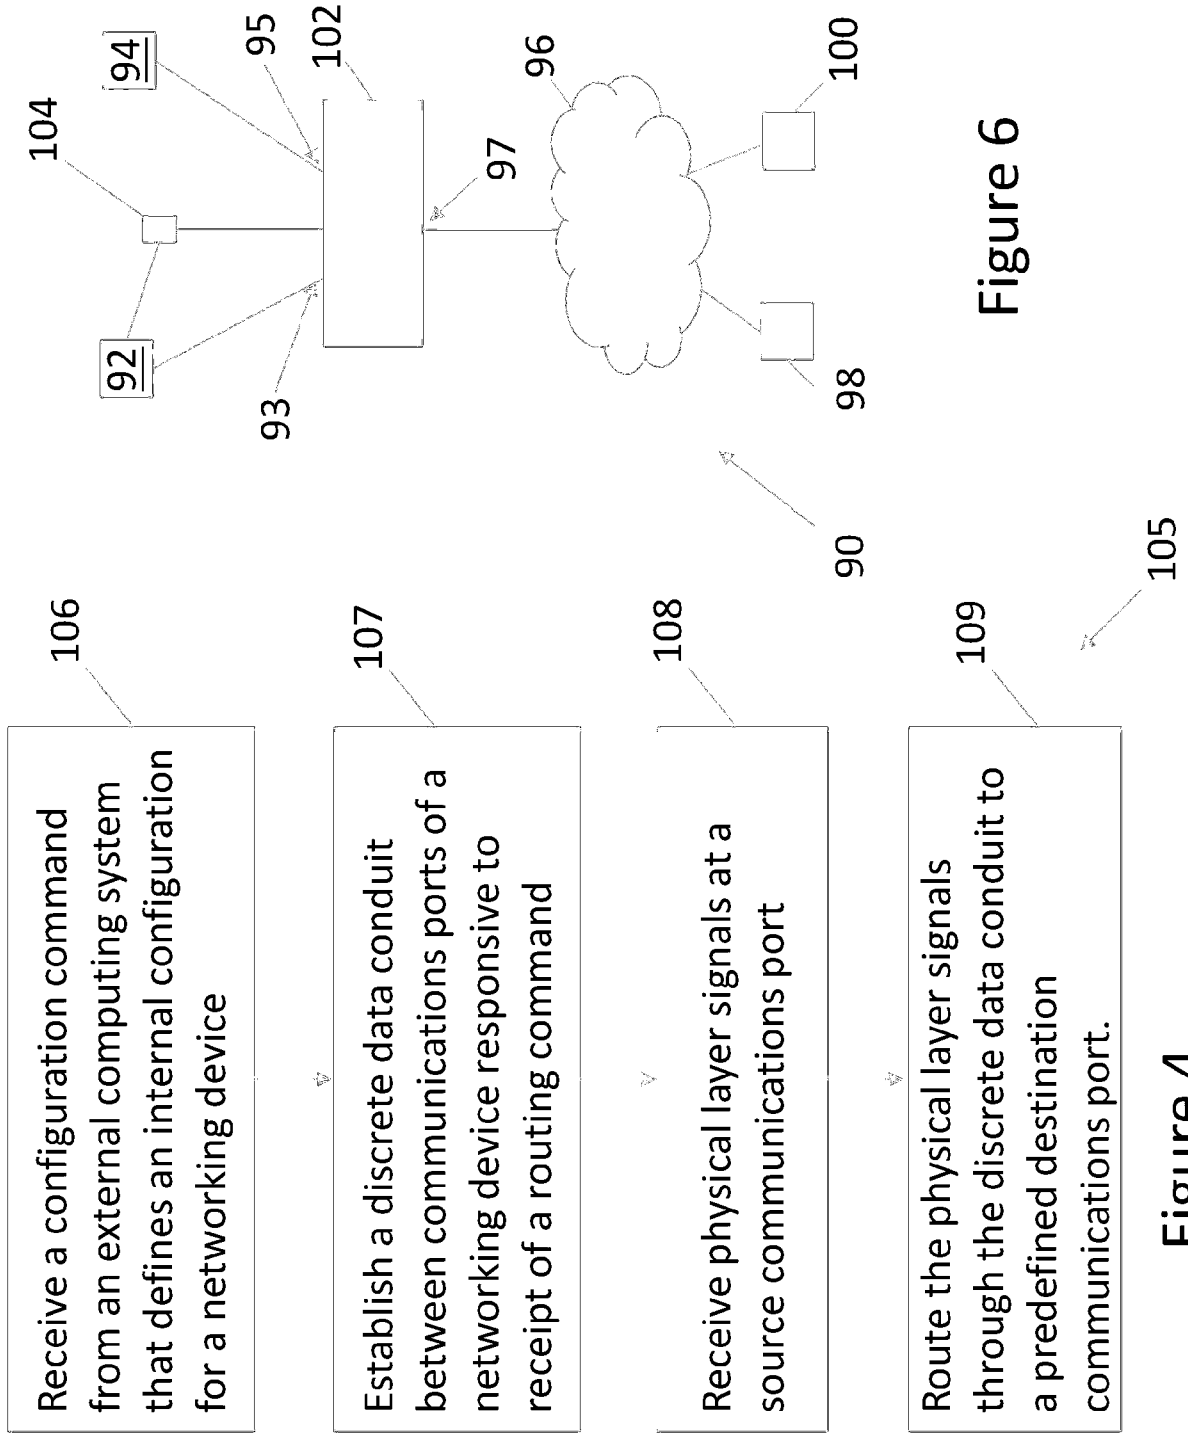


Figure 3






Figure 7




Figure 8b




Figure 8a

INTERNATIONAL SEARCH REPORT

International application No.

PCT/AU2015/000311

A. CLASSIFICATION OF SUBJECT MATTER H04L 12/931 (2013.01) H04L 12/00 (2006.01) G06F 13/38 (2006.01)		
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)		
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched		
Electronic data base consulted during the international search (name of data base and, where practicable, search terms used) Databases: EPODOC, WPIAP, INSPEC, Google, Google Patents, TXPEA, TXPEB, TXPEC, TXPEE, TXPEF, TXPEH, TXPEI, TXPEP, TXPES, TXPEPEA, TXPUSE0A, TXPUSE1A, TXPUSEA, TXPUSEB, TXPW0EA with keywords - (physical layer signals, dynamic routing, data conduits, reconfigurable, external routing commands, transceiver, low latency signal transmission, AC couplings, crosspoint switch, switch fabric, service ports, physical layer routing, transputer, C004 chips, race++, data switching matrix, finance, high frequency trading, port-to-port pass through routing, network protocols, transmission protocols, internal routing loop, clock recovery, signal regeneration, SFP/GBIC/XFP/XAUI transceivers, VITESSE, VSC3172) and like terms Google Prior Art Finder with WO 2013/155566 A1 and WO 2013/134810 A1 Applicant(s) and Inventor(s) search performed in AusPat, Espacenet and Internal Databases provided by IP Australia		
C. DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
	Documents are listed in the continuation of Box C	
<input checked="" type="checkbox"/> Further documents are listed in the continuation of Box C <input checked="" type="checkbox"/> See patent family annex		
* "A"	Special categories of cited documents: document defining the general state of the art which is not considered to be of particular relevance	"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
"E"	earlier application or patent but published on or after the international filing date	"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
"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)	"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
"O"	document referring to an oral disclosure, use, exhibition or other means	"&" document member of the same patent family
"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 28 August 2015	Date of mailing of the international search report 28 August 2015	
Name and mailing address of the ISA/AU AUSTRALIAN PATENT OFFICE PO BOX 200, WODEN ACT 2606, AUSTRALIA Email address: pct@ipaustalia.gov.au	Authorised officer Kusha Panta AUSTRALIAN PATENT OFFICE (ISO 9001 Quality Certified Service) Telephone No. 0262832833	

INTERNATIONAL SEARCH REPORT		International application No.
C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		PCT/AU2015/000311
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	WO 2013/155566 A1 (ZOMOJO PTY LTD) 24 October 2013 Whole document, in particular, abstract, page 1 line 23 - page 4 line 7; page 5 line 14 - page 7 line 27; page 8 line 8 - page 9 line 23; page 10 line 10 - page 14 line 20; page 15 line 1 - page 16 line 18; FIGs. 1-7	1 - 35
X	WO 2013/134810 A1 (ZEPTOIP PTY LTD et al.) 19 September 2013 Whole document, in particular, abstract; page 1 line 21 – page 6 line 29; page 8 line 11 - page 10 line 10; page 12 line 21 - page 13 line 22, page 14 line 15 - page 16 line 19; FIGs. 1-4	1 - 35
X	US 2013/0121692 A1 (PATEL) 16 May 2013 Whole document, in particular, abstract; paragraphs [0006]-[0012], [0019], [0025]-[0026], [0043]-[0047] and [0075]; FIGs. 1-3	33 - 35

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/AU2015/000311

This Annex lists known patent family members relating to the patent documents cited in the above-mentioned international search report. The Australian Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

Patent Document/s Cited in Search Report		Patent Family Member/s	
Publication Number	Publication Date	Publication Number	Publication Date
WO 2013/155566 A1	24 October 2013	WO 2013155566 A1	24 Oct 2013
		AU 2013248939 A1	23 Jan 2014
		AU 2013248939 B2	07 May 2015
		CA 2870483 A1	24 Oct 2013
		EP 2839621 A1	25 Feb 2015
		JP 2015518687 A	02 Jul 2015
		KR 20150003261 A	08 Jan 2015
		US 2015046613 A1	12 Feb 2015
		US 2015046613 A1	12 Feb 2015
WO 2013/134810 A1	19 September 2013	WO 2013134810 A1	19 Sep 2013
		AU 2012373188 A1	04 Sep 2014
		US 2015043593 A1	12 Feb 2015
US 2013/0121692 A1	16 May 2013	US 2013121692 A1	16 May 2013
		US 8891963 B2	18 Nov 2014
		WO 2013033832 A1	14 Mar 2013

End of Annex

Due to data integration issues this family listing may not include 10 digit Australian applications filed since May 2001.

Form PCT/ISA/210 (Family Annex)(July 2009)