



(19) **United States**

(12) **Patent Application Publication**

LEE et al.

(10) **Pub. No.: US 2008/0225723 A1**

(43) **Pub. Date: Sep. 18, 2008**

(54) **OPTICAL IMPAIRMENT AWARE PATH COMPUTATION ARCHITECTURE IN PCE BASED NETWORK**

Related U.S. Application Data

(60) Provisional application No. 60/895,283, filed on Mar. 16, 2007.

(75) Inventors: **Young LEE**, Plano, TX (US);
Linda DUNBAR, Plano, TX (US)

Publication Classification

(51) **Int. Cl.**
G06F 11/00 (2006.01)
(52) **U.S. Cl.** **370/235**
(57) **ABSTRACT**

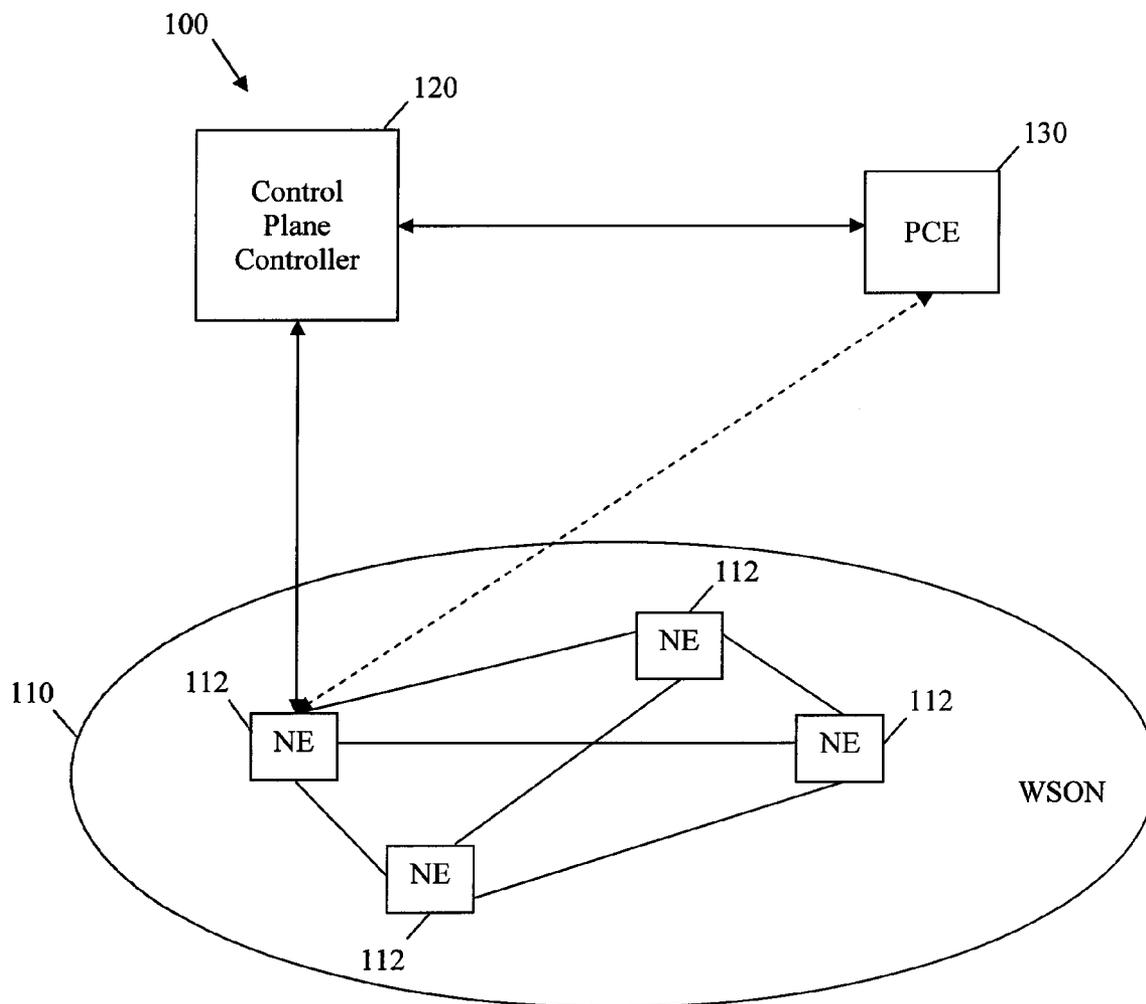
Correspondence Address:
CONLEY ROSE, P.C.
5601 GRANITE PARKWAY, SUITE 750
PLANO, TX 75024 (US)

An apparatus comprising at least one processor configured to implement a method comprising receiving a path computation request comprising at least one path computation constraint, and determining whether there is a path through an optical network that satisfies the path computation constraints. Also disclosed is an apparatus configured to process a data structure comprising a flags field comprising at least one flag having one of an active state or an inactive state, wherein each flag is representative of an optical quality constraint.

(73) Assignee: **FUTUREWEI TECHNOLOGIES, INC.**, Plano, TX (US)

(21) Appl. No.: **12/046,557**

(22) Filed: **Mar. 12, 2008**



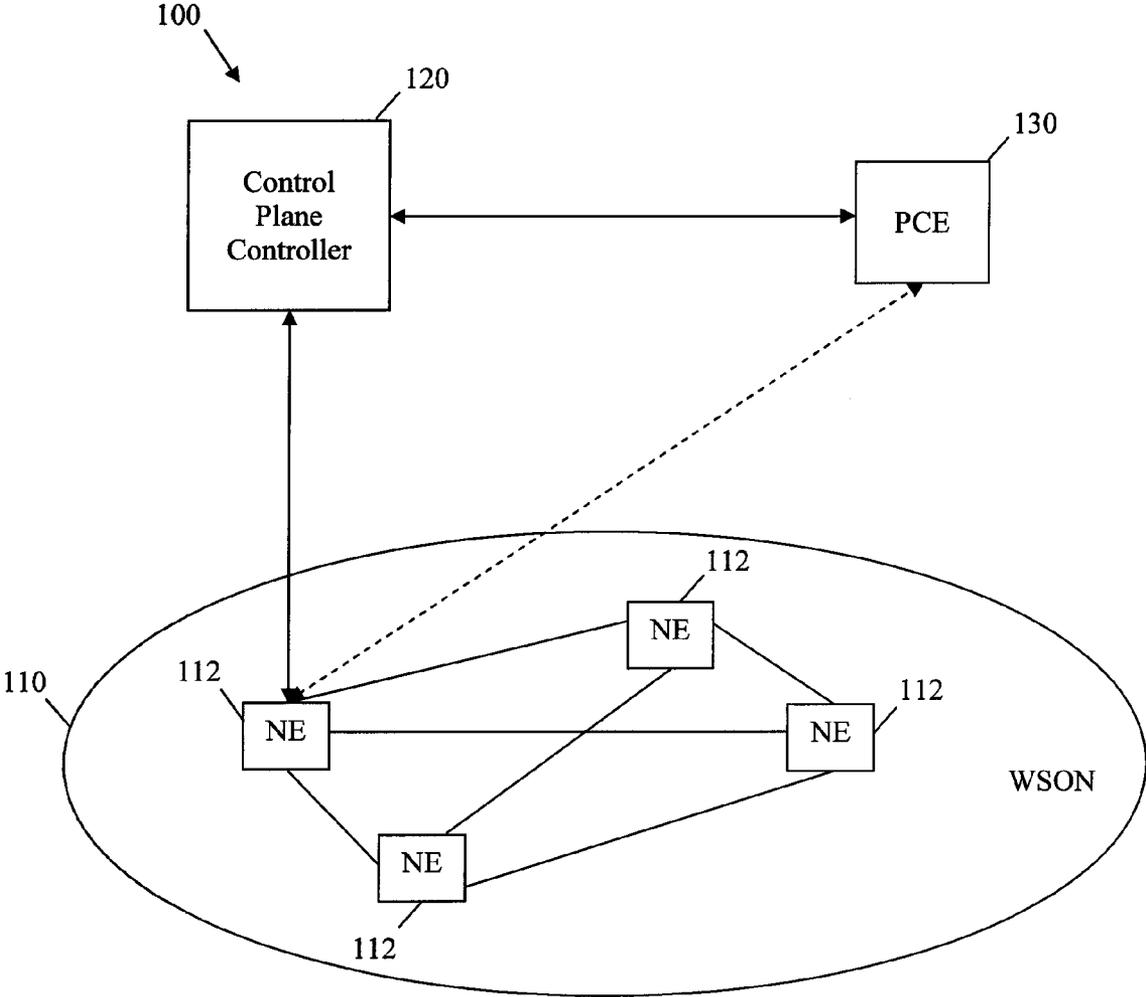


FIG. 1

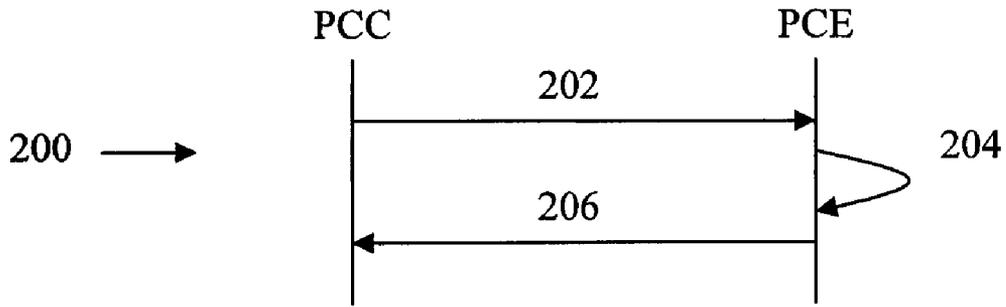


FIG. 2

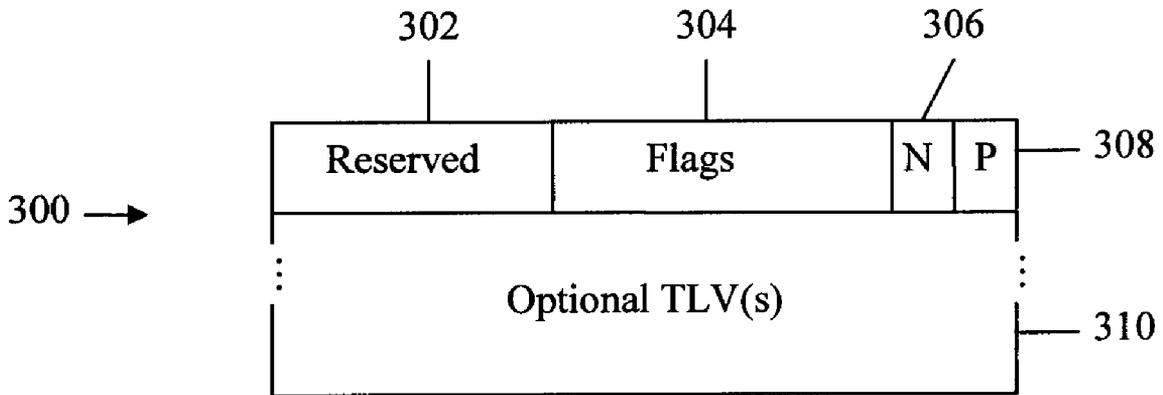


FIG. 3

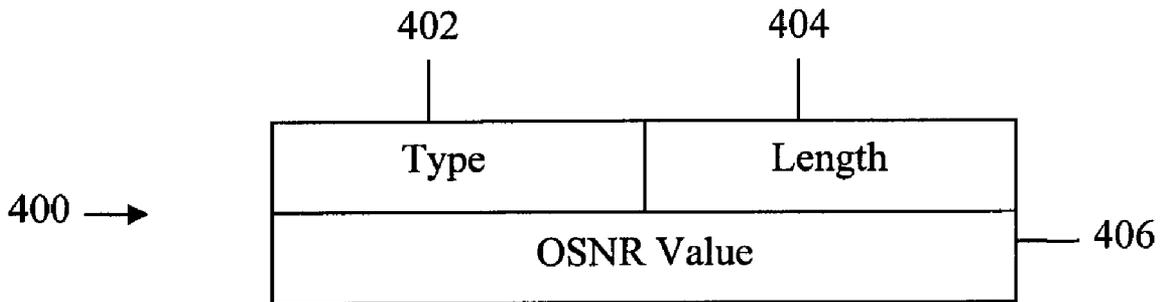


FIG. 4

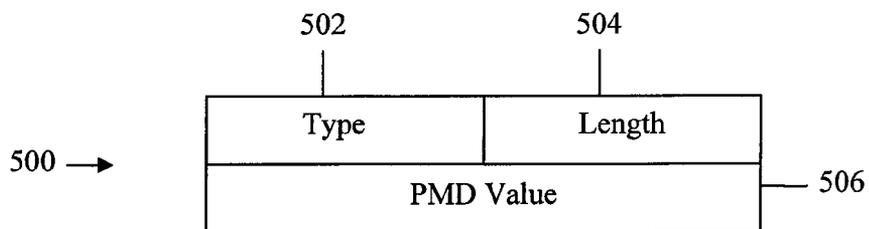


FIG. 5

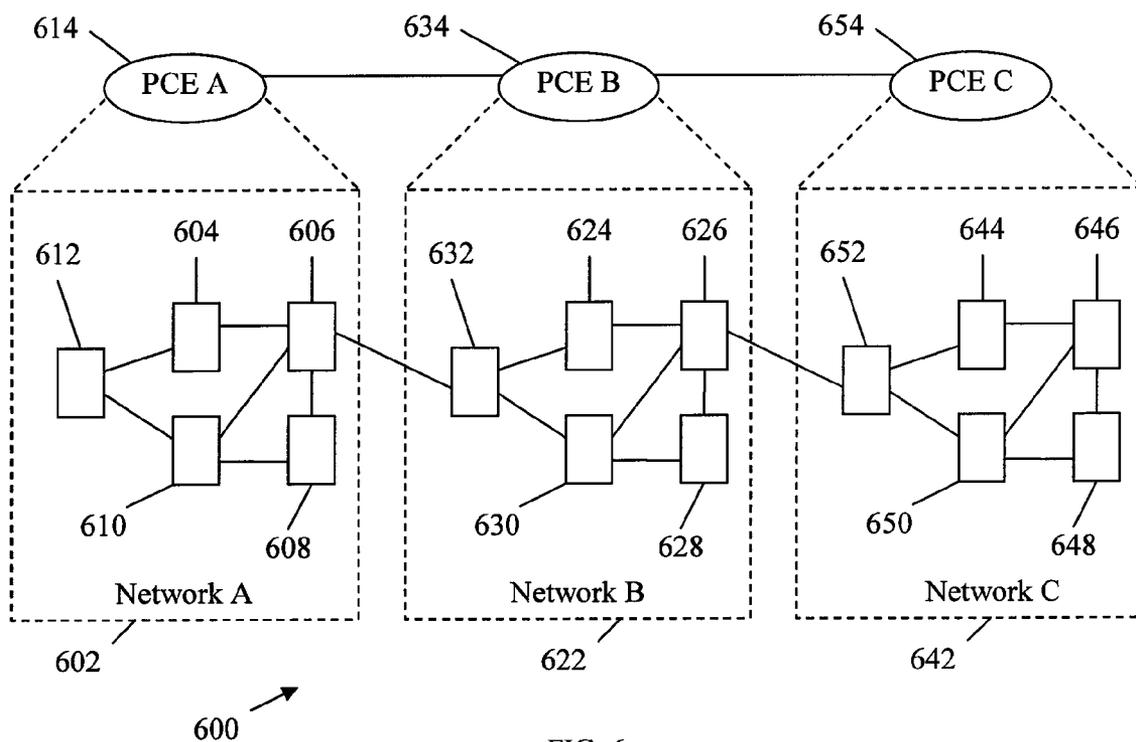


FIG. 6

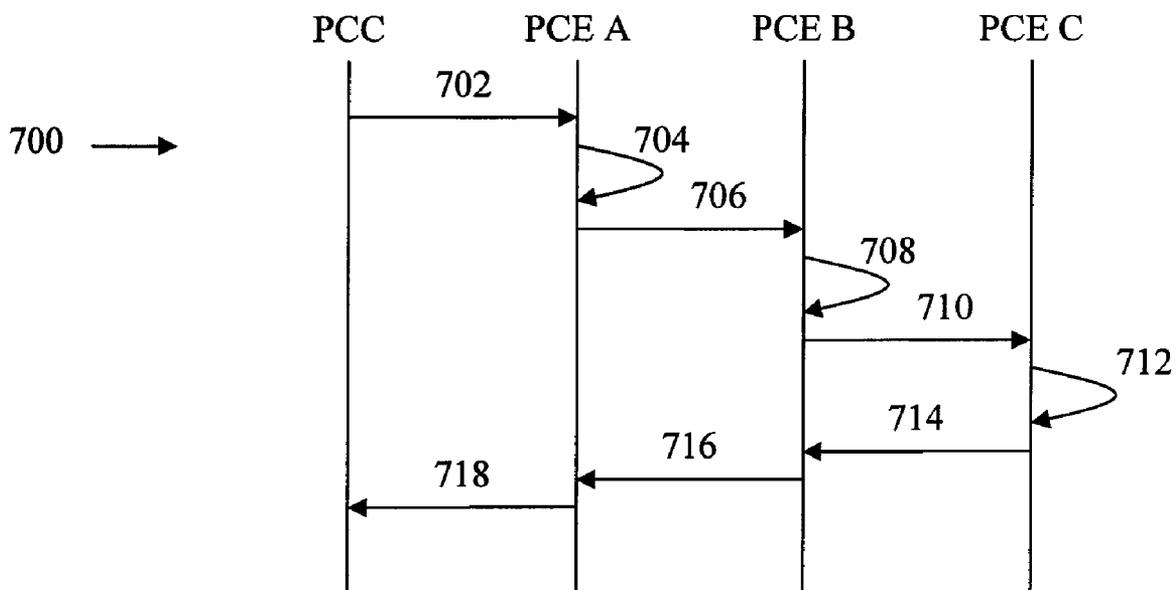


FIG. 7

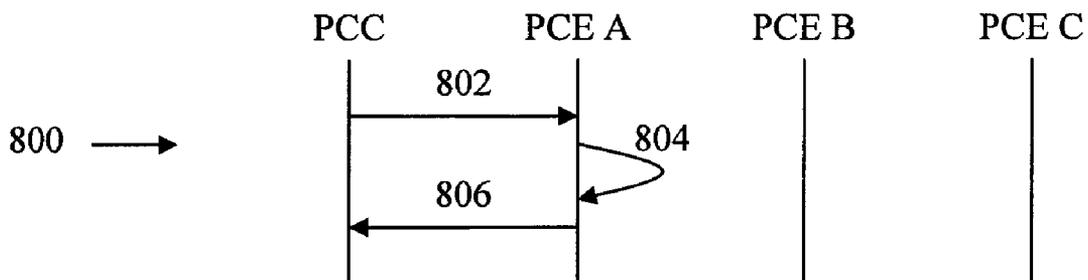


FIG. 8

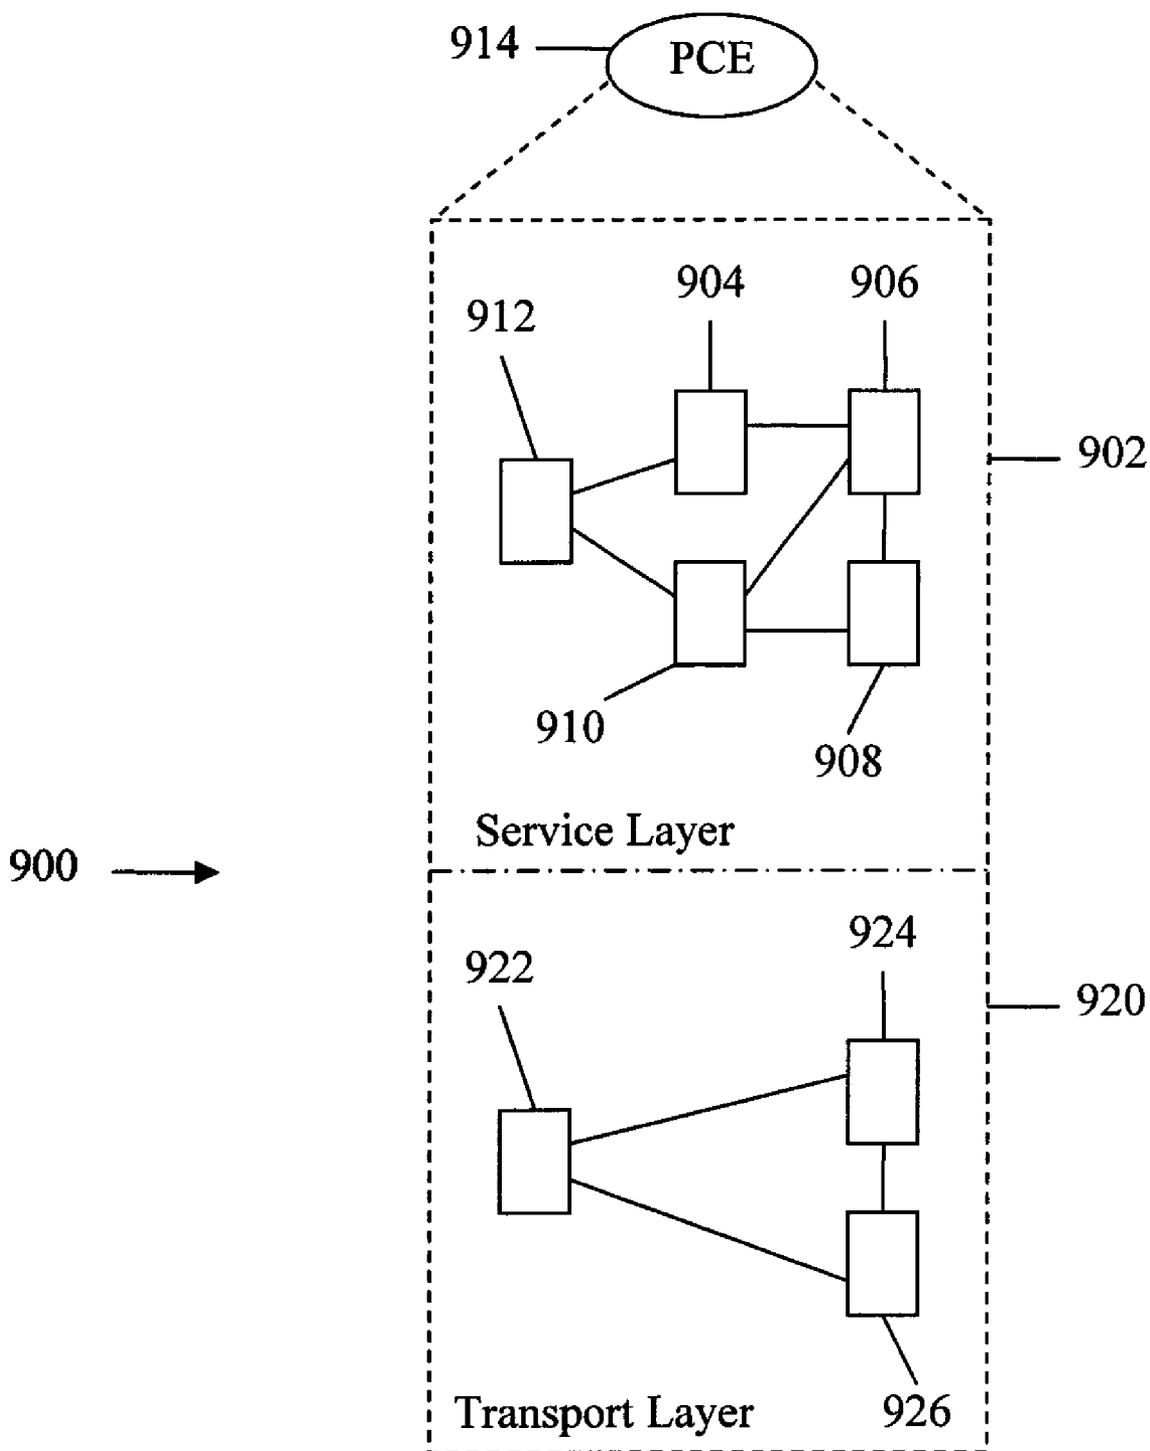


FIG. 9

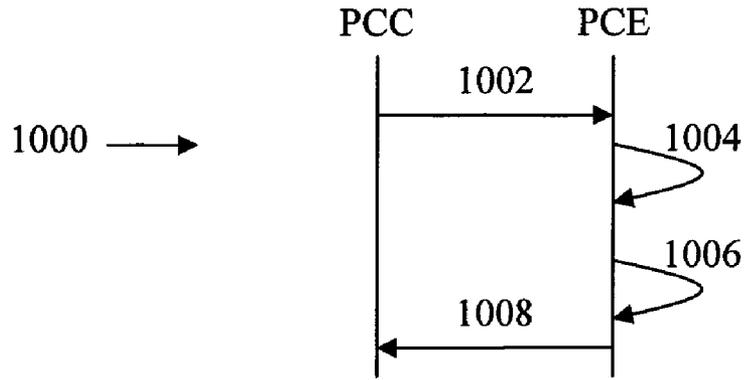


FIG. 10

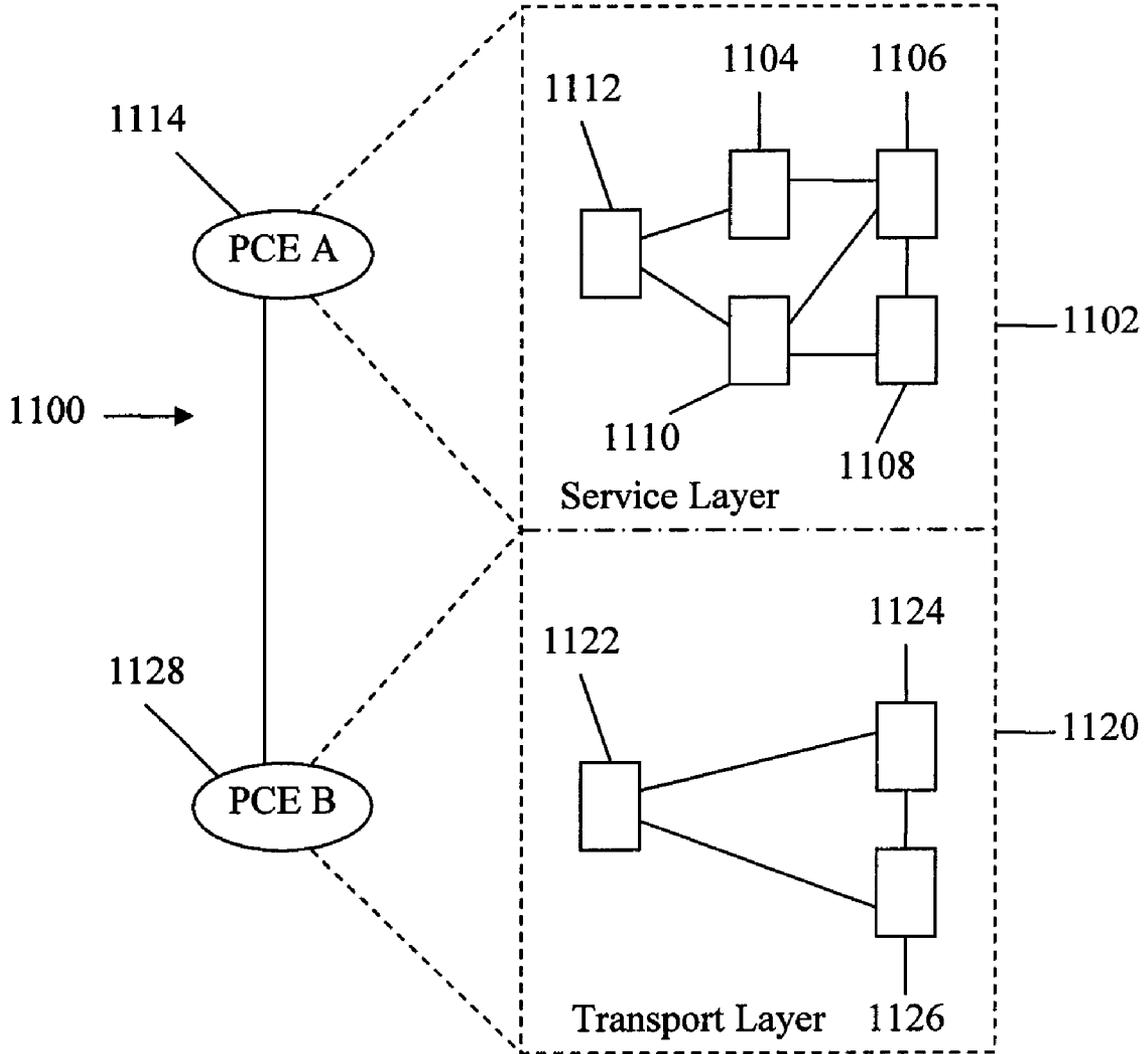


FIG. 11

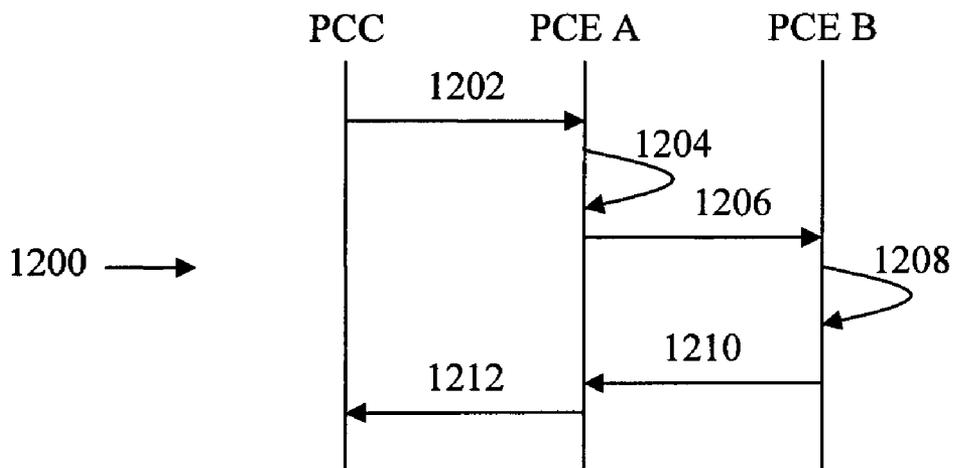


FIG. 12

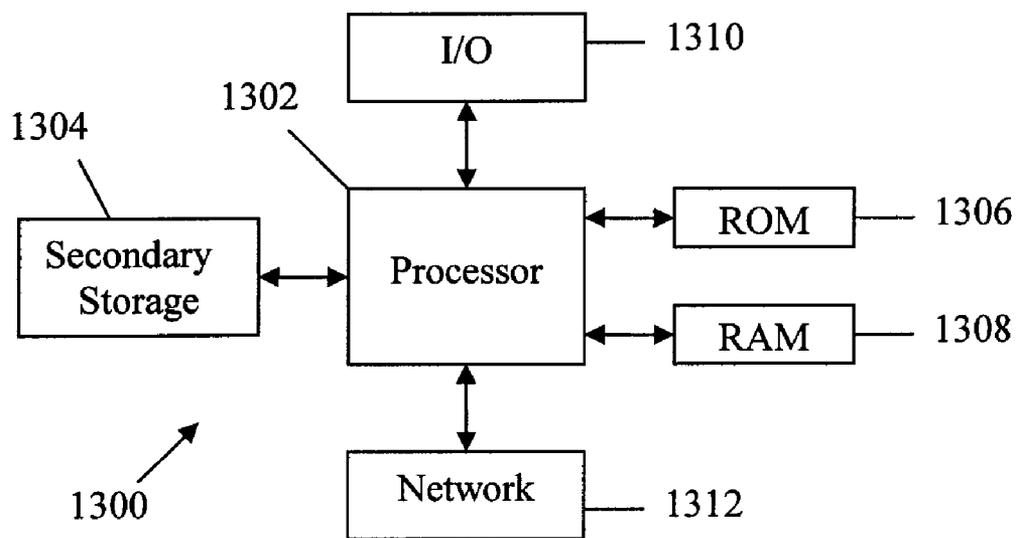


FIG. 13

OPTICAL IMPAIRMENT AWARE PATH COMPUTATION ARCHITECTURE IN PCE BASED NETWORK

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] The present application claims priority to U.S. Provisional Patent Application Ser. No. 60/895,283 filed Mar. 16, 2007 by Dunbar et al. and entitled "System for Optical Impairment Aware Path Computation Architecture in PCE Based Network", which is incorporated herein by reference as if reproduced in its entirety.

STATEMENT REGARDING FEDERALLY SPONSORED RESEARCH OR DEVELOPMENT

[0002] Not applicable.

REFERENCE TO A MICROFICHE APPENDIX

[0003] Not applicable.

BACKGROUND

[0004] Wavelength division multiplexing (WDM) is one technology that is envisioned to increase bandwidth capability and enable bidirectional communications in optical networks. In WDM networks, multiple data signals can be transmitted simultaneously between network elements (NEs) using a single fiber. Specifically, the individual signals may be assigned different transmission wavelengths so that they do not interfere or collide with each other. The path that the signal takes through the network is referred to as the lightpath. One type of WDM network, a wavelength switched optical network (WSON), seeks to switch the optical signals with fewer optical-electrical-optical (OEO) conversions along the lightpath, e.g. at the individual NEs, than existing optical networks.

[0005] One of the challenges in implementing WDM networks is the determination of the path for the various signals that are being transported through the network at any given time. Unlike traditional circuit-switched and connection-oriented packet-switched networks that merely have to determine a route for the data stream across the network, WDM networks are burdened with the additional constraint of having to ensure that the same wavelength is not simultaneously used by two signals over a single fiber. This constraint is compounded by the fact that WDM networks typically use specific optical bands comprising a finite number of usable optical wavelengths. Path computations can also be constrained due to other issues, such as excessive optical noise, along the lightpath.

SUMMARY

[0006] In one embodiment, the disclosure includes an apparatus comprising at least one processor configured to implement a method comprising receiving a path computation request comprising at least one path computation constraint, and determining whether there is a path through an optical network that satisfies the path computation constraints.

[0007] In another embodiment, the disclosure includes an apparatus configured to process a data structure comprising a flags field comprising at least one flag having one of an active state or an inactive state, wherein each flag is representative of an optical quality constraint.

[0008] These and other features will be more clearly understood from the following detailed description taken in conjunction with the accompanying drawings and claims.

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] For a more complete understanding of this disclosure, reference is now made to the following brief description, taken in connection with the accompanying drawings and detailed description, wherein like reference numerals represent like parts.

[0010] FIG. 1 is a schematic diagram of an embodiment of a WSON system.

[0011] FIG. 2 is a protocol diagram of an embodiment of the communications between a path computation element (PCE) and a path computation client (PCC).

[0012] FIG. 3 is a schematic diagram of an embodiment of an optical quality constraint (OQC) object.

[0013] FIG. 4 is a schematic diagram of an embodiment of a type, length, and value (TLV) sub-object.

[0014] FIG. 5 is a schematic diagram of another embodiment of a TLV sub-object.

[0015] FIG. 6 is a schematic diagram of an embodiment of a multi-domain network architecture.

[0016] FIG. 7 is a protocol diagram of an embodiment of the communications in a multi-domain network architecture.

[0017] FIG. 8 is a protocol diagram of another embodiment of the communications in a multi-domain network architecture.

[0018] FIG. 9 is a schematic diagram of an embodiment of a multi-layer network architecture.

[0019] FIG. 10 is a protocol diagram of an embodiment of the communications in a multi-layer network architecture.

[0020] FIG. 11 is a schematic diagram of another embodiment of a multi-layer network architecture.

[0021] FIG. 12 is a protocol diagram of another embodiment of the communications in a multi-layer network architecture.

[0022] FIG. 13 is a schematic diagram of one embodiment of a general-purpose computer system.

DETAILED DESCRIPTION

[0023] It should be understood at the outset that although an illustrative implementation of one or more embodiments are provided below, the disclosed systems and/or methods may be implemented using any number of techniques, whether currently known or in existence. The disclosure should in no way be limited to the illustrative implementations, drawings, and techniques illustrated below, including the exemplary designs and implementations illustrated and described herein, but may be modified within the scope of the appended claims along with their full scope of equivalents.

[0024] Disclosed herein is a method and system for including one or more path computation constraints in the PCE protocol (PCEP). Specifically, various types of path computation constraints are disclosed, as well as a communication protocol by which the PCE can consider the path computation constraints when performing its path computation for the PCC. Various alternatives are proposed whereby a plurality of PCEs each having authority over a different network communicate with each other and provide a unified path computation to the PCC. Alternatively, one or more PCEs may analyze a network at different layers, such as the service layer and the transport layer, to ensure that the path computation complies with the path computation constraints at each layer.

[0025] FIG. 1 illustrates one embodiment of a WSON system 100. The system 100 may comprise a WSON 110, a

control plane controller **120**, and a PCE **130**. The WSON **110**, control plane controller **120**, and PCE **130** may communicate with each other via optical, electrical, or wireless means. The WSON **110** may comprise a plurality of NEs **112** coupled to one another using optical fibers. In an embodiment, the optical fibers may also be considered NEs **112**. The optical signals may be transported through the WSON **110** over lightpaths that may pass through some of the NEs **112**. In addition, some of the NEs **112**, for example those at the ends of the WSON **110**, may be configured to convert between electrical signals from external sources and the optical signals used in the WSON **110**. Although four NEs **112** are shown in the WSON **110**, the WSON **110** may comprise any number of NEs **112**.

[0026] The WSON **110** may be any optical network that uses active or passive components to transport optical signals. The WSON **110** may implement WDM to transport the optical signals through the WSON **110**, and may comprise various types of optical components. For example, the WSON **110** may comprise optical cross connects (OXC), photonic cross connects (PXC), reconfigurable optical add/drop multiplexers (ROADM), wavelength selective switches (WSS), fixed optical add/drop multiplexers (FOADM), and so forth. The WSON **110** may be part of a long haul network, a metropolitan network, or an access network.

[0027] The control plane controller **120** may coordinate activities within the WSON **110**. Specifically, the control plane controller **120** may receive optical connection requests and provide lightpath signaling to the WSON **110** via Multi-protocol Label Switching Traffic Engineering (MPLS-TE) or Generalized Multi-Protocol Label Switching (GMPLS), thereby coordinating the NEs **112** such that data signals are routed through the WSON **110** with little or no contention. In addition, the control plane controller **120** may communicate with the PCE **130** using PCEP to provide the PCE **130** with information that may be used for the path computation, and/or receive the path computation from the PCE **130** and forward the path computation to the NEs **112**. The control plane controller **120** may be located in a component outside of the WSON **110**, such as an external server, or may be located in a component within the WSON **110**, such as a NE **112**.

[0028] The PCE **130** may perform all or part of the path computation for the WSON system **100**. Specifically, the PCE **130** may determine the routes of Label Switched Paths (LSPs) through the network. As such, the PCE **130** may receive the path computation constraints that may be used for the path computation from the control plane controller **120**, from the NEs **112**, or both. The PCE **130** may use the path computation constraints when computing the routes, e.g. lightpaths, for the optical signals. The path computation may include at least one route for each incoming signal and optionally at least one wavelength associated with each route. The PCE **130** may then send the path computation to the control plane controller **120** or directly to the NEs **112**. To assist the PCE **130** in this process, the PCE **130** may comprise a global traffic-engineering database (TED), a path computation information database, an optical performance monitor (OPM), a physical layer constraint (PLC) information database, or combinations thereof. The PCE **130** may be located in a component outside of the WSON **110**, such as an external server, or may be located in a component within the WSON **110**, such as a NE **112**.

[0029] The NEs **112** may be coupled to each other via optical fibers. The optical fibers may be used to establish optical links and transport the optical signals between the NEs **112**. The optical fibers may comprise standard single mode fibers (SMFs) as defined in ITU-T standard G.652, dispersion shifted SMFs as defined in ITU-T standard G.653,

cut-off shifted SMFs as defined in ITU-T standard G.654, non-zero dispersion shifted SMFs as defined in ITU-T standard G.655, wideband non-zero dispersion shifted SMFs as defined in ITU-T standard G.656, or combinations thereof. These fiber types may be differentiated by their optical impairment characteristics, such as attenuation, chromatic dispersion, polarization mode dispersion (PMD), four wave mixing, or combinations thereof. These effects may be dependent upon wavelength, channel spacing, input power level, or combinations thereof. The optical fibers may be used to transport WDM signals, such as coarse WDM (CWDM) signals as defined in ITU-T G.694.2 or dense WDM (DWDM) signals as defined in ITU-T G.694.1. All of the standards described herein are incorporated herein by reference.

[0030] In some embodiments, the PCE **130** may receive a path computation request from a PCC. The PCC may be any client application requesting a path computation to be performed by the PCE **130**. The PCC may also be any network component that makes such a request, such as the control plane controller **120**, or any NE **112**, such as a ROADM or a FOADM. Generally, the PCC communicates with the PCE **130** using PCEP, although other acceptable communications protocol may be used as well.

[0031] There may be many types of path computation constraints that can affect the path computation. In one embodiment, the path computation constraints include optical quality constraints. Examples of such include the optical signal-to-noise ratio (OSNR), amplifier spontaneous emission (ASE), PMD, polarization-dependent loss (PDL), coherent optical crosstalk, incoherent optical crosstalk, effective pass-band, gain non-uniformity, gain transients, and chromatic dispersion. In some embodiments, the path computation constraints may be classified as linear in that their effects are independent of the optical signal power and they affect the wavelengths individually. Alternatively, the path computation constraints may be classified as nonlinear in that their effects are dependent of the optical signal power, generate dispersion on a plurality of wavelength channels, and induce crosstalk between wavelength channels. Regardless, the path computation constraints are communicated to the PCE **130** so that the PCE **130** may consider them when computing a signal's path through the WSON **100**.

[0032] FIG. 2 illustrates an embodiment of a path computation communication method **200** between the PCC and the PCE. The method **200** may be implemented using any suitable protocol, such as the PCEP. The method **200** begins when the PCC sends a path computation request **202** to the PCE. The request **202** may comprise the OQC object described below. In some embodiments, the request **202** may comprise an indication of the need for a given source-destination (S-D) path through the network. At **204**, the PCE calculates a path through the network that meets the path computation constraints and any other network constraints. The PCE then sends a path computation reply **206** to the PCC. The reply **206** may comprise one or more paths through the network, which may be embodied as the OQC object and TLVs described below. In some embodiments, the reply **206** may comprise an indication whether the requested S-D path satisfies the path computation constraints. If the PCE is not able to calculate a path through the network that satisfies the path computation constraints, then the reply **206** may contain an error message that indicates that the PCE is not able to calculate a path through the network that satisfies the path computation constraints. Alternatively, if the PCE is not allowed to calculate a path through the network, then the reply **206** may contain an

error message that indicates that the PCE is not allowed to calculate a path through the network, for example, due to policy reasons.

[0033] FIG. 3 illustrates an embodiment of the OQC object 300 that may be included in the request and reply messages described herein. The OQC object 300 may comprise a reserved field 302, a flags field 304 comprising one or more flags, such as an N flag 306 and a P flag 308, and optionally one or more TLVs 310. The reserved field 302 may comprise the first about 10 bits of the OQC object 300 and may be reserved for uses by the PCC and PCE unrelated to the path computation constraints. The flags field 304 may comprise the subsequent about 22 bits of the OQC object 300 and may comprise one or more flags. Each flag may be related to one or more path computation constraints, such as the optical quality constraints. The flags may be any length and may be positioned anywhere and in any order in the flags field 304, but in an embodiment, each flag is about one bit in length and the flags are justified to the right of the flags field 304 in a predetermined order. In such a case, the presence of a one bit in a particular location may indicate that a particular path computation constraint exists or is applicable to the path computation request or reply. Alternatively, a zero bit in a particular location may indicate that the particular path computation constraint does not exist or is not applicable to the path computation request or reply. The presence of the one bit or the zero bit in the flags field 304 may also have differing meaning depending on whether the one bit or the zero bit is part of the request or the reply. The TLVs 310 may be associated with the flags in the flags field 304, and may provide more information regarding the path computation constraints.

[0034] For example and as shown in FIG. 3, the 31st bit in the OQC object 300 may be the N flag 306. The N flag 306 may be associated with the OSNR associated with the path. When the N flag 306 is set to zero in the request, it may indicate that the OSNR does not have to be considered in the path computation. When the N flag 306 is set to one in the request, it may indicate that the OSNR should be considered in the path computation. In such a case, the OQC object 300 may comprise a TLV 310 that contains further details regarding the OSNR, such as a value for the upper limit of the OSNR. When the N flag 306 is set to zero in the reply, it may indicate that the computed path does not comply with the OSNR constraint, further details of which may be included in a TLV 310. When the N flag 306 is set to one in the reply, it may indicate that the computed path complies with the OSNR constraint.

[0035] Similarly, the 32nd bit in the OQC object 300 may be the P flag 308. The P flag 308 may be associated with the PMD associated with the path. When the P flag 308 is set to zero in the request, it may indicate that the PMD does not have to be considered in the path computation. When the P flag 308 is set to one in the request, it may indicate that the PMD should be considered in the path computation. In such a case, the OQC object 300 may comprise a TLV 310 that contains further details regarding the PMD, such as a value for the upper limit of the PMD. When the P flag 308 is set to zero in the reply, it may indicate that the computed path does not comply with the PMD constraint, further details of which may be included in a TLV 310. When the P flag 308 is set to one in the reply, it may indicate that the computed path complies with the PMD constraint.

[0036] FIG. 4 illustrates one embodiment of a TLV 400 that may be associated with the N flag in the OQC object. The TLV 400 may comprise a type field 402, a length field 404, and a value field 406. The type field 402 may comprise the first

about 16 bits of the TLV 400 and may associate the TLV 400 with the N flag in the OQC object. The length field 404 may be the subsequent about 16 bits and may indicate the size of the value field 406 in bytes. The value field 406 may be any size, but in some embodiments is the subsequent about 32 bits on the TLV 400. The value field 406 may contain information related to the OSNR, such as the OSNR upper limit for the path or other OSNR impairment factor details.

[0037] FIG. 5 illustrates one embodiment of a TLV 500 that may be associated with the P flag in the OQC object. The TLV 500 may comprise a type field 502, a length field 504, and a value field 506. The type field 502 may comprise the first about 16 bits of the TLV 500 and may associate the TLV 500 with the P flag in the OQC object. The length field 504 may be the subsequent about 16 bits and may indicate the size of the value field 506 in bytes. The value field 506 may be any size, but in some embodiments is the subsequent about 32 bits on the TLV 500. The value field 506 may contain information related to the PMD, such as the PMD upper limit for the path or other PMD impairment factor details.

[0038] FIG. 6 illustrates an embodiment of a multi-domain network architecture 600 that may be used to illustrate the concepts described herein. The multi-domain network architecture 600 may comprise three networks: network A 602, network B 622, and network C 642, which may each be similar to the WSON described above. Network A 602 may comprise PCE A 614 and nodes 604, 606, 608, 610, and 612 (collectively 604-612). Similarly, Network B 622 may comprise PCE B 634 and nodes 624, 626, 628, 630, and 632 (collectively 624-632), and Network C 642 may comprise PCE C 654 and nodes 644, 646, 648, 650, and 652 (collectively 644-652). PCE A 614 can communicate with each of the nodes 604-612 and can calculate paths subject to path computation constraints within network A 602. Similarly, PCE B 634 can communicate with each of the nodes 624-632 and can calculate paths subject to path computation constraints within network B 622, and PCE C 654 can communicate with each of the nodes 644-652 and can calculate paths subject to path computation constraints within network C 642. In addition, the various components within the multi-domain network architecture 600 are coupled to and can communicate with each other as indicated by the solid lines in FIG. 6. It will be appreciated that while FIG. 6 is described in the context of a multi-domain network architecture, the concepts described herein are also applicable to situations where multiple PCEs are responsible for certain sections of a single network.

[0039] FIG. 7 illustrates one embodiment of a path computation communication method 700 between the PCC and the PCEs for the three networks illustrated in FIG. 6. Specifically, FIG. 7 illustrates the case where the path computation is performed by multiple PCEs, and each PCE may indicate whether the path computation constraints have been satisfied within the network or network portion of the path for which the PCE is responsible. The method 700 may be implemented using any suitable protocol, such as the PCEP. The method 700 begins when the PCC sends a path computation request 702 to one of the PCEs, such as PCE A. The request 702 may comprise the OQC object described above. In some embodiments, the request 702 may comprise an indication of the need for a path computation for a given S-D path in one or more of networks A, B, and C. At 704, PCE A calculates a path through network A that meets the path computation constraints and any other network constraints. For example, PCE A may determine that a path through nodes 612, 604, and 606 meets the path computation constraints for network A. Assuming that PCE A can calculate a path that meets the path

computation constraints for network A, PCE A forwards the request **706** to PCE B. At **708**, PCE B calculates a path through network B that meets the path computation constraints and any other network constraints. For example, PCE B may determine that a path through nodes **632**, **630**, and **626** meets the path computation constraints for network B. Assuming that PCE B can calculate a path that meets the path computation constraints for network B, PCE B forwards the request **710** to PCE C.

[0040] At **712**, PCE C calculates a path through network C that meets the path computation constraints and any other network constraints. In some cases, PCE C may be able to calculate a path that meets the path computation constraints for network C. For example, PCE C may determine that a path through nodes **652**, **650**, **648**, and **646** meets the path computation constraints for network C. In such cases, PCE C may add the path and any information related to the path to a reply **714**, and sends the reply **714** to PCE B. In other cases, PCE C may not be able to calculate a path that meets the path computation constraints for network C. For example, PCE C may determine that no path through network C meets the path computation constraints for network C. In such cases, PCE C may add one of the error messages described above and an optional explanation to the reply **714**, and may send the reply **714** to PCE B. In either case, the reply **714** may comprise the OQC object described above, which may include at least one path, at least one error message, and/or at least one indication whether the requested S-D path satisfies the path computation constraints. Upon receiving the reply **714**, PCE B may include the path through network B and any other related information in the reply **716** sent to PCE A. Upon receiving the reply **716**, PCE A may include the path through network A and any other related information in the reply **718** sent to the PCC. Thus, upon receipt of the reply **718**, the PCC knows whether a path through the networks A, B, and C exists, and if not, the location of and reason for the path computation failure. It will also be appreciated that in some embodiments the various PCEs may add one or more of their paths to the request as it is propagated through the networks, rather than adding their paths to the reply as it is returned through the networks.

[0041] FIG. **8** illustrates another embodiment of a path computation communication method **800** between the PCC and the PCEs for the three networks illustrated in FIG. **6**. Specifically, FIG. **8** illustrates the case where the path computation is performed by multiple PCEs, but the path computation fails at the first PCE, e.g. PCE A. The method **800** may be implemented using any suitable protocol, such as the PCEP. The method **800** begins when the PCC sends a path computation request **802** to one of the PCEs, such as PCE A. The request **802** may comprise the OQC object described above. In some embodiments, the request **802** may comprise an indication of the need for a path computation for a given S-D path in one or more of networks A, B, and C. At **804**, PCE A attempts to calculate a path through network A that meets the path computation constraints and any other network constraints. In this case, PCE A may not be able to calculate a path that meets the path computation constraints for network A. For example, PCE A may determine that no path through network A meets the path computation constraints for network A. Thus, PCE A may add one of the error messages described above and an optional explanation to the reply **806**, and sends the reply **806** to the PCC. Thus, the method **800** does not forward requests through the various networks when no path meeting the path computation constraints can be found in the first network. It will also be appreciated that the communications between the PCC and the PCEs may be any

combination of methods **700** and **800** and/or modified by some policy, for example where the request is forwarded to the next PCE unless more than one of the path computation constraints cannot be met.

[0042] FIG. **9** illustrates an embodiment of a multi-layer network architecture **900** that may be used to illustrate the concepts described herein. The multi-layer network architecture **900** may be similar to the WSON described above and may comprise a PCE **914** and two layers: a service layer **902** and a transport layer **920**. The service layer **902** may comprise nodes **904**, **906**, **908**, **910**, and **912** (collectively **904-912**), and the transport layer **920** may comprise nodes **922**, **924**, and **926** (collectively **922-926**). In some embodiments, at least some of the nodes **904-912** may be the same physical structure as at least some of the nodes **922-926**, but wherein the nodes **904-912** are separated from the nodes **922-926** by one or more logical partitions. The PCE **914** can communicate with each of the nodes **904-912** and **922-926** and can calculate paths subject to path computation constraints within the service layer **902** and the transport layer **920**. In addition, the various components within the multi-layer network architecture **900** have the topologies as shown in that they are coupled to and can communicate with each other as indicated by the solid lines in FIG. **9**. It will be appreciated that while FIG. **9** is described in the context of a single multi-layer network architecture, the concepts described herein are also applicable to situations where one or more PCEs are responsible for a plurality of path computations in a plurality of multi-layer networks.

[0043] FIG. **10** illustrates an embodiment of a path computation communication method **1000** between the PCC and the PCE for the network illustrated in FIG. **9**. Specifically, FIG. **10** illustrates the case where the path computation for multiple layers within a single network is performed by a single PCE. The method **1000** may be implemented using any suitable protocol, such as the PCEP. The method **1000** begins when the PCC sends a path computation request **1002** to the PCE. The request **1002** may comprise the OQC object described above. In some embodiments, the request **1002** may comprise an indication of the need for a path computation for a given S-D path in the service layer, the transport layer, or both. At **1004**, the PCE calculates a path through the service layer that meets the path computation constraints and any other network constraints. For example, the PCE may determine that a path through nodes **912**, **904**, and **906** meets the path computation constraints for the service layer. Assuming that the PCE can calculate a path that meets the path computation constraints for the service layer, the PCE verifies whether the path selected at **1004** meets the path computation constraints and any other network constraints of the transport layer at **1006**. In some cases, the path selected at **1004** meets the path computation constraints and any other network constraints for the transport layer. For example, the PCE may determine that a path through nodes **922** and **924** meets the path computation constraints for the transport layer. In such cases, the PCE adds the path and any information related to the path to a reply **1008**, and sends the reply **1008** to the PCC.

[0044] In other cases, the path selected at **1004** may not meet one or more of the path computation constraints and any other network constraints for the transport layer. For example, the PCE may determine that a path through nodes **922** and **924** does not meet the path computation constraints for transport layer. In such cases, the PCE will analyze alternative and perhaps less optimal paths through the service layer and the transport layer to determine if there are any paths that satisfy the path computation constraints for both the service layer and the transport layer. For example, the path through nodes

912, **910**, and **908** may satisfy the service layer path computation constraints while also satisfying the transport layer path computation constraints for the transport layer nodes **922**, **926**, and **924**. If such a path is found, the PCE may include the path and any other related information in the reply **1008** sent to the PCC. If such a path is not found, the PCE may include one of the error messages described above and an optional explanation in the reply **1008** sent to the PCC. In any event, the reply **1008** may comprise the OQC object described above, which may comprise at least one path, at least one error message, and/or at least one indication whether the requested S-D path satisfies the path computation constraints. Thus, upon receipt of the reply **1008**, the PCC knows whether a path through the service layer and transport layer exists, and if not, the location of and reason for the path computation failure.

[0045] FIG. 11 illustrates an embodiment of a multi-layer network architecture **1100** that may be used to illustrate the concepts described herein. The multi-layer network architecture **1100** may be similar to the WSON described above, and may comprise two PCEs, PCE A **1114** and PCE B **1128**, and two layers, a service layer **1102** and a transport layer **1120**. The service layer **1102** may comprise nodes **1104**, **1106**, **1108**, **1110**, and **1112** (collectively **1104-1112**), and transport layer **1120** may comprise nodes **1122**, **1124**, and **1126** (collectively **1122-1126**). In some embodiments, at least some of the nodes **1104-1112** may be the same physical structure as at least some of the nodes **1122-1126**, but wherein the nodes **1104-1112** are separated from the nodes **1122-1126** by one or more logical partitions. PCE A **1114** can communicate with each of the nodes **1104-1112** and can calculate paths subject to path computation constraints within the service layer **1102**, while PCE B **1128** can communicate with each of the nodes **1122-1126** and can calculate paths subject to path computation constraints within transport layer **1120**. In addition, the various components within the multi-layer network architecture **1100** have the topologies as shown in that they are coupled to and can communicate with each other as indicated by the solid lines in FIG. 11. It will be appreciated that while FIG. 11 is described in the context of a single multi-layer network architecture, the concepts described herein are also applicable to situations where a single PCE is responsible for one or more layers in multiple networks.

[0046] FIG. 12 illustrates another embodiment of a path computation communication method **1200** between the PCC and the two PCEs for the network illustrated in FIG. 11. Specifically, FIG. 12 illustrates the case where the path computation for each layer is performed by a separate PCE. The method **1200** may be implemented using any suitable protocol, such as the PCEP. The method **1200** begins when the PCC sends a path computation request **1202** to PCE A. The request **1202** may comprise the OQC object described above. In some embodiments, the request **1202** may comprise an indication of the need for a path computation for a given S-D path in the service layer, the transport layer, or both. At **1204**, PCE A calculates a path through the service layer that meets the path computation constraints and any other network constraints. For example, PCE A may determine that a path through nodes **1112**, **1104**, and **1106** meets the path computation constraints for the service layer. If there is not a path through the service layer that meets the path computation constraints and any other network constraints, then the PCE A can send a reply message to the PCC with one of the error messages described above. Assuming that PCE A can calculate a path that meets the path computation constraints and any other network constraints, PCE A includes the path in the request **1206** sent to PCE B. Upon receipt of the request **1206**, PCE B verifies whether the path selected at **1204** meets the path computation

constraints and any other network constraints for the transport layer at **1208**. In some cases, the path selected at **1204** meets the path computation constraints and any other network constraints for the transport layer. For example, PCE B may determine that a path through nodes **1122** and **1124** meets the path computation constraints for transport layer. In such cases, PCE B includes the path and any other related information in the reply **1210** sent to PCE A. PCE A then forwards the reply **1212** back to the PCC.

[0047] In other cases, the path selected at **1204** may not meet one or more of the path computation constraints and any other network constraints for the transport layer. For example, PCE B may determine that the path through nodes **1122** and **1124** does not meet the path computation constraints for transport layer. In such cases, PCE B will analyze alternative and perhaps less optimal paths through the transport layer to determine if there are any paths that satisfy the path computation constraints for the transport layer. For example, the path through nodes **1122**, **1126**, and **1124** may satisfy the transport layer path computation constraints. If such a path is not found, PCE B may include one of the error messages described above and an optional explanation in the reply **1210** sent to PCE A. If such a path is found, PCE B may include the path and any related information in the reply **1210** sent to PCE A. Upon receipt of a reply **1210** with an alternate path, PCE A verifies whether the alternate path selected at **1208** meets the path computation constraints and any other network constraints for the service layer. If the alternate path selected at **1208** meets the path computation constraints and any other network constraints for the service layer, PCE A may adopt the alternate path as the service layer path, and may include the alternate path in the reply **1212** sent to the PCC. If the path selected at **1208** does not meet the path computation constraints and any other network constraints of the service layer, then PCEs A and B can repeat the steps **1204**, **1206**, **1208**, and **1210** until either a path satisfying the path computation constraints for both the service layer and the transport layer is found, or it is determined that there is not a path that satisfies the path computation constraints for both the service layer and the transport layer. After such, the path or an error message is included in the reply **1212** sent to the PCC. In some embodiments, PCE A may add a plurality of paths to the request **1206** so that PCE B can verify each of the paths and indicate such in the reply **1210**. If there are multiple valid paths, PCE A can choose to keep one or more of these paths in the reply **1210**, and forward the reply **1212** to the PCC. In any event, the reply **1212** may comprise the OQC object described above, which may comprise at least one path, at least one error message, and/or at least one indication whether the requested S-D path satisfies the path computation constraints. Thus, upon receipt of the reply **1212**, the PCC knows whether a path through the service layer and transport layer exists, and if not, the location of and reason for the path computation failure.

[0048] The network components described above may be implemented on any general-purpose network component, such as a computer or network component with sufficient processing power, memory resources, and network throughput capability to handle the necessary workload placed upon it. FIG. 13 illustrates a typical, general-purpose network component suitable for implementing one or more embodiments of the components disclosed herein. The network component **1300** includes a processor **1302** (which may be referred to as a central processor unit or CPU) that is in communication with memory devices including secondary storage **1304**, read only memory (ROM) **1306**, random access memory (RAM) **1308**, input/output (I/O) devices **1310**, and network connec-

tivity devices **1312**. The processor may be implemented as one or more CPU chips, or may be part of one or more application specific integrated circuits (ASICs).

[0049] The secondary storage **1304** is typically comprised of one or more disk drives or tape drives and is used for non-volatile storage of data and as an over-flow data storage device if RAM **1308** is not large enough to hold all working data. Secondary storage **1304** may be used to store programs that are loaded into RAM **1308** when such programs are selected for execution. The ROM **1306** is used to store instructions and perhaps data that are read during program execution. ROM **1306** is a non-volatile memory device that typically has a small memory capacity relative to the larger memory capacity of secondary storage **1304**. The RAM **1308** is used to store volatile data and perhaps to store instructions. Access to both ROM **1306** and RAM **1308** is typically faster than to secondary storage **1304**.

[0050] While several embodiments have been provided in the present disclosure, it should be understood that the disclosed systems and methods might be embodied in many other specific forms without departing from the spirit or scope of the present disclosure. The present examples are to be considered as illustrative and not restrictive, and the intention is not to be limited to the details given herein. For example, the various elements or components may be combined or integrated in another system or certain features may be omitted, or not implemented.

[0051] In addition, techniques, systems, subsystems, and methods described and illustrated in the various embodiments as discrete or separate may be combined or integrated with other systems, modules, techniques, or methods without departing from the scope of the present disclosure. Other items shown or discussed as coupled or directly coupled or communicating with each other may be indirectly coupled or communicating through some interface, device, or intermediate component whether electrically, mechanically, or otherwise. Other examples of changes, substitutions, and alterations are ascertainable by one skilled in the art and could be made without departing from the spirit and scope disclosed herein.

What is claimed is:

1. An apparatus comprising:
at least one processor configured to implement a method comprising:
receiving a path computation request comprising at least one path computation constraint; and
determining whether there is a path through an optical network that satisfies the path computation constraints.
2. The apparatus of claim 1, wherein the method further comprises sending a reply comprising the path when there is at least one path through the optical network that satisfies the path computation constraints.
3. The apparatus of claim 2, wherein the method further comprises sending a reply comprising an error message when there is not a path through the optical network that satisfies the path computation constraints.
4. The apparatus of claim 3, wherein the method further comprises forwarding the request to a path computation element (PCE) associated with another optical network domain.
5. The apparatus of claim 4, wherein the request forwarded to the PCE does not comprise the path.
6. The apparatus of claim 4, wherein the reply comprises at least one path for the second optical network.

7. The apparatus of claim 3, wherein the optical network comprises a plurality of layers, wherein determining whether there is a path through the optical network comprises analyzing less than all of the layers, and wherein the method further comprises forwarding the request to a path computation element (PCE) associated with at least some of the unanalyzed layers.

8. The apparatus of claim 7, wherein the request forwarded to the PCE comprises the path that satisfies the path computation constraints for the analyzed layers.

9. The apparatus of claim 7, wherein there is not a path through the optical network that satisfies the path computation constraints when there is not a path that satisfies the path computation constraints for all of the layers.

10. The apparatus of claim 3, wherein the optical network comprises a plurality of layers, wherein determining whether there is a path through the optical network comprises analyzing all of the layers.

11. The apparatus of claim 10, wherein the optical network comprises a plurality of network elements (NEs), and wherein the topology between the NEs is not identical in all of the layers.

12. The apparatus of claim 1, wherein the path computation constraint are selected from the group consisting of: optical signal-to-noise ratio (OSNR), amplifier spontaneous emission (ASE), polarization mode dispersion (PMD), polarization-dependent loss (PDL), coherent optical crosstalk, incoherent optical crosstalk, effective pass-band, gain non-uniformity, gain transients, chromatic dispersion, and combinations thereof.

13. The apparatus of claim 1, wherein the path is present in an optical quality constraint object.

14. An apparatus configured to process a data structure comprising:

a flags field comprising at least one flag having one of an active state or an inactive state, wherein each flag is representative of an optical quality constraint.

15. The apparatus of claim 14, wherein the data structure further comprises a type-length-value (TLV) data structure for at least some of the flags having the active state.

16. The apparatus of claim 14, wherein each flag is about one bit in length, is in the active state when the bit is a one, and is in the inactive state when the bit is a zero.

17. The apparatus of claim 14, wherein the last flag in the flags field is a polarization mode dispersion (PMD) flag.

18. The apparatus of claim 14, wherein the second-to-last flag in the flags field is an optical signal-to-noise ratio (OSNR) flag.

19. The apparatus of claim 14, wherein the flags having active states have different meanings depending on whether the flag is in a request message or a reply message.

20. The apparatus of claim 14, wherein the optical quality constraints are selected from the group consisting of: optical signal-to-noise ratio (OSNR), amplifier spontaneous emission (ASE), polarization mode dispersion (PMD), polarization-dependent loss (PDL), coherent optical crosstalk, incoherent optical crosstalk, effective pass-band, gain non-uniformity, gain transients, chromatic dispersion, and combinations thereof.

* * * * *