

Exhibit E30 – Invalidity of U.S. Patent No. 9,450,884 in view of Peterson – US Patent No. 8,274,887

INVALIDITY CLAIM CHART

Invalidity of U.S. Patent No. 9,450,884 in view of Peterson – US Patent No. 8,274,887

US Patent No. 8,274,887 (“Peterson”) was published on September 25, 2012, from an application filed on May 14, 2010. Peterson therefore constitutes prior art under at least 35 U.S.C. §§ 102(a)(1) and 102(a)(2). Peterson anticipates and/or renders obvious the Asserted Claims of the U.S. Patent No. 9,450,884, either alone or in combination with one or more of the references identified in Defendant’s Cover Pleading (at least as Defendant understands WSOU’s application of the Asserted Claims in an effort to show infringement).

To the extent WSOU argues that any element below is not disclosed by Peterson, a person of ordinary skill in the art would have found it obvious to combine the teachings of Peterson with the background knowledge of a person of ordinary skill in the art and/or the additional references, and exemplary teachings, set forth in Appendix E.

The chart below provides representative examples of where each element of each claim is found within Peterson. Citations are meant to be exemplary, not exhaustive, and Defendant reserves the right to identify and discuss additional portions of the reference in support of its contentions and/or to rebut arguments made by WSOU. Citations to figures, drawings, tables, and the like include reference to any accompanying or related text. All internal cross references are meant to incorporate the cross-referenced material as if fully set forth therein.

It is Defendant’s position that WSOU’s Preliminary Infringement Contentions have not established that any accused product or service infringes any valid claim. Thus, Defendant’s statements below should not be treated as an admission, implication, or suggestion that Defendant agrees with WSOU regarding either the scope, construction, or interpretation of any of the Asserted Claims or the infringement theories advanced by WSOU in its Preliminary Infringement Contentions, including whether any Asserted Claim satisfies 35 U.S.C. §§ 101 or 112. In certain cases, Defendant specifies non-limiting examples of where its application of the prior art is based on WSOU’s apparent application of the claim element. These statements are not intended to suggest that Defendant agrees with WSOU’s application of any claim term, suggest a proposed construction at this stage of the case, or suggest that construction is needed, as the parties are not required to exchange terms for construction or proposed constructions until a later date.

WSOU has yet to identify any limitation of the Asserted Claims that it contends is not anticipated and/or rendered obvious by Peterson. Defendant therefore expressly reserves the right to respond to any such contention, including by identifying additional obviousness combinations, if WSOU makes any such contention.

Where Defendant states that Peterson “discloses” a limitation, that disclosure may be express, implicit, and/or inherent.

Exhibit E30 – Invalidity of U.S. Patent No. 9,450,884 in view of Peterson – US Patent No. 8,274,887

Claim	Disclosure
[1pre] A method of adjusting bandwidth allocation by a network switching element in a communications network, the network switching element including a target port, the method comprising:	To the extent the preamble is limiting, Peterson discloses this preamble. See Peterson at Abstract: “According to an example embodiment, a total offered traffic load for a shared resource within a network switching system may be determined, the total offered traffic load may include, for example, a sum of offered traffic loads from one or more active virtual output queues (VOQs) of the network switching system. A capacity of the shared resource within the network switching system may be determined. A transmission rate from one or more of the active VOQs over the shared resource may be adjusted such that the total traffic from the active VOQs does not exceed the capacity of the shared resource.” See Peterson at 1:19-27: “A variety of network switching systems have been developed, such as routers, switches and the like. A switching system may commonly include multiple devices, such as switches, line cards or other devices. Each device may include one or more ports. Depending on the direction of traffic, a port may act as an input (or ingress) port or an output (or egress) port. One function of a network switching system is to receive packets (messages) on an ingress port, and forward or route the packet via a selected egress port.” See Peterson at 1:56-2:35: “According to an example embodiment, a network switching system may be provided. The network switching system may include a plurality of ingress devices. One or more of the ingress devices may include a plurality of ports and a plurality of traffic queues (such as VOQs) to store traffic to be transmitted via one or more shared resources to an egress device of the network switching system. One or more of the ingress devices may be configured to exchange (e.g., transmit and receive) traffic queue state information with one or more other ingress devices within the network switching system. One or more of the ingress devices may also be configured to adjust a transmitted traffic load from one or more of the traffic queues (or VOQs) based on the queue state information received from other ingress devices. According to another example embodiment, a network switching system may include a plurality of network devices, each network device including a plurality of virtual output queues (VOQs) for storing traffic to be forwarded to a destination. The network switching system may include at least one intermediary device coupled between the network devices. One or more of the ingress devices being configured to receive queue state messages indicating a state of the VOQs of one or more of

Exhibit E30 – Invalidity of U.S. Patent No. 9,450,884 in view of Peterson – US Patent No. 8,274,887

Claim	Disclosure
	the other network devices of the network switching system. One or more of the network devices also being configured to adjust a transmission rate from the VOQs based on the received queue state information received from one or more other ingress devices. According to an example embodiment, a total offered traffic load for a shared resource within a network switching system may be determined, the total offered traffic load may include, for example, a sum of offered traffic loads from one or more active virtual output queues (VOQs) of the network switching system. A capacity of the shared resource within the network switching system may be determined. A transmission rate from one or more of the active VOQs over the shared resource may be adjusted such that the total traffic load from the active VOQs does not exceed the capacity of the shared resource. A network device (e.g., switch) within the network switching system may receive queue status messages from the other network devices, indicating a queue state (e.g., empty/non-empty state, a level of fullness) for one or more VOQs. In this manner, each network device in the network switching system may, for example, determine a global or overall offered traffic load from multiple VOQs or multiple network devices, and each network device may adjust a transmission rate from one or more its VOQs, e.g., to avoid exceeding a capacity of the shared resource.” <i>See also, e.g., Peterson at Fig. 10.</i>
[1a] monitoring, by the network switching element, a data flow traversing the target port of the network switching element;	Peterson discloses monitoring, by the network switching element, a data flow traversing the target port of the network switching element. <i>See Peterson at Abstract:</i> “According to an example embodiment, a total offered traffic load for a shared resource within a network switching system may be determined, the total offered traffic load may include, for example, a sum of offered traffic loads form one or more active virtual output queues (VOQs) of the network switching system. A capacity of the shared resource within the network switching system may be determined. A transmission rate from one or more of the active VOQs over the shared resource may be adjusted such that the total traffic from the active VOQs does not exceed the capacity of the shared resource.” <i>See Peterson at 1:19-27:</i> “A variety of network switching systems have been developed, such as routers, switches and the like. A switching system may commonly include multiple devices, such as switches, line cards or other devices. Each device may include one or more ports. Depending on the direction of traffic, a port may act as an input (or ingress) port or an output (or egress) port. One function of a network

Exhibit E30 – Invalidity of U.S. Patent No. 9,450,884 in view of Peterson – US Patent No. 8,274,887

Claim	Disclosure
	switching system is to receive packets (messages) on an ingress port, and forward or route the packet via a selected egress port.” See Peterson at 1:56-2:35: “According to an example embodiment, a network switching system may be provided. The network switching system may include a plurality of ingress devices. One or more of the ingress devices may include a plurality of ports and a plurality of traffic queues (such as VOQs) to store traffic to be transmitted via one or more shared resources to an egress device of the network switching system. One or more of the ingress devices may be configured to exchange (e.g., transmit and receive) traffic queue state information with one or more other ingress devices within the network switching system. One or more of the ingress devices may also be configured to adjust a transmitted traffic load from one or more of the traffic queues (or VOQs) based on the queue state information received from other ingress devices. According to another example embodiment, a network switching system may include a plurality of network devices, each network device including a plurality of virtual output queues (VOQs) for storing traffic to be forwarded to a destination. The network switching system may include at least one intermediary device coupled between the network devices. One or more of the ingress devices being configured to receive queue state messages indicating a state of the VOQs of one or more of the other network devices of the network switching system. One or more of the network devices also being configured to adjust a transmission rate from the VOQs based on the received queue state information received from one or more other ingress devices. According to an example embodiment, a total offered traffic load for a shared resource within a network switching system may be determined, the total offered traffic load may include, for example, a sum of offered traffic loads from one or more active virtual output queues (VOQs) of the network switching system. A capacity of the shared resource within the network switching system may be determined. A transmission rate from one or more of the active VOQs over the shared resource may be adjusted such that the total traffic load from the active VOQs does not exceed the capacity of the shared resource. A network device (e.g., switch) within the network switching system may receive queue status messages from the other network devices, indicating a queue state (e.g., empty/non-empty state, a level of fullness) for one or more VOQs. In this manner, each network device in the network switching system may, for example, determine a global or overall offered traffic load from multiple VOQs or multiple network devices, and each network device may adjust a transmission rate from one or more its VOQs, e.g., to avoid exceeding a capacity of the shared resource.”

Exhibit E30 – Invalidity of U.S. Patent No. 9,450,884 in view of Peterson – US Patent No. 8,274,887

Claim	Disclosure
	See Peterson at 3:24-42: “According to an example embodiment, each ingress device (or network device) may determine an overall or global traffic load offered for one or more destinations, path segments, or resources in a network switching system, e.g., based on received queue state messages received from the other ingress devices in the network switching system. The queue state messages may indicate, for example, a queue state (e.g., an empty/non-empty state or a level of fullness of the queue. A resource (e.g., path segment, intermediary device, or switch fabric device) may be shared among multiple active (e.g., non-empty) queues. The queues may be, for example, virtual output queues (VOQs). Based on a total offered traffic load for a shared resource, e.g., including traffic from multiple VOQs, one or more network devices in the network switching system may adjust a transmission rate from one or more of its active queues or VOQs over the shared resource such that the total traffic load from a plurality of queues does not exceed the bandwidth or traffic capacity (e.g., bit rate capacity) of the shared resource.” See Peterson at 3:43-51: “Thus, for example, by exchanging queue state information between network devices in a network switching system, congestion in the switching system may be avoided in some cases (e.g., anticipated and transmission rates adjusted to avoid congestion). By exchanging queue state information between ingress devices, an overall or global offered traffic load for a resource (e.g., path segment, fabric device or other resource) may be determined by each network device in the network switching system.” See Peterson at 4:1-20: “FIG. 1 is a diagram of a network switching system according to an example embodiment. Network switching system 100 may include one or more network devices, such as network devices 102, 104, etc. Network devices 102, 104 may be network switches, routers, line cards, or other network devices. Each network device may be a bidirectional device, and may be considered an ingress (or input) device and/or an egress (or output) device, depending on the direction of travel of a particular packet, for example. Each network device may include one or more network ports, such as network ports 110 for network device 102 and network ports 112 for network device 104. Network devices 102 and 104 of network switching system 100 may communicate with external devices (not shown) using network ports 110 and 112 via a network interface 103. Network interface 103 may be, for example, an Ethernet interface, or other network interface. Network devices 102 and 104 may

Exhibit E30 – Invalidity of U.S. Patent No. 9,450,884 in view of Peterson – US Patent No. 8,274,887

Claim	Disclosure
	communicate with each other within network switching system 100 via an internal switching system interface that is referred to herein as a fabric interface 114.” <i>See Peterson at 4:44-63:</i> “Overall, network switching system 100 may receive and process a received packet (e.g., Ethernet packet) via an ingress port of an ingress device (e.g., the network device that receives the packet), and may forward the packet to another (e.g., egress) port of the same or different network device. The egress device of the network switching system may forward the packet on to another device (not shown) outside the network switching system, via a network interface. The port which received the incoming packet may be referred to as an ingress port on the ingress device for this particular packet, and the output port for this particular packet may be referred to as the egress port on the egress device. The packet may, for example, be received via an ingress port (of an ingress network device), and may be transmitted to an external device (not shown) via an egress port (of an egress network device). Each port (and also each network device) may be bidirectional, and thus, the use of the term ‘ingress’ or ‘egress’ may, for example, be with respect to a particular packet or a particular traffic flow direction, according to an example embodiment.” <i>See also, e.g., id. at 4:64-5:64.</i> <i>See Peterson at 5:65-6:11:</i> “FIG. 5 is a diagram illustrating a network device according to an example embodiment. Network device 500 may be any type of network device, such as a switch, router, line card, or other device. Network device 500 may include a network interface circuit 508 providing a network (e.g., Ethernet) interface 103 via one or more network ports. A fabric interface circuit 510 may provide a fabric interface 114 via one or more fabric interface ports. Thus, according to an example embodiment, network device 500 may include both network (e.g., Ethernet) ports and fabric interface ports. Network device 500 may include a controller 502 to provide overall control of network device 500, a memory 504 to store data, instructions or other information, and a path and bandwidth database 506.” <i>See Peterson at 6:25-41:</i> “Network device 500 may also include a plurality of queues to store packets and/or cells to be transmitted or forwarded across a network switching system. According to an example embodiment, virtual output queues (VOQs) may be used, such as VOQ1, VOQ2, VOQ3, ... VOQN for N VOQs. There may be, for example, one virtual output queue (VOQ) at an ingress network device storing cells or packets to be transmitted to each destination. Thus, within a network device

Exhibit E30 – Invalidity of U.S. Patent No. 9,450,884 in view of Peterson – US Patent No. 8,274,887

Claim	Disclosure
	500 (as an ingress device), there may be an association between a virtual output queue (VOQ) and a destination, in an example embodiment. In an example embodiment, there may be one VOQ in an ingress device per destination, for example. A destination may include, for example, within a network switching system, an egress device, a port or physical port (e.g., egress port) of an egress device, a virtual port of an egress device, etc. The use of VOQs may, for example, avoid, or at least decrease, head of line blocking.” <i>See Peterson at 9:22-36:</i> “A second fabric header set 604 may further include a source field (e.g., identifying a source such as an ingress device or ingress port) and a length field indicating a length of the fabric packet). A third fabric header set 606 illustrates an example header set that may be assembled when two network switches or two network processors are exchanging packets within the network switching system, e.g., across one or more intermediary devices. The header set 606 may include, for example, a flow field identifying a flow (or group of associated fabric packets), a format field to identify a format for the fabric packet, and one or more processing parameters. The flow, format and processing parameter fields may, for example, be used by an egress device within the network switching system to perform the required processing during egress handling of the packet.” <i>See Peterson at 12:10-21:</i> “A fabric packet may include a number of different headers as part of a header set. These headers may include, for example, a destination header, a source aggregate header, a length header, a format header, a flow header, a proxy parameters header, processing parameters, and a bandwidth management header (e.g., to provide or generate a queue path message, path state message and path rate message, as examples). These are merely some examples of the types of headers that may be included in a fabric packet or cell. In some cases, one or more of these headers may be optional, and may be provided in some different combinations. Some of these headers will be briefly described.” <i>See Peterson at 14:3-12:</i> “A flow header may also be used to identify a flow for the current packet. A flow may be, for example, a group of related packets, or one or more packets from a same source or application, and/or directed to a same destination. For example, a flow may be a source, destination and priority combination that defines a connection or aggregate of connections and/or that may originate from and/or are destined to a same or collocated network locations. Other types of flows may be used as

Exhibit E30 – Invalidity of U.S. Patent No. 9,450,884 in view of Peterson – US Patent No. 8,274,887

Claim	Disclosure
	well. The flow header may include a next header field to identify a next header after the current fabric header.” See Peterson at 15:12-39: “The queue state message 810 may be used by ingress devices to convey or indicate the state of one or more VOQs (virtual output queues) at the ingress device. The queue state message 810 may be, for example, a 1-bit value indicating an empty or non-empty state of each VOQ, or a multi-bit value indicating a level of fullness of each VOQ. Thus, the queue state message 810 may include a field indicating a state of each VOQ (or for one or more VOQs) of the reporting ingress device, for example. Other information for each VOQ may also be transmitted in a queue state message 810, such as a requested data transmission rate for the VOQ, a traffic priority for queued traffic or the VOQ, and other information. The queue state message 810 may include a queue state header, that may include node ID field identifying the ingress node that is reporting, an indication of the number of VOQs being reported (or an identification of which VOQs are being reported), and a queue state for each VOQ being reported, for example. This is merely an example format, and other formats may be used. The queue state message 810 may be transmitted (e.g., to other devices in the network switching system) when there is a change in queue state (e.g., empty to non-empty, or increases in fullness), or may be sent out at specific time periods or intervals regardless whether there has been a change in queue state. Or, both techniques may be used to send out queue state messages (e.g., both when a change in a queue state occurs and sending out a queue state message at specific times or intervals regardless of a change).” See Peterson at 16:36-60: “Congestion in a network switching system may occur when an offered load of network traffic exceeds the traffic capacity or bit rate of one or more resources of the system. According to an example embodiment, each of one or more ingress devices in a network switching system may have one or more traffic queues (e.g., VOQs). According to an example embodiment, each ingress device may have a plurality of VOQs, with each VOQ being associated with a destination (e.g., egress device, egress port, or virtual port on an egress device). According to an example embodiment, the ingress devices may exchange queue state information with the other ingress devices of the network switching system. The queue state information may be exchanged, for example, by each ingress device transmitting queue state messages to the other devices indicating a state of its VOQs (e.g., empty/non-empty, or indicating a level of fullness of each VOQ, or other VOQ information). For

Exhibit E30 – Invalidity of U.S. Patent No. 9,450,884 in view of Peterson – US Patent No. 8,274,887

Claim	Disclosure
	example, the queue state messages may allow one or more of the network devices to obtain a global or overall view of the total offered traffic load over one or more shared resources in the network switching system, for example. Each network device, based on the queue state messages, may determine a total offered traffic load for a shared resource within the network switching system, where the total offered traffic load may include a sum of offered traffic loads from each of one or more VOQs.” <i>See Peterson at 18:54-62:</i> “Each VOQ (VOQs 902 and 904) would like to transmit 30 Gbps. However, ingress device 1 transmits a queue state message (via an intermediary device to ingress device 2) indicating a queue state of VOQ 902 of non-empty, or full. Ingress device 2, likewise, may transmit a queue state message to ingress device 1 indicating that VOQ 904 has a queue state of non-empty or full. Thus, both ingress devices 1 and 2 may determine that there are two active VOQs which share the same resources, including path segment 907.” <i>See Peterson at 19:20-28:</i> “With reference to FIG. 9B, ingress devices 1 and 2 may exchange queue state messages so that each ingress device may obtain or determine a total offered traffic load for segment 922, which is a resource shared by both VOQs 902 and 904. Both VOQs 902 and 904 would like to transmit at a data rate of 30 Gbps each, for example (or at other rates that collectively may exceed the capacity of segment 922). However, segment 922 has a traffic capacity of 20 Gbps, which is shared by both VOQs 902 and 904.” <i>See Peterson at 20:19-28:</i> “According to an example embodiment, in-band control messages, such as queue state messages 810, path state messages 820, and path rate messages 830 may be passed between the ingress devices 1 and 2 (FIGS. 9A and 9B) to convey the state of the VOQs 902 and 904. For example, an empty state, a full state, an intermediate state, or a more precise representation of the capacity of the VOQ and how many packets, cells, requested traffic rate, traffic priority (or traffic class) for each queue, or other data associated with the VOQ may be shared between the ingress devices 1 and 2.” <i>See Peterson at 22:61-23:8:</i> “For example, operation 1010 may include each of a plurality of ingress devices of the network switching system reporting a state of one or more of its VOQs such that one or more ingress devices of the network switching system obtains a global view of a demand (or offered traffic load)

Exhibit E30 – Invalidity of U.S. Patent No. 9,450,884 in view of Peterson – US Patent No. 8,274,887

Claim	Disclosure
	upon one or more shared resources of the network switching system. For example, queue state messages (e.g., indicating empty/non-empty state or a level of fullness) may be exchanged between network devices of the network switching system to allow one or more network devices of the system to determine a state (e.g., active or non-empty) of one or more VOQs. Each of the VOQs may be associated with a destination, and a path from an ingress device to an egress device may include one or more resources that may be shared among VOQs of the network switching system.”
[1b] determining, by the network switching element, a bandwidth allocation for the target port, the bandwidth allocation for the target port being a bandwidth that is currently allocated for the data flow;	Peterson discloses determining, by the network switching element, a bandwidth allocation for the target port, the bandwidth allocation for the target port being a bandwidth that is currently allocated for the data flow. <i>See Peterson at Abstract:</i> “According to an example embodiment, a total offered traffic load for a shared resource within a network switching system may be determined, the total offered traffic load may include, for example, a sum of offered traffic loads from one or more active virtual output queues (VOQs) of the network switching system. A capacity of the shared resource within the network switching system may be determined. A transmission rate from one or more of the active VOQs over the shared resource may be adjusted such that the total traffic from the active VOQs does not exceed the capacity of the shared resource.” <i>See Peterson at 2:16-35:</i> “According to an example embodiment, a total offered traffic load for a shared resource within a network switching system may be determined, the total offered traffic load may include, for example, a sum of offered traffic loads from one or more active virtual output queues (VOQs) of the network switching system. A capacity of the shared resource within the network switching system may be determined. A transmission rate from one or more of the active VOQs over the shared resource may be adjusted such that the total traffic load from the active VOQs does not exceed the capacity of the shared resource. A network device (e.g., switch) within the network switching system may receive queue status messages from the other network devices, indicating a queue state (e.g., empty/non-empty state, a level of fullness) for one or more VOQs. In this manner, each network device in the network switching system may, for example, determine a global or overall offered traffic load from multiple VOQs or multiple network devices, and each network device may adjust a transmission rate from one or more its VOQs, e.g., to avoid exceeding a capacity of the shared resource.”

Exhibit E30 – Invalidity of U.S. Patent No. 9,450,884 in view of Peterson – US Patent No. 8,274,887

Claim	Disclosure
	<i>See Peterson at 3:24-42:</i> “According to an example embodiment, each ingress device (or network device) may determine an overall or global traffic load offered for one or more destinations, path segments, or resources in a network switching system, e.g., based on received queue state messages received from the other ingress devices in the network switching system. The queue state messages may indicate, for example, a queue state (e.g., an empty/non-empty state or a level of fullness of the queue. A resource (e.g., path segment, intermediary device, or switch fabric device) may be shared among multiple active (e.g., non-empty) queues. The queues may be, for example, virtual output queues (VOQs). Based on a total offered traffic load for a shared resource, e.g., including traffic from multiple VOQs, one or more network devices in the network switching system may adjust a transmission rate from one or more of its active queues or VOQs over the shared resource such that the total traffic load from a plurality of queues does not exceed the bandwidth or traffic capacity (e.g., bit rate capacity) of the shared resource.” <i>See Peterson at 3:43-59:</i> “Thus, for example, by exchanging queue state information between network devices in a network switching system, congestion in the switching system may be avoided in some cases (e.g., anticipated and transmission rates adjusted to avoid congestion). By exchanging queue state information between ingress devices, an overall or global offered traffic load for a resource (e.g., path segment, fabric device or other resource) may be determined by each network device in the network switching system. Based on this global demand of a shared resource, each network device may, for example, adjust a transmission rate for one or more of its active VOQs that are using the resource to use up to, for example, approximately a fair portion (or other amount) of the bandwidth or traffic (e.g., bit rate) capacity, such as approximately 1/Nth of the bandwidth or traffic capacity for the shared resource for N VOQs sharing the resource, according to an example embodiment.” <i>See Peterson at 5:65-6:11:</i> “FIG. 5 is a diagram illustrating a network device according to an example embodiment. Network device 500 may be any type of network device, such as a switch, router, line card, or other device. Network device 500 may include a network interface circuit 508 providing a network (e.g., Ethernet) interface 103 via one or more network ports. A fabric interface circuit 510 may provide a fabric interface 114 via one or more fabric interface ports. Thus, according to an example embodiment, network device 500 may include both network (e.g., Ethernet) ports and fabric interface ports. Network device 500 may include a controller 502 to provide overall control of

Exhibit E30 – Invalidity of U.S. Patent No. 9,450,884 in view of Peterson – US Patent No. 8,274,887

Claim	Disclosure
	network device 500, a memory 504 to store data, instructions or other information, and a path and bandwidth database 506.” <i>See Peterson at 9:22-36:</i> “A second fabric header set 604 may further include a source field (e.g., identifying a source such as an ingress device or ingress port) and a length field indicating a length of the fabric packet). A third fabric header set 606 illustrates an example header set that may be assembled when two network switches or two network processors are exchanging packets within the network switching system, e.g., across one or more intermediary devices. The header set 606 may include, for example, a flow field identifying a flow (or group of associated fabric packets), a format field to identify a format for the fabric packet, and one or more processing parameters. The flow, format and processing parameter fields may, for example, be used by an egress device within the network switching system to perform the required processing during egress handling of the packet.” <i>See Peterson at 12:10-21:</i> “A fabric packet may include a number of different headers as part of a header set. These headers may include, for example, a destination header, a source aggregate header, a length header, a format header, a flow header, a proxy parameters header, processing parameters, and a bandwidth management header (e.g., to provide or generate a queue path message, path state message and path rate message, as examples). These are merely some examples of the types of headers that may be included in a fabric packet or cell. In some cases, one or more of these headers may be optional, and may be provided in some different combinations. Some of these headers will be briefly described.” <i>See Peterson at 14:3-12:</i> “A flow header may also be used to identify a flow for the current packet. A flow may be, for example, a group of related packets, or one or more packets from a same source or application, and/or directed to a same destination. For example, a flow may be a source, destination and priority combination that defines a connection or aggregate of connections and/or that may originate from and/or are destined to a same or collocated network locations. Other types of flows may be used as well. The flow header may include a next header field to identify a next header after the current fabric header.” <i>See Peterson at 15:12-39:</i> “The queue state message 810 may be used by ingress devices to convey or indicate the state of one

Exhibit E30 – Invalidity of U.S. Patent No. 9,450,884 in view of Peterson – US Patent No. 8,274,887

Claim	Disclosure
	or more VOQs (virtual output queues) at the ingress device. The queue state message 810 may be, for example, a 1-bit value indicating an empty or non-empty state of each VOQ, or a multi-bit value indicating a level of fullness of each VOQ. Thus, the queue state message 810 may include a field indicating a state of each VOQ (or for one or more VOQs) of the reporting ingress device, for example. Other information for each VOQ may also be transmitted in a queue state message 810, such as a requested data transmission rate for the VOQ, a traffic priority for queued traffic or the VOQ, and other information. The queue state message 810 may include a queue state header, that may include node ID field identifying the ingress node that is reporting, an indication of the number of VOQs being reported (or an identification of which VOQs are being reported), and a queue state for each VOQ being reported, for example. This is merely an example format, and other formats may be used. The queue state message 810 may be transmitted (e.g., to other devices in the network switching system) when there is a change in queue state (e.g., empty to non-empty, or increases in fullness), or may be sent out at specific time periods or intervals regardless whether there has been a change in queue state. Or, both techniques may be used to send out queue state messages (e.g., both when a change in a queue state occurs and sending out a queue state message at specific times or intervals regardless of a change).” See Peterson at 16:36-17:6: “Congestion in a network switching system may occur when an offered load of network traffic exceeds the traffic capacity or bit rate of one or more resources of the system. According to an example embodiment, each of one or more ingress devices in a network switching system may have one or more traffic queues (e.g., VOQs). According to an example embodiment, each ingress device may have a plurality of VOQs, with each VOQ being associated with a destination (e.g., egress device, egress port, or virtual port on an egress device). According to an example embodiment, the ingress devices may exchange queue state information with the other ingress devices of the network switching system. The queue state information may be exchanged, for example, by each ingress device transmitting queue state messages to the other devices indicating a state of its VOQs (e.g., empty/non-empty, or indicating a level of fullness of each VOQ, or other VOQ information). For example, the queue state messages may allow one or more of the network devices to obtain a global or overall view of the total offered traffic load over one or more shared resources in the network switching system, for example. Each network device, based on the queue state messages, may determine a total offered traffic load for a shared resource within the network switching system, where the total offered traffic load may include a sum of offered traffic loads from each of one or

Exhibit E30 – Invalidity of U.S. Patent No. 9,450,884 in view of Peterson – US Patent No. 8,274,887

Claim	Disclosure
	more VOQs. Each network device may then adjust a transmission rate or traffic load from one or more of its VOQs that are using the shared resource, e.g., such that a traffic capacity (such as bit rate) of the shared resource is not exceeded by the total offered traffic load across the shared resource. For example, the shared resource may be a path segment or intermediary device that is being used by multiple VOQs (e.g., across different network devices). Based on this global view, a network device in the network switching system may then adjust its transmission rate for one or more of its VOQs so that a total offered traffic load preferably does not exceed the capacity of one or more shared resources within the network switching system, according to an example embodiment.” <i>See Peterson at 18:54-66:</i> “Each VOQ (VOQs 902 and 904) would like to transmit 30 Gbps. However, ingress device 1 transmits a queue state message (via an intermediary device to ingress device 2) indicating a queue state of VOQ 902 of non-empty, or full. Ingress device 2, likewise, may transmit a queue state message to ingress device 1 indicating that VOQ 904 has a queue state of non-empty or full. Thus, both ingress devices 1 and 2 may determine that there are two active VOQs which share the same resources, including path segment 907. Based on the offered traffic load of two active VOQs and a capacity of segment 907 of 18 Gbps, ingress devices 1 and 2 may adjust or decrease the transmission rate or traffic load from VOQs 902 and 904, respectively.” <i>See Peterson at 19:20-37:</i> “With reference to FIG. 9B, ingress devices 1 and 2 may exchange queue state messages so that each ingress device may obtain or determine a total offered traffic load for segment 922, which is a resource shared by both VOQs 902 and 904. Both VOQs 902 and 904 would like to transmit at a data rate of 30 Gbps each, for example (or at other rates that collectively may exceed the capacity of segment 922). However, segment 922 has a traffic capacity of 20 Gbps, which is shared by both VOQs 902 and 904. Therefore, for example, based on the known capacity of segment 922 (e.g., determined from path rate messages or from path and bandwidth database 506), and based on the number of active VOQs that share segment 922 (two active VOQs in this case), ingress device 1 and 2 may each adjust an offered traffic load for VOQ 902 and 904, respectively, such that the total offered traffic load from all ingress devices or from all VOQs in the network switching system that share the resource (segment 922 in this example) does not exceed the capacity of segment 922.” <i>See Peterson at 20:19-41:</i> “According to an example embodiment, in-band control messages, such as queue state messages

Exhibit E30 – Invalidity of U.S. Patent No. 9,450,884 in view of Peterson – US Patent No. 8,274,887

Claim	Disclosure
	810, path state messages 820, and path rate messages 830 may be passed between the ingress devices 1 and 2 (FIGS. 9A and 9B) to convey the state of the VOQs 902 and 904. For example, an empty state, a full state, an intermediate state, or a more precise representation of the capacity of the VOQ and how many packets, cells, requested traffic rate, traffic priority (or traffic class) for each queue, or other data associated with the VOQ may be shared between the ingress devices 1 and 2. In this manner, one or more (or even all) of the ingress devices and/or one or more of the VOQs 902 and 904 in the network switching system may develop a global or aggregate view of the states of the other VOQs and/or their forwarding requirements. Once the aggregate state of all of the VOQs in the network switching system is known, the various network devices of the system may know the demand on the switch fabric or the stacking interconnection, path segments or other resources. The system and/or each network device may also know where the congestion points are (e.g., based on traffic capacity of each resource or segment) and how much more data the VOQs 902 and 904 in the ingress devices 1 and 2 can hold.” See Peterson at 21:4-16: “This per-ingress-device awareness of the global queue state may be combined with a map of the various paths (or path segments) through the network switching system and the capacity of those segments such that each ingress device is able to determine the total requested (or offered) load on each path segment or resource. This requested (or total offered) traffic load may then be compared to each segment’s capacity and each ingress device may make an independent determination of the fraction of the available capacity that it may be going to claim for itself. The allocation algorithms may ensure that the sum of the individual allocations will never exceed the capacity of the slowest shared path. Thus, congestion may be avoided.” See Peterson at 22:61-23:8: “For example, operation 1010 may include each of a plurality of ingress devices of the network switching system reporting a state of one or more of its VOQs such that one or more ingress devices of the network switching system obtains a global view of a demand (or offered traffic load) upon one or more shared resources of the network switching system. For example, queue state messages (e.g., indicating empty/non-empty state or a level of fullness) may be exchanged between network devices of the network switching system to allow one or more network devices of the system to determine a state (e.g., active or non-empty) of one or more VOQs. Each of the VOQs may be associated with a destination, and a path from an ingress device to an egress device may include one or more resources that may be shared among VOQs of the network switching system.”

Exhibit E30 – Invalidity of U.S. Patent No. 9,450,884 in view of Peterson – US Patent No. 8,274,887

Claim	Disclosure
	See Peterson at 23:17-20: “‘At 1030, a transmission rate from one or more of the active VOQs over the shared resource may be adjusted such that the total traffic load from the active VOQs does not exceed the capacity of the shared resource.’”
[1c] determining, by the network switching element, a fair-share bandwidth allocation for the target port, the fair-share bandwidth allocation being a proportional allocation of a total bandwidth of the network switching element; and	Peterson discloses determining, by the network switching element, a fair-share bandwidth allocation for the target port, the fair-share bandwidth allocation being a proportional allocation of a total bandwidth of the network switching element. See Peterson at 3:24-42: “‘According to an example embodiment, each ingress device (or network device) may determine an overall or global traffic load offered for one or more destinations, path segments, or resources in a network switching system, e.g., based on received queue state messages received from the other ingress devices in the network switching system. The queue state messages may indicate, for example, a queue state (e.g., an empty/non-empty state or a level of fullness of the queue. A resource (e.g., path segment, intermediary device, or switch fabric device) may be shared among multiple active (e.g., non-empty) queues. The queues may be, for example, virtual output queues (VOQs). Based on a total offered traffic load for a shared resource, e.g., including traffic from multiple VOQs, one or more network devices in the network switching system may adjust a transmission rate from one or more of its active queues or VOQs over the shared resource such that the total traffic load from a plurality of queues does not exceed the bandwidth or traffic capacity (e.g., bit rate capacity) of the shared resource.’” See Peterson at 3:43-67: “‘Thus, for example, by exchanging queue state information between network devices in a network switching system, congestion in the switching system may be avoided in some cases (e.g., anticipated and transmission rates adjusted to avoid congestion). By exchanging queue state information between ingress devices, an overall or global offered traffic load for a resource (e.g., path segment, fabric device or other resource) may be determined by each network device in the network switching system. Based on this global demand of a shared resource, each network device may, for example, adjust a transmission rate for one or more of its active VOQs that are using the resource to use up to, for example, approximately a fair portion (or other amount) of the bandwidth or traffic (e.g., bit rate) capacity, such as approximately 1/Nth of the bandwidth or traffic capacity for the shared resource for N VOQs sharing the resource, according to an example embodiment. For example, if the traffic capacity of a resource (e.g., path segment) is 9 Mbps and there are 3 VOQs

Exhibit E30 – Invalidity of U.S. Patent No. 9,450,884 in view of Peterson – US Patent No. 8,274,887

Claim	Disclosure
	sharing that resource, then each VOQ may adjust its transmission rate to 1/3rd of the traffic capacity (1/3rd of 9 Mbps), or to 3 Mbps per VOQ. In this manner, the group of VOQs or network devices transmitting across a shared resource may, at least in some cases, avoid congestion by typically not exceeding the bandwidth or traffic (e.g., bit rate) capacity of the resource, for example.” <i>See Peterson at 4:21-36:</i> “Although interface 114 may be referred to as a fabric interface, the fabric interface 114 may provide an interface to allow network devices to communicate with each other within a network switching system, both in the presence of a switch fabric (or other intermediary devices) and in the absence of a switch fabric or other intermediary device coupled between the network devices. Therefore, network devices within switching system 100 may communicate via a fabric interface 114, either directly as shown in FIG. 1, or via one or more intermediary devices (e.g., via switch fabric devices, shared memory switches, buffered cross-bar switches or other intermediary devices), not shown in FIG. 1. Thus, the term fabric interface is not limited to a switching system that employs a fabric or other intermediary devices, but may include the case where network devices may be coupled directly together, as shown in FIG. 1, for example.” <i>See Peterson at 5:43-64:</i> “Also, although in this example embodiment, each switch or device may be interconnected to other switches via one hop (e.g., via one intermediary device), in other example embodiment, one or more network devices may be coupled to one or more other network devices via multiple hops (e.g., via multiple switch fabric devices, or multiple crossbar switches or shared memory switches). Thus, one or more paths may be provided between ports of each switch as an ingress device, via one or more intermediary devices, to another switch as an egress device. The intermediary devices and path segments (or links) between intermediary devices (or switch fabric devices) or between an intermediary device and a switch may be shared among multiple paths (or shared among multiple destinations). Thus, the term shared resource may be used to generally describe any resource (e.g., intermediary device, switch fabric device, crossbar switch, shared memory switch or other device) that may be shared, e.g., shared between multiple paths or shared between multiple network devices or switches, for example. The network switching system illustrated in FIGS. 1-4 are merely some example system configurations, and many other types or configurations of network switching systems may be used.” <i>See Peterson 17:45-58:</i> “For example, based on a capacity of the shared resource (e.g., indicated by path and bandwidth

Exhibit E30 – Invalidity of U.S. Patent No. 9,450,884 in view of Peterson – US Patent No. 8,274,887

Claim	Disclosure
	database 506 and/or path rate messages 830), an ingress device may then adjust an offered traffic load or transmission rate for one or more of its active VOQs (e.g., which are using the shared resource) such that the total offered traffic load from the active VOQs (e.g., across multiple ingress devices) may not exceed the capacity of the shared resource. For example, an ingress device may adjust a transmission rate from one (or more) of its VOQs transmitting across the shared resource to be approximately 1/Nth of the total capacity of the shared resource (or 1/Nth of the lowest shared resource in a path), where N may be the number of active VOQs offering traffic load over the shared resource, for example.” <i>See Peterson at 18:54-19:7:</i> “Each VOQ (VOQs 902 and 904) would like to transmit 30 Gbps. However, ingress device 1 transmits a queue state message (via an intermediary device to ingress device 2) indicating a queue state of VOQ 902 of non-empty, or full. Ingress device 2, likewise, may transmit a queue state message to ingress device 1 indicating that VOQ 904 has a queue state of non-empty or full. Thus, both ingress devices 1 and 2 may determine that there are two active VOQs which share the same resources, including path segment 907. Based on the offered traffic load of two active VOQs and a capacity of segment 907 of 18 Gbps, ingress devices 1 and 2 may adjust or decrease the transmission rate or traffic load from VOQs 902 and 904, respectively. For example, the ingress devices may adjust the transmission rate for each VOQ to a fair portion of the capacity of the shared resource, or 1/Nth of the capacity, where N maybe the total number of VOQs sharing the resource (N=2 in this example). In this example, N=2, and the capacity of the path segment is 18 Gbps. Therefore, ingress devices 1 and 2 may adjust the traffic rate from VOQs 902 and 904, respectively, to 9 Gbps each, as shown in FIG. 9A.” <i>See Peterson at 19:38-47:</i> “For example, referring to FIG. 9B, ingress devices 1 and 2 may adjust a traffic rate of VOQs 902 and 904, respectively, up to a fair portion or ½ of the traffic capacity of segment 922. In this example, ingress device 1 may adjust the transmission rate or traffic load from VOQ 902 to approximately (or up to) 10 Gbps, and ingress device 2 may adjust the transmission rate or traffic load from VOQ 904 to approximately (or up to) 10 Gbps. Thus, the total offered traffic load from the VOQs that share this resource (path segment 922) will typically not exceed the capacity of segment 922.” <i>See Peterson at 20:56-21:3:</i> “Depending on the in-band control messages 810, 820 and 830, the VOQ state information may

Exhibit E30 – Invalidity of U.S. Patent No. 9,450,884 in view of Peterson – US Patent No. 8,274,887

Claim	Disclosure
	cause the ingress devices to adjust their behavior or allocation to the egress device or destination. For example, the information may be used independently by each ingress device to adjust its VOQs' shaper settings (data transmission rates) so that congestion may be avoided or at least decreased. There may be more than 2 ingress devices, but in the current example, each of the ingress devices 1 and 2 may set their capacity to 50% of the capacity of the shared segment or resource, for example. Likewise, if there were four ingress devices (e.g., one VOQ per ingress device) to a single egress device, the VOQ transmission rate may be scaled for each of the four VOQs may be scaled back to 25% of the traffic capacity of the shared segment.” <i>See Peterson at 21:17-22:</i> “The available traffic capacity of a segment or resource may be allocated to multiple VOQs using a variety of different techniques or algorithms. The capacity of a resource may be allocated using a fair allocation or fair apportionment, e.g., 1/Nth of the capacity for each of the N VOQs that use the shared resource.” <i>See Peterson at 22:32-41:</i> “Also, in another embodiment, one or more ingress devices being configured to adjust a transmission rate may include the device being configured to adjust a transmission rate from each of the VOQs that use the shared resource to be approximately 1/Nth of the capacity of the shared resource based on N active VOQs offering traffic load over the shared resource. The shared resource may include a path segment between a VOQ and a destination of the network switching system, the destination may include one of an egress network device, an egress port, or a virtual port on an egress network device.” <i>See Peterson at 23:17-20:</i> “At 1030, a transmission rate from one or more of the active VOQs over the shared resource may be adjusted such that the total traffic load from the active VOQs does not exceed the capacity of the shared resource. For example, operation 1030 may include adjusting a transmission rate from each (or one or more) of the VOQs that use the shared resource to be approximately a fair share (e.g., 1/Nth) of the capacity of the shared resource based on a number of active VOQs offering traffic load over the shared resource and the capacity of the shared resource.” <i>See also, e.g., Peterson at claim 7 and claim 15.</i>

Exhibit E30 – Invalidity of U.S. Patent No. 9,450,884 in view of Peterson – US Patent No. 8,274,887

Claim	Disclosure
[1d] adjusting, by the network switching element, the bandwidth allocation for the target port based on the fair-share bandwidth allocation.	Peterson discloses adjusting, by the network switching element, the bandwidth allocation for the target port based on the fair-share bandwidth allocation. See Peterson at 3:24-42: “According to an example embodiment, each ingress device (or network device) may determine an overall or global traffic load offered for one or more destinations, path segments, or resources in a network switching system, e.g., based on received queue state messages received from the other ingress devices in the network switching system. The queue state messages may indicate, for example, a queue state (e.g., an empty/non-empty state or a level of fullness of the queue. A resource (e.g., path segment, intermediary device, or switch fabric device) may be shared among multiple active (e.g., non-empty) queues. The queues may be, for example, virtual output queues (VOQs). Based on a total offered traffic load for a shared resource, e.g., including traffic from multiple VOQs, one or more network devices in the network switching system may adjust a transmission rate from one or more of its active queues or VOQs over the shared resource such that the total traffic load from a plurality of queues does not exceed the bandwidth or traffic capacity (e.g., bit rate capacity) of the shared resource.” See Peterson at 3:43-67: “Thus, for example, by exchanging queue state information between network devices in a network switching system, congestion in the switching system may be avoided in some cases (e.g., anticipated and transmission rates adjusted to avoid congestion). By exchanging queue state information between ingress devices, an overall or global offered traffic load for a resource (e.g., path segment, fabric device or other resource) may be determined by each network device in the network switching system. Based on this global demand of a shared resource, each network device may, for example, adjust a transmission rate for one or more of its active VOQs that are using the resource to use up to, for example, approximately a fair portion (or other amount) of the bandwidth or traffic (e.g., bit rate) capacity, such as approximately 1/Nth of the bandwidth or traffic capacity for the shared resource for N VOQs sharing the resource, according to an example embodiment. For example, if the traffic capacity of a resource (e.g., path segment) is 9 Mbps and there are 3 VOQs sharing that resource, then each VOQ may adjust its transmission rate to 1/3rd of the traffic capacity (1/3rd of 9 Mbps), or to 3 Mbps per VOQ. In this manner, the group of VOQs or network devices transmitting across a shared resource may, at least in some cases, avoid congestion by typically not exceeding the bandwidth or traffic (e.g., bit rate) capacity of the resource, for example.”

Exhibit E30 – Invalidity of U.S. Patent No. 9,450,884 in view of Peterson – US Patent No. 8,274,887

Claim	Disclosure
	See Peterson at 4:21-36: “Although interface 114 may be referred to as a fabric interface, the fabric interface 114 may provide an interface to allow network devices to communicate with each other within a network switching system, both in the presence of a switch fabric (or other intermediary devices) and in the absence of a switch fabric or other intermediary device coupled between the network devices. Therefore, network devices within switching system 100 may communicate via a fabric interface 114, either directly as shown in FIG. 1, or via one or more intermediary devices (e.g., via switch fabric devices, shared memory switches, buffered cross-bar switches or other intermediary devices), not shown in FIG. 1. Thus, the term fabric interface is not limited to a switching system that employs a fabric or other intermediary devices, but may include the case where network devices may be coupled directly together, as shown in FIG. 1, for example.” See Peterson at 5:43-64: “Also, although in this example embodiment, each switch or device may be interconnected to other switches via one hop (e.g., via one intermediary device), in other example embodiment, one or more network devices may be coupled to one or more other network devices via multiple hops (e.g., via multiple switch fabric devices, or multiple crossbar switches or shared memory switches). Thus, one or more paths may be provided between ports of each switch as an ingress device, via one or more intermediary devices, to another switch as an egress device. The intermediary devices and path segments (or links) between intermediary devices (or switch fabric devices) or between an intermediary device and a switch may be shared among multiple paths (or shared among multiple destinations). Thus, the term shared resource may be used to generally describe any resource (e.g., intermediary device, switch fabric device, crossbar switch, shared memory switch or other device) that may be shared, e.g., shared between multiple paths or shared between multiple network devices or switches, for example. The network switching system illustrated in FIGS. 1-4 are merely some example system configurations, and many other types or configurations of network switching systems may be used.” See Peterson 17:45-58: “For example, based on a capacity of the shared resource (e.g., indicated by path and bandwidth database 506 and/or path rate messages 830), an ingress device may then adjust an offered traffic load or transmission rate for one or more of its active VOQs (e.g., which are using the shared resource) such that the total offered traffic load from the active VOQs (e.g., across multiple ingress devices) may not exceed the capacity of the shared resource. For example, an ingress device may adjust a transmission rate from one (or more) of its VOQs transmitting across the shared resource to

Exhibit E30 – Invalidity of U.S. Patent No. 9,450,884 in view of Peterson – US Patent No. 8,274,887

Claim	Disclosure
	be approximately 1/Nth of the total capacity of the shared resource (or 1/Nth of the lowest shared resource in a path), where N may be the number of active VOQs offering traffic load over the shared resource, for example.” <i>See Peterson at 18:54-19:7:</i> “Each VOQ (VOQs 902 and 904) would like to transmit 30 Gbps. However, ingress device 1 transmits a queue state message (via an intermediary device to ingress device 2) indicating a queue state of VOQ 902 of non-empty, or full. Ingress device 2, likewise, may transmit a queue state message to ingress device 1 indicating that VOQ 904 has a queue state of non-empty or full. Thus, both ingress devices 1 and 2 may determine that there are two active VOQs which share the same resources, including path segment 907. Based on the offered traffic load of two active VOQs and a capacity of segment 907 of 18 Gbps, ingress devices 1 and 2 may adjust or decrease the transmission rate or traffic load from VOQs 902 and 904, respectively. For example, the ingress devices may adjust the transmission rate for each VOQ to a fair portion of the capacity of the shared resource, or 1/Nth of the capacity, where N maybe the total number of VOQs sharing the resource (N=2 in this example). In this example, N=2, and the capacity of the path segment is 18 Gbps. Therefore, ingress devices 1 and 2 may adjust the traffic rate from VOQs 902 and 904, respectively, to 9 Gbps each, as shown in FIG. 9A.” <i>See Peterson at 19:38-47:</i> “For example, referring to FIG. 9B, ingress devices 1 and 2 may adjust a traffic rate of VOQs 902 and 904, respectively, up to a fair portion or ½ of the traffic capacity of segment 922. In this example, ingress device 1 may adjust the transmission rate or traffic load from VOQ 902 to approximately (or up to) 10 Gbps, and ingress device 2 may adjust the transmission rate or traffic load from VOQ 904 to approximately (or up to) 10 Gbps. Thus, the total offered traffic load from the VOQs that share this resource (path segment 922) will typically not exceed the capacity of segment 922.” <i>See Peterson at 20:56-21:3:</i> “Depending on the in-band control messages 810, 820 and 830, the VOQ state information may cause the ingress devices to adjust their behavior or allocation to the egress device or destination. For example, the information may be used independently by each ingress device to adjust its VOQs’ shaper settings (data transmission rates) so that congestion may be avoided or at least decreased. There may be more than 2 ingress devices, but in the current example, each of the ingress devices 1

Exhibit E30 – Invalidity of U.S. Patent No. 9,450,884 in view of Peterson – US Patent No. 8,274,887

Claim	Disclosure
	and 2 may set their capacity to 50% of the capacity of the shared segment or resource, for example. Likewise, if there were four ingress devices (e.g., one VOQ per ingress device) to a single egress device, the VOQ transmission rate may be scaled for each of the four VOQs may be scaled back to 25% of the traffic capacity of the shared segment.” <i>See Peterson at 21:17-22:</i> “The available traffic capacity of a segment or resource may be allocated to multiple VOQs using a variety of different techniques or algorithms. The capacity of a resource may be allocated using a fair allocation or fair apportionment, e.g., 1/Nth of the capacity for each of the N VOQs that use the shared resource.” <i>See Peterson at 22:32-41:</i> “Also, in another embodiment, one or more ingress devices being configured to adjust a transmission rate may include the device being configured to adjust a transmission rate from each of the VOQs that use the shared resource to be approximately 1/Nth of the capacity of the shared resource based on N active VOQs offering traffic load over the shared resource. The shared resource may include a path segment between a VOQ and a destination of the network switching system, the destination may include one of an egress network device, an egress port, or a virtual port on an egress network device.” <i>See Peterson at 23:17-20:</i> “At 1030, a transmission rate from one or more of the active VOQs over the shared resource may be adjusted such that the total traffic load from the active VOQs does not exceed the capacity of the shared resource. For example, operation 1030 may include adjusting a transmission rate from each (or one or more) of the VOQs that use the shared resource to be approximately a fair share (e.g., 1/Nth) of the capacity of the shared resource based on a number of active VOQs offering traffic load over the shared resource and the capacity of the shared resource.” <i>See also, e.g., Peterson at claim 7 and claim 15.</i>
[11pre] A method of adjusting bandwidth allocation by a network controller in a communications network, the network controller being	To the extent the preamble is limiting, Peterson discloses this preamble. <i>See claim element [1pre] supra.</i>

Exhibit E30 – Invalidity of U.S. Patent No. 9,450,884 in view of Peterson – US Patent No. 8,274,887

Claim	Disclosure
configured to control a network switching element, the network switching element including a target port, the method comprising:	
[11a] receiving, by the network controller, data flow information from the network switching element, the data flow information including information about data flows traversing of the target port of the network switching element;	<i>See claim element [1a] supra.</i>
[11b] determining, by the network controller, a bandwidth allocation of the target port based on the data flow information, the bandwidth allocation of the target port being a bandwidth that is currently allocated for each of the data flows traversing the target port;	<i>See claim element [1b] supra.</i>
[11c] determining, by the network controller, an over-subscription ratio, the over-subscription ratio being a ratio of the bandwidth allocation of the target port to a number of data flows traversing the target port;	Peterson discloses determining, by the network controller, an over-subscription ratio, the over-subscription ratio being a ratio of the bandwidth allocation of the target port to a number of data flows traversing the target part. <i>See claim element [1b] supra.</i> <i>See Peterson at 1:36-49:</i> “Congestion, in some cases, may be a problem for networking switching systems. Congestion may occur when an offered load of network traffic exceeds the processing or forwarding capacity of the affected element (e.g., switch, port, line card). Congestion may result from, for example, speed

Exhibit E30 – Invalidity of U.S. Patent No. 9,450,884 in view of Peterson – US Patent No. 8,274,887

Claim	Disclosure
	mismatches (e.g., fast ports forwarding to slow ports), oversubscription (e.g., multiple ports forwarding to a single port or group of ports of lesser total bandwidth), etc. In some network switching systems, FIFOs (first in, first out buffers) are used to absorb excess traffic near a congestion point. When a FIFO fullness at an egress device exceeds a congestion threshold, an indication is usually provided to the traffic source(s). This is an example of a system that detects congestion, and then reacts to the congestion.” <i>See Peterson at 19:64-20:18:</i> “In a network switching system with multiple ports, congestion may be a possibility. Congestion may occur in communications systems when the offered load of network traffic may exceed the processing or forwarding capacity of the affected element. For example, port speed mismatches (fast ports forwarding to slow ports) and oversubscription (multiple ports forwarding to a single port of lesser total bandwidth) may cause congestion. Buffers, such as VOQs, may be used to absorb the excess traffic upstream of the congestion point. Buffers, FIFOs or other structured are of finite size and may, eventually, be overrun by bursts of congestion if those bursts are sufficiently long or frequent. To prevent congestion, according to an example embodiment, each source of traffic may be able to modulate or adjust the rate at which it offers traffic to egress devices, such as switch fabric devices, so that the total amount of traffic addressed to each fabric egress port does not exceed the capacity of those ports. When the ports on ingress devices reside on multiple discrete devices, then the management of the congestion may involve the communication of congestion-related information between the affected network devices within the network switching system.” <i>See Peterson at 20:19-41:</i> “According to an example embodiment, in-band control messages, such as queue state messages 810, path state messages 820, and path rate messages 830 may be passed between the ingress devices 1 and 2 (FIGS. 9A and 9B) to convey the state of the VOQs 902 and 904. For example, an empty state, a full state, an intermediate state, or a more precise representation of the capacity of the VOQ and how many packets, cells, requested traffic rate, traffic priority (or traffic class) for each queue, or other data associated with the VOQ may be shared between the ingress devices 1 and 2. In this manner, one or more (or even all) of the ingress devices and/or one or more of VOQs 902 and 904 in the network switching system may develop a global or aggregate view of the states of the other VOQs and/or their forwarding requirements. Once the aggregate state of all of the VOQs in the network switching system is known, the various network devices of the system may know the demand on the switch fabric or the stacking interconnection, path segments or other resources. The

Exhibit E30 – Invalidity of U.S. Patent No. 9,450,884 in view of Peterson – US Patent No. 8,274,887

Claim	Disclosure
	system and/or each network device may also know where the congestion points are (e.g., based on a traffic capacity of each resource or segment) and how much more data the VOQs 902 and 904 in the ingress devices 1 and 2 can hold.” <i>See Peterson at 20:50-55:</i> “In multicast scenarios, each of the VOQs 902, 904 may require the use of several paths. These paths may also be used by unicast queues (VOQs) and by other multicast VOQs. Therefore, the rate adjustment of the individual queue may take these combinations into account so that none of the paths become oversubscribed, for example.”
[11d] transmitting, by the network controller, the over-subscription ratio to the network switching element based on the over-subscription ratio and a threshold value;	<i>See claim element [1c] supra.</i>
[11e] determining, by the network controller, a fair-share bandwidth allocation for the target port based on the over-subscription ratio and the threshold value, the fair-share bandwidth allocation being a proportional allocation of a total bandwidth of the network switching element; and	<i>See claim element [1c] supra.</i>
[11f] adjusting, by the network controller, the bandwidth allocation for the target port based on the fair-share bandwidth allocation.	<i>See claim element [1d] supra.</i>
[17pre] An edge switch for adjusting bandwidth allocation	To the extent the preamble is limiting, Peterson discloses this preamble.

Exhibit E30 – Invalidity of U.S. Patent No. 9,450,884 in view of Peterson – US Patent No. 8,274,887

Claim	Disclosure
in a communications network, the edge switch including a target port, the edge switch configured to:	<i>See claim element [1pre] supra.</i>
[17a] monitor a data flow traversing the target port;	<i>See claim element [1a] supra.</i>
[17b] determine a bandwidth allocation for the target port, the bandwidth allocation for the target port being a bandwidth that is currently allocated for the data flow;	<i>See claim element [1b] supra.</i>
[17c] determine a fair-share bandwidth allocation for the target port, the fair-share bandwidth allocation being a proportional allocation of a total bandwidth of the network switching element; and	<i>See claim element [1c] supra.</i>
[17d] adjust the bandwidth allocation for the target port based on the fair-share bandwidth allocation.	<i>See claim element [1d] supra.</i>
[20pre] A Software Defined Networking (SDN) controller for adjusting bandwidth allocation in a communications network, the SDN controller being configured to control an aggregation switch, the aggregation switch including a target port, the SDN controller	To the extent the preamble is limiting, Peterson discloses this preamble. <i>See claim element [1pre] supra.</i>

Exhibit E30 – Invalidity of U.S. Patent No. 9,450,884 in view of Peterson – US Patent No. 8,274,887

Claim	Disclosure
is configured to:	
[20a] receive data flow information from the network switching element, the data flow information including information about data flows traversing of the target port;	<i>See claim element [1a] supra; see claim element [11a] supra.</i>
[20b] determine a bandwidth allocation of the target port based on the data flow information, the bandwidth allocation of the target port being a bandwidth that is currently allocated for each of the data flows traversing the target port;	<i>See claim element [1b] supra; see claim element [11b] supra.</i>
[20c] determine an over-subscription ratio, the over-subscription ratio being a ratio of the bandwidth allocation of the target port to a number of data flows traversing the target port; and	<i>See claim element [11c] supra.</i>
[20d] transmit the over-subscription ratio to the aggregation switch based on the over-subscription ratio and a threshold value;	<i>See claim element [1c] supra; see claim element [11d] supra.</i>
[20e] determine a fair-share bandwidth allocation for the target port based on the over-subscription ratio and the threshold value, the fair-share bandwidth allocation being a	<i>See claim element [1c] supra; see claim element [11e] supra.</i>

Exhibit E30 – Invalidity of U.S. Patent No. 9,450,884 in view of Peterson – US Patent No. 8,274,887

Claim	Disclosure
proportional allocation of a total bandwidth of the network switching element; and	
[20f] adjust the bandwidth allocation for the target port based on the fair-share bandwidth allocation.	<i>See claim element [1d] supra; see claim element [11f] supra.</i>