

(19) World Intellectual Property Organization  
International Bureau



(43) International Publication Date  
1 October 2009 (01.10.2009)

(10) International Publication Number  
**WO 2009/118050 A1**

PCT

(51) International Patent Classification:  
**H04L 12/56** (2006.01)

(21) International Application Number:  
PCT/EP2008/053772

(22) International Filing Date:  
28 March 2008 (28.03.2008)

(25) Filing Language: English

(26) Publication Language: English

(71) Applicant (for all designated States except US): **TELEFONAKTIEBOLAGET LM ERICSSON (PUBL)** [SE/SE]; S-164 83 Stockholm (SE).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **IOVANNA, Paola** [IT/IT]; Via E. Ferri 3, 1-00046 Rome (IT). **ZEMA, Cristiano** [IT/IT]; Via Checco Durante 15 A/6, 1-00477 Rome (IT).

(74) Agent: **LIND, Robert**; Marks & Clerk LLP, 4220 Nash Court, Oxford Business Park South, Oxford Oxfordshire OX4 2RU (GB).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MT, NL, NO, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

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

(54) Title: END-TO-END INTER-DOMAIN ROUTING

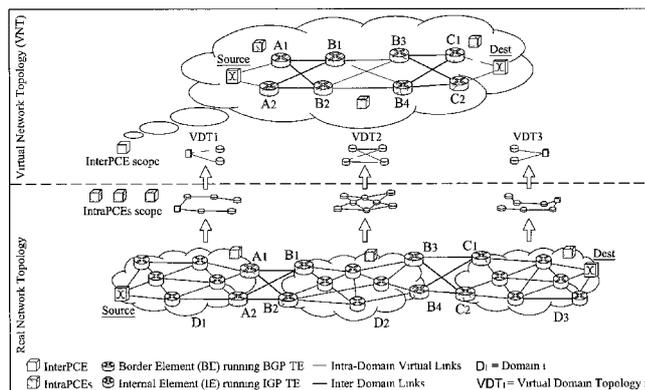


Figure 1

(57) Abstract: According to a first aspect of the present invention there is provided a method of constructing an end-to-end routing path across multiple packet switched domains. The method comprises generating a domains topology identifying all possible domains, the border elements of those domains and the inter-domain links. When a request is received for computation of a path from a source entity in a first domain to a destination entity in a second domain, all domains and inter-domain links that cannot or are unlikely to feature in the requested path are removed from the domains topology. Candidate intra-domain paths for each remaining domain are then computed, with each candidate intra-domain path being mapped to a logical link defined by its two ends, at least one end being a border element and the other being a border element or the source or destination entity. Attributes for each logical link are determined, together with attributes of the remaining inter-domain links. A virtual network topology is generated by combining the logical links and their associated attributes with the remaining inter-domain links and their associated attributes. This allows a best end-to-end route to be determined by analysing the virtual network topology.



WO 2009/118050 A1

## END-TO-END INTER-DOMAIN ROUTING

Technical Field

5 The present invention relates to end-to-end traffic engineered multi-domain routing in connection-oriented packet based architectures. It is applicable in particular, though not necessarily, to Multi-Protocol Label Switching (MPLS) and its extensions (i.e. GMPLS, T-MPLS) and Traffic Engineered Provider Backbone Bridging (PBB-TE).

10 Background

The Internet is a conglomeration of Autonomous Systems (AS) or domains that define the administrative authority and the routing policies of different organizations. These domains consist of routers that run Interior Gateway Protocols (IGPs) such as Routing  
15 Information Protocol (RIP), Enhanced Interior Gateway Routing Protocol (EIGRP), Open Shortest Path First (OSPF), and Intermediate System-to-Intermediate System (IS-IS) within their boundaries. Neighbouring domains are interconnected via an Exterior Gateway Protocol (EGP); the current Internet standard EGP is the Border Gateway Protocol Version 4 (BGP-4) defined in RFC 4271.

20

Exterior routing protocols were created to control the expansion of routing tables and to provide a more structured view of the Internet by segregating it into separate administrations, or domains, each with their own independent routing policies and unique IGPs.

25

These routing protocols define how routers determine their 'map' of the network and from which they can compute the shortest path to a destination, allowing routing to be a largely automatic process. However, the shortest path is not always the fastest or the best. Traffic Engineering (TE) is the process where data is routed through the  
30 network according to the availability of resources and the current and expected traffic. The required quality of service (QoS) can also be factored into this process.

Traffic Engineering may be under the control of operators whereby they monitor the state of the network and route the traffic, or provision additional resources, to compensate for problems as they arise. Alternatively, Traffic Engineering may be automated. Traffic Engineering helps the network provider make the best use of available resources, spreading the load over the layer 2 links, and allowing some links to be reserved for certain classes of traffic or for particular customers.

New technologies such as Traffic Engineered Provider Backbone Bridging (PBB-TE), and more standardized ones, such as Multi-Protocol Label Switching (MPLS) and its extensions (i.e. GMPLS, T-MPLS), provide efficient TE solutions within a single domain (i.e. intra-domain) thanks to their connection oriented nature. However, in order to support all available services end-to-end, future packet switched network architectures need to guarantee both suitable QoS and efficient use of resources between networks. This requires routing solutions that can apply Traffic Engineering for the entire path, end-to-end, both within and between domains.

The traditional approach for inter-domain TE routing is based on BGP-TE (*see IETF draft-fedyk-bgp-te-attribute-03 - "Traffic Engineering Attribute"*). However, this method is only capable of applying TE constraints to computation of the inter-domain paths and to the intra-domain paths of the source and destination domains. BGP-TE does not consider detailed information relating to the intra-domain paths within any transit domains, which must be traversed to reach the destination. In addition, in a BGP-based approach, complex policy constraints can be configured, but are mostly targeted to peering agreements and economical or administrative choices. TE constraints such as maximum overload on a link and congestion occurring on a transit link are still not supported.

US09/981,138 discloses a system for end-to-end path computation taking into consideration detailed network resource information in the source and destination domains, while transit domains are selected according to traditional BGP-TE. BGP-TE floods the network with aggregate information in terms of routes and TE weights and, as a result, the lack of detailed information on transit domains does not allow for

efficient end-to-end TE path computation.

The computation of an end-to-end path taking into account several constraints (such as QoS, bandwidth, priority, protection, etc.) in a multi-domain scenario is a very difficult issue. Detailed information is required from each domain in order to compute efficient paths using TE; the more detailed this information the more path computation meets TE requirements. This has obvious problems of scalability when expanding to incorporate more domains. In addition, a lot of intra-domain information (e.g. state of links, topology, administrative policies, etc.) is preferably confidential such that the owners/administrators of a domain would not want to share such details with other domains or administrative entities external to the domain. For these reasons the use of sophisticated intra-domain TE strategies is limited when applied to the computation of end-to-end inter-domain paths.

Standardization bodies, such as the IETF, are working on the definition of communication protocols and related architectures to deal with the multi-domain network scenario. These are based on a client-server architecture where a generic entity named the Path Computation Client (PCC) represents the client performing the path request, while the Path Computation Element (PCE) is the entity receiving the request and performing the path computation (*see IETF RFC 4655 - "A Path Computation Element (PCE)-based architecture"*). The standards define a communication protocol and the requirements to be satisfied (such as protection type, priority, colour, etc.), while the solution to satisfy such requirements in a multi-domain scenario is outside of the scope of the standardization work.

25

### Summary

According to a first aspect of the present invention there is provided a method of constructing an end-to-end routing path across multiple packet switched domains supporting one or more connection oriented technologies. The method comprises generating a domains topology identifying all possible domains, the border elements of those domains and the inter-domain links. When a request is received for

30

computation of a path from a source entity in a first domain to a destination entity in a second domain, all domains and inter-domain links that cannot or are unlikely to feature in the requested path are removed from the domains topology.

5 Candidate intra-domain paths for each remaining domain are then computed, with each candidate intra-domain path being mapped to a logical link defined by its two ends, at least one end being a border element and the other being a border element or the source or destination entity. Attributes for each logical link are determined, together with attributes of the remaining inter-domain links.

10

A virtual network topology is generated by combining the logical links and their associated attributes with the remaining inter-domain links and their associated attributes. This allows a best end-to-end route to be determined by analysing the virtual network topology.

15

In a particular implementation, the step of generating a domains topology further comprises, at an intra-domain path computation entity in each domain, generating inter-domain data including details of the domain border elements and their inter-domain links and sending said inter-domain data to an end-to-end path computation  
20 entity. At the end-to-end path computation entity, said inter-domain data is used to generate the domains topology.

The method may involve, at the end-to-end path computation entity, generating a virtual Domains Topology request for each domain including each of the border  
25 elements of each domain that may feature in the end-to-end path, any attributes required by said end-to-end path and, for the source and destination domains, the source element and the destination element respectively and sending said virtual Domains Topology request to the intra-domain path computation entity of the source domain, the destination domain and of all possible transit domains substantially  
30 simultaneously. Upon receipt of a virtual Domains Topology request at each intra-domain path computation entity, the candidate intra-domain paths for the request are computed.

The step of determining the attributes for each logical link may comprise determining a cost parameter that represents the availability/suitability of the logical link with respect to the other logical links computed within that domain.

5

An intra-domain path computation entity may send a virtual Domains Topology reply to the end-to-end path computation entity, the reply including the identified logical links, the border elements that feature in the logical links and the associated summarised attributes and, for the source and destination domains, also the source and destination element respectively. Upon receipt of the virtual Domains Topology reply at the end-to-end path computation entity, the logical links and their associated attributes may be combined with the remaining inter-domain links and their associated attributes. An attribute for a link may be given as a percentage of the total network resources that are available within the domain for that link.

15

The step of generating a virtual Domains Topology request for each domain may comprise including within the virtual Domains Topology request, an identification of the attribute(s) to be determined by the intra-domain path computation entity. The virtual Domains Topology request may further include an identification of the routing strategy or algorithm to be used by the intra-domain path computation entity to compute the candidate intra-domain paths.

In order to make the system dynamic, an intra-domain path computation entity may monitor its inter-domain data and send updates of the inter-domain data to the end-to-end path computation entity. Updates of inter-domain data may be sent to the end-to-end path computation entity when changes in the inter-domain data cross some defined threshold(s).

According to a second aspect of the present invention there is provided a method of facilitating construction of an end-to-end routing path across multiple domains supporting packet based connection oriented technologies. The method comprises receiving inter-domain data from border elements of said domain, said inter-domain

30

data including information relating to each border element of said domain and the links between those border elements of the domain and border elements of other domains. Said inter-domain data is then sent to an end-to-end path computation entity. Subsequently, a virtual Domains Topology request is received from said end-to-end path computation entity, said virtual Domains Topology request including the border elements of that domain that may feature in the end-to-end path and any attributes required by said path. The candidate intra-domain paths are computed, and attributes determined for each of the computed candidate intra-domain paths. Each candidate intra-domain path is then mapped to a logical link defined by its two ends, at least one end being a border element and the other being a border element or a source or destination entity. Finally, a virtual Domains Topology reply is sent to the end-to-end path computation entity, including the logical links and their associated attributes.

Considering the process of determining attributes, consider the example where the path computation criteria of the E2E-PCE is the minimum hop. The cost associated with the candidate path is therefore the number of hops to go from the head element of the path to the tail element (e.g. ten). That number is therefore the cost to be associated with the corresponding logical link (i.e. the cost of traversing the domain using that logical link is the same cost as that involving a path of ten hops).

According to a third aspect of the present invention there is provided a method of constructing an end-to-end routing path across multiple domains supporting packet based connection oriented technologies. The method comprises receiving inter-domain data from intra-domain path computation entities, said inter-domain data including information relating to each border element of each domain and the links between those border elements of each domain and border elements of other domains. A domains topology is generated, identifying all possible domains, the border elements of those domains and the inter-domain links. An end-to-end path request is received from a source entity, said end-to-end path request including details of the desired source and destination of a path and any attributes required for that path.

All domains and inter-domain links that cannot or are unlikely to feature in the path are removed from the domains topology, and a virtual Domains Topology request(s) is generated, including each of the border elements of each domain that may feature in the path and any attributes required by said path. The virtual Domains Topology request(s) is sent to an intra-domain path computation entity of the source domain, the destination domain and of all possible transit domains simultaneously.

A virtual Domains Topology reply is received from each said domain, the virtual Domains Topology reply including logical links representing each candidate intra-domain path and its associated attributes. The attributes of each of the remaining inter-domain links are determined, and a virtual network topology generated by combining the logical links and their associated attributes with the remaining inter-domain links and their associated attributes. Finally, a best end-to-end route is determined by analysing said virtual network topology.

15

#### Brief Description of the Drawings

Figure 1 illustrates a simplified diagram of a PCE-based network topology and its translation to a virtual network topology;

20 Figure 2 illustrates a simplified diagram of a PCE-based network comprised of seven domains;

Figure 3 illustrates domains D4 and D5 of the network of Figure 2, their InterLinks, and a type of information collected by an InterPCE in order to generate a Domains Topology (DT);

25 Figure 4 illustrates the Domains Topology of the network of Figure 2;

Figure 5 illustrates the pruning of the Domains Topology of Figure 4 according to the Path Request of Table 2, and the resultant Pruned Domains Topology (PDT);

Figure 6a illustrates the candidate paths in the intra and inter domains of the network of Figure 2;

30 Figure 6b illustrates the Virtual Links of the network of Figure 2 combined into a Virtual Domains Topology;

Figure 6c illustrates the Virtual Network Topology of the network of Figure 2;

Figure 7 is a flow diagram illustrating the process of creating and updating a Domains Topology;

Figure 8 is a flow diagram illustrating the process of computing an end-to-end best path;

5 Figure 9 illustrates schematically a source entity according to an embodiment of the invention;

Figure 10 illustrates schematically an intra-domain path computation entity according to an embodiment of the invention; and

10 Figure 11 illustrates schematically an end-to-end path computation entity according to an embodiment of the invention.

### Detailed Description

An end-to-end route determination process for packet based connection oriented  
15 technologies is presented here which makes use of a heterogeneous virtual topology including all information from both the intra-domain and inter-domain elements of the route. The topology of the network is separated into two layers, the intra-domain layer and the inter-domain layer. The inter-domain layer relates to a heterogeneous virtual topology that is comprised of both real and virtual elements. Its elements are  
20 the Border Elements (BEs) of the source domain, destination domain and any transit domains, the physical inter-domain links (InterLinks) connecting different domains, the source node of the source domain, the destination node of the destination domain, the Logical Links (LLs) connecting all the BEs of each transit domain, and the logical links connecting the BEs of the source/destination domain to the source/destination  
25 node. The intra-domain layer relates to the real topology of each domain. It is comprised of the source domain, the destination domain and each of the transit domains. Each domain is further comprised of a number of interconnected intra-domain nodes, including the source and destination nodes in the source and destination domains respectively, and the domain BEs.

30

For the purpose of the following discussion, the network architecture is PCE-based, each domain having at least one intra-domain PCE (Intra-PCE) responsible for

calculating the intra-domain routing and there being at least one end-to-end PCE (E2E-PCE) responsible for calculating the entire end-to-end route, combining intra-domain and inter-domain routing.

5 Figure 1 shows a simplified diagram of a PCE-based network topology and its translation to a virtual network topology. The complete topology of the network is shown below the dashed line and features a source within a source domain D1, a destination within a destination domain D3 and one transit domain D2. Each of the domains includes a number of interconnected nodes or Internal Elements (IEs),  
10 Border Elements and an Intra-PCE. The Border Elements of domain D1 are connected to Border Elements of the transit domain D2. Other Border Elements of domain D2 are connected to the Border Elements of the destination domain D3.

Each Intra-PCE is aware of the topology of its domain network and is capable of  
15 calculating TE routing within the domain according to known TE routing methods (e.g. according to intra-domain path selection criteria, metrics, policies, rules and target functions). The Intra-PCE uses these methods to determine the 'best' candidate paths between intra-domain nodes and the Border Elements or, when the domain is a transit domain to be traversed, between separate Border Elements of the domain.  
20 Candidate paths computed by different domains are defined as "best" according to different criteria, so the local cost parameter associated to each best path within each different domain has a different meaning.

These candidate paths are then translated into Logical Links (LLs), defining the path  
25 through a domain by its two end points. These Logical Links are combined with the inter-domain links, with source and destination nodes and with BEs to form the overall virtual network topology as shown above the dashed line in Figure 1.

In addition, the Intra-PCE translates the costs of candidate paths into one or more  
30 costs for each of the Logical Links. These costs are chosen as the parameters (or a summary of them) needed by the E2E-PCE to compute the End-to-End Best Path (E2E-BP) according to its TE routing methods (e.g. according to inter-domain path

selection criteria, metrics, policies, rules and target functions). This allows the costs of LLs within different domains to be consistent, and allows them to be consistent with inter-domain link costs.

5 Traffic balancing, which is a challenging issue, can be achieved by choosing a LL cost parameter that represents the availability/suitability of a LL with respect to the other LLs computed within that domain. The parameters of these costs are determined according to the desired routing objectives. As an example, the attributes of a logical link could be given as a percentage of the total network resources within  
10 the domain that are available for the candidate path corresponding to that logical link, termed "Equivalent Available Resources" (EAR). Such a parameter would be easily comparable across all domains. The E2E-PCE is then able to use the associated attributes of the links to compute an appropriate end-to-end path across multiple domains.

15

By way of example, consider the Virtual Domains Topology VDT2 in Figure 6b. This consists of four LLs (LL1, LL2, LL3, LL4). Now assume further that the Intra-PCE of the domain (D2) assigns EAR values to LLs according to the formula:

$$\text{EAR} = [\text{AABW} + (100 - \text{ANL})] / 2$$

20 where AABW is the Average Available Bandwidth of all links and ANL is the Average Node Load of all nodes along the candidate path corresponding to a logical link. Table 2 below illustrates example percentage AABW and ANL values for the four LLs. In such a situation, if all other conditions are equal, the E2E-PCE will choose LL1 to traverse D2, because LL1 has a higher EAR, so traffic will be better  
25 balanced using that LL.

The application of the solution requires two phases; the creation of a Domains Topology (DT), within which the domains are seen as nodes and their border elements as node ports, and the computation of the End-to-End Best Path (E2E-BP).

30

#### *Phase 1: Domains Topology Generation*

The generation of the Domains Topology is independent of the specific path request

as its objective is to give the E2E-PCE a global, inter-domain view of the network. This phase is therefore carried out off-line.

With reference to the flow diagram of Figure 7, the Phase 1 process commences at step 100. Initially, step 101, each Intra-PCE sends to the E2E-PCE a set of information concerning its domain (Domain Info) and a set of information concerning the Border Elements of the domain and their inter-domain links (InterLinks Info). The Domain Info can be parameters such as the administrative attributes to go from one domain to another (according to Service Level Agreements and/or Service Level Specifications). The InterLink information can be parameters concerning both the InterLinks and Border Elements such as link bandwidth or load on the Border Elements. Using this information the E2E-PCE creates the Domains Topology, step 102.

Each Intra-PCE monitors, step 103, the intra-domain network resources (according to IGP-TE/EGP-TE e.g. OSPF-TE/BGP-TE) and updates its Domain Info and InterLink Info accordingly. To reduce the traffic on the network the Intra-PCEs will only notify the E2E-PCE when these changes are significant, as defined by exceeding some threshold, steps 104 and 105. The DT can then be updated if the E2E-PCE receives new parameters from an Intra-PCE. The process ends at step 106.

#### *Phase 2: End to End Best Path Computation*

This process is illustrated in the flow diagram of Figure 8, and commences at step 200. Consider a source node wishing to send data to a destination node. The source node generates a message, or Path Request, containing details of the destination and any constraints for the path. This Path Request is sent either directly to the E2E-PCE or via the Intra-PCE of the source domain.

After receipt of a Path Request from a source node, the E2E-PCE uses the Domains Topology information and any constraints specified within the Path Request to 'prune' those domains and/or InterLinks that cannot feature in the path, step 201.

The E2E-PCE then simultaneously sends Path Requests to the Intra-PCEs of the remaining domains, step 202, requesting them to provide the intra-domain Logical Links that may be suitable to the specific path request, and the associated attributes of each link.

5

Each Intra-PCE calculates the possible candidate paths for the specific path request using any appropriate TE routing method(s), step 203. The source domain Intra-PCE computes the candidate paths connecting the source with the Border Elements of the domain, the transit domain Intra-PCEs compute the candidate paths interconnecting its Border Elements, and the destination domain Intra-PCE computes the candidate paths connecting its Border Elements to the destination. If the path request included certain constraints (such as colour as defined in RFC2702, disjointness or path separation, protection, etc) the Intra-PCE also determines the suitability of each path for meeting these requirements. Each path is then translated into a Logical Link and provided, along with the appropriate attributes parameter, to the E2E-PCE, step 204. The E2E-PCE then applies any suitable TE routing method to a combination of these Logical Links and the InterLinks to select the best end-to-end path, considering the associated attributes of all of the links, step 206.

10

15

20

25

30

The end-to-end path provided by the E2E-PCE will be optimum, according to a TE routing strategy, if both the Intra-PCE and E2E-PCE apply the same TE routing strategy. However, if the inter-domain and intra-domain routing strategies do differ, then the Intra-PCE can provide attribute parameters suitable for use with the TE routing algorithm used by the E2E-PCE. This could be implemented by the incorporation of a request/reply mechanism wherein the Path Request sent from the E2E-PCE to the Intra-PCE could specify the routing strategy that should be used by the Intra-PCE and/or the parameters of the attributes of the logical links required by the E2E-PCE. This makes the present solution extremely flexible, enabling the two layers to use independent routing strategies whilst maintaining efficient end-to-end routing.

The solution will now be further discussed by way of example, with reference again to

the steps of the flow diagram of Figure 8. Figure 2 shows a simplified diagram of a PCE-based network comprising seven domains. For the purpose of the following discussion, it is assumed the entity that sends the Path Request (PathReq) is a Path Computation Client (PCC) that can be any element of any domain and therefore has  
5 been represented as not belonging to any domain. Domains D1 to D7 each have a number of Internal Elements (not shown for simplicity), Border Elements and an Intra-PCE. Once again, the Border Elements of the domains are interconnected by a number of InterLinks.

10 Figure 3 illustrates domains D4 and D5 of Figure 2 and their InterLinks, Lx and Ly. The BEs of D4 and D5, that is E2 and F1 respectively, provide the information relating to InterLink Lx to their Intra-PCEs, which in turn provide this information and the Domain Info to the E2E-PCE. All Intra-PCEs send their Domain Info and InterLink Info parameters to the E2E-PCE, via the PCE communication Protocol  
15 (PCEP), with extensions where needed. Using this information the E2E-PCE creates the Domains Topology within which the domains are seen as nodes and their border elements as node ports, as shown in Figure 4.

A Path Request is then generated by a PCC (which could be the source node) and sent  
20 to the E2E-PCE. The Path Request must specify a Source Address (SA) and a Destination Address (DA) that are used to discover the Source Domain (SD) and the Destination Domain (DD). The Path Request also contains further fields for specifying constraints for the path. The following list is an example of some of the constraints that could be specified within the Path Request:

- 25
- Service Type (e.g. unicast, multicast or broadcast)
  - LSP Type (e.g. unidirectional or bidirectional, E-LSP or L-LSP)
  - Bandwidth Upstream
  - Bandwidth Downstream
  - Max\_Node\_Load (i.e. the maximum allowed percentage of node load along  
30 E2E path)
  - QoS (e.g. classes 0 to 7 of DiffServ)
  - Protection (e.g. 1+1, 1:1, 1:n)

- Straight or loose Link Disjointness along E2E path
  - Straight or loose Node Disjointness along E2E path
  - Straight or loose Domain Disjointness
  - Straight or loose Inter-domain Link Disjointness
- 5
- Objective Function (e.g. minimize Node Load, maximize the residual Bandwidth along E2E path, perform traffic balancing, etc.)
  - DomExcl (i.e. Explicit Domains Exclusion from the requested path)

Table 2 below shows an example Path Request with a set of constraints that must be  
 10 satisfied by the computed path. Applying the Path Request of Table 2 to the topology  
 of Figure 2 leads to the creation of the Pruned Domains Topology, the E2E-PCE  
 performing the following pruning:

- The DomExcl field specifies D6 as an excluded domain and therefore D6 and its related InterLinks are pruned.
- 15
- D7 is a leaf-connected domain that is neither the SD nor the DD. No path from the SD to the DD will traverse this domain and therefore D7 and its related InterLinks are pruned.
  - Assuming InterLink Lx is congested, either it has insufficient bandwidth or one or both of its BEs has a load exceeding that specified in the PathReq, then  
 20 Lx is also pruned.

Figure 5 illustrates this pruning of the network topology of Figure 2 according to the  
 Path Request of Table 2, as represented by those circled domains and crossed-out  
 InterLinks, and the resultant Pruned Domains Topology (PDT)

25 Following creation of the Pruned Domains Topology the E2E-PCE sends a Virtual  
 Domains Topology Computation Request (VDTCReq), with the same constraints as  
 the Path Request, to the Intra-PCEs of the Source Domain, the Destination Domain  
 and all possible Transit Domains (TDs). Different VDTCReqs are sent to the Source  
 Domain, to the Destination Domain and to all the possible Transit Domains because  
 30 these domains have to provide different kinds of Virtual Domain Topologies (VDTs)  
 (i.e. made by different elements), as explained above.

The Intra-PCEs compute all candidate paths that meet the VDTCReq according to local, intra-domain path selection criteria, metrics, policies, rules and target functions. As previously discussed, these criteria used by the Intra-PCE can differ to those applied by the E2E-PCE. In such circumstances, once the candidate paths that satisfy the Path Request have been computed, the attributes associated with each path are determined in accordance with those parameters needed by the E2E-PCE to compute the End-to-End Best Path (E2E-BP). The Intra-PCE then forms a Virtual Domains Topology (VDT) of the domain, including a set of Logical Links (LLs) representing each possible candidate path connection and their associated attributes. A Virtual Domains Topology for the Source Domain provides Logical Links for the best paths connecting the Source Element (SE) to each of those Border Elements of the domain that could possibly feature in the route. A Virtual Domains Topology for a Transit Domain provides the Logical Links for the best paths between those Border Elements of the domain that could possibly feature in the route. A Virtual Domains Topology for the Destination Domain provides the Logical Links for the best paths connecting each of the Border Elements of the domain that could possibly feature in the route to the Destination Element (DE).

Figure 6a illustrates the candidate paths (i.e. the best paths among all the possible ones connecting the same couple of elements) in the intra and inter domains following domain pruning of the network topology of Figure 2. The candidate paths through each domain are shown separately at the top of Figure 6a. Each of these paths is then translated into a Logical Link and all of them are combined into a Virtual Domains Topology as shown at the bottom of Figure 6b.

Once the Virtual Domains Topology is created, each IntraPCE sends a Virtual Domains Topology Computation Reply (VDTCRep), carrying the Virtual Domains Topology, to the E2E-PCE.

Following receipt of the VDTCRep from all of the requested domains the E2E-PCE can create a Virtual Network Topology (VNT), in which nodes are the Source Element, Destination Element and Border Elements of all domains of the Pruned

Domains Topology, and the links are the intra-domain Logical Links and inter-domain links (each one with its attributes), as shown in Figure 6c

Using the Virtual Network Topologies, the E2E-PCE can compute the End-to-End Best Path (E2E-BP) using a known TE routing method, taking into consideration the constraints specified in the Path Request and both the intra and inter-domain network resources. The E2E-BP comprises a set of inter-domain InterLinks and a set of Best Path Segments (BPSs), each Best Path Segment corresponding to a selected Logical Link in the Virtual Domain Topologies.

10

The E2E-PCE then sends a Best Path Segment Communication (BPSCom) to the Intra-PCEs of the domains involved in the Best Path, step 207. The Best Path Segment Communication carries the Best Path Segment (i.e. the LL selected from those provided in that domains VDT) and all the information needed by the inter-domain signalling protocols, according to the chosen signalling technique. Each Intra-PCE that receives a BPSCom, maps the Logical Link to the corresponding intra-domain path and triggers its set up, step 208.

Once the intra-domain path setup is completed, each Intra-PCE sends a Best Path Segment Setup Acknowledgment (BPSSet\_ACK) to the E2E-PCE, step 209. Following receipt of a BPSSet\_ACK from all the domains in the path, the E2E-PCE sends a Path Reply (PathRep) to the PCC to inform it about the result of end-to-end inter-domain path computation, step 210.

Figure 9 illustrates schematically a source (or destination) entity which might be an endpoint for a communication session established via a plurality of discrete domain. The entity comprises a processor 1, a transmitter 2, and a receiver 3. Both the transmitter and receiver are interfaced to an access domain of the entity. The entity may be configured to generate and send a Path Request as discussed above. Figure 10 illustrates schematically an Intra Domain PCE, again comprising a processor 4, transmitter 5, and receiver 6. Both the processor and receiver are interfaced to a local domain of the PCE, which might be a transit domain for a particular session, or the

source or destination domain. Similarly, Figure 11 shows an E2E-PCE comprising processor 7, and a transmitter 8 and receiver 9. The E2E-PCE may be a standalone node, or may also implement Intra Domain PCE functionality for a given domain.

- 5 It will be appreciated by the person of skill in the art that various modifications may be made to the above described embodiments without departing from the scope of the present invention.

<b>LL</b>	<b>% of Average Available BW (AABW)</b>	<b>% of Average Node Load (ANL)</b>	<b>EAR</b>
<b>LL 1</b>	90%	20%	85%
<b>LL 2</b>	80%	70%	55%
<b>LL 3</b>	30%	20%	55%
<b>LL 3</b>	10%	80%	15%

Table 1

5

<b>Path Request</b>	
SA	118.226.205.011
DA	215.116.212.113
ServType	Unicast
LSPType	Unidirectional
BWUp	10 MB
BWDwn	5 MB
Max_Node_Load	40%
QoS	6
Link Disj.	Yes
Obj_Func	MinLoad
DomExcl.	D6

Table 2

10

Claims

1. A method of constructing an end-to-end routing path across multiple packet switched domains supporting one or more connection oriented technologies, the method comprising:

- generating a domains topology identifying all possible domains, the border elements of those domains and the inter-domain links;
- receiving a request for computation of a path from a source entity in a first domain to a destination entity in a second domain;
- removing from the domains topology all domains and inter-domain links that cannot or are unlikely to feature in the requested path;
- computing candidate intra-domain paths for each remaining domain;
- mapping each candidate intra-domain path to a logical link defined by its two ends, one end being a border element and the other being a border element or the source or destination entity;
- determining attributes for each logical link;
- determining attributes of the remaining inter-domain links;
- generating a virtual network topology by combining the logical links and their associated attributes with the remaining inter-domain links and their associated attributes; and
- determining a best end-to-end route by analysing said virtual network topology.

2. A method according to claim 1, wherein said step of generating a domains topology further comprises:

- at an intra-domain path computation entity in each domain, generating inter-domain data including details of the domain border elements and their inter-domain links and sending said inter-domain data to an end-to-end path computation entity; and
- at the end-to-end path computation entity, using said inter-domain data to generate the domains topology.

3. A method according to claim 1 or 2, wherein said step of requesting computation of a path from a source entity in a first domain to a destination entity in a second domain further comprises:

generating an end-to-end path request including details of the desired destination of the path and any attributes required for that path and sending this to the end-to-end path computation entity.

4. A method according to claim 2 or 3, wherein said step of removing from the domains topology all domains and inter-domain links that cannot or are unlikely to feature in the requested path is carried out at said end-to-end path computation entity upon receipt of said end-to-end path request.

5. A method according to any one of the preceding claims, wherein said step of computing candidate intra-domain paths for each remaining domain further comprises:

- at the end-to-end path computation entity, generating a virtual Domains Topology request for each domain including each of the border elements of each domain that may feature in the end-to-end path, any attributes required by said end-to-end path and, for the source and destination domains, also the source element and the destination element respectively and sending said virtual Domains Topology request to the intra-domain path computation entity of the source domain, the destination domain and of all possible transit domains substantially simultaneously; and
- upon receipt of a virtual Domains Topology request at each intra-domain path computation entity, computing the candidate intra-domain paths for the request.

6. A method according to any one of the preceding claims, wherein said step of determining the attributes for each logical link further comprises determining a cost parameter that represents the availability/suitability of the logical link with respect to

the other logical links computed within that domain

7. A method according to claim 5 or to claim 6 when appended to claim 5, wherein said step of generating a virtual network topology further comprises:

- at an intra-domain path computation entity, sending a virtual Domains Topology reply to the end-to-end path computation entity, the reply including the identified logical links, the border elements that feature in the logical links and the associated attributes and, for the source and destination domains, also the source and destination element respectively ; and
- upon receipt of the virtual Domains Topology reply at the end-to-end path computation entity, combining the logical links and their associated attributes with the remaining inter-domain links and their associated attributes.

8. A method according to any one of the preceding claims, wherein an attribute for a link is given as a percentage of the total network resources that are available within the domain for that link.

9. A method according to any one of the preceding claims and comprising, at an intra-domain path computation entity, monitoring its inter-domain data and sending updates of the inter-domain data to the end-to-end path computation entity.

10. A method according to claim 9 and comprising sending updates of inter-domain data to the end-to-end path computation entity when changes in the inter-domain data cross some defined threshold(s).

11. A method according to claim 5 or to any one of claims 6 to 10 when appended to claim 5, wherein said step of generating a virtual Domains Topology request for each domain comprises including within the virtual Domains Topology request, an identification of the attribute(s) to be determined by the intra-domain path computation entity.

12. A method according to claim 5 or to any one of claims 6 to 11 when appended to claim 5, wherein said step of generating a virtual Domains Topology request for each domain further comprises including within the virtual Domains Topology request, an identification of the routing strategy or algorithm to be used by the intra-domain path computation entity to compute the candidate intra-domain paths.

13. A method of facilitating constructing of an end-to-end routing path across multiple domains supporting packet based connection oriented technologies and comprising:

- receiving inter-domain data from border elements of said domain, said inter-domain data including information relating to each border element of said domain and the links between those border elements of the domain and border elements of other domains;
- sending said inter-domain data to an end-to-end path computation entity;
- receiving a virtual Domains Topology request from said end-to-end path computation entity, said virtual Domains Topology request including the border elements of that domain that may feature in the end-to-end path and any attributes required by said path;
- computing candidate intra-domain paths, determining attributes for each of the computed candidate intra-domain paths;
- mapping each candidate intra-domain path to a logical link defined by its two ends, one end being a border element and the other being a border element or a source or destination entity;
- sending a virtual Domains Topology reply including the logical links and their associated attributes to the end-to-end path computation entity.

14. Apparatus configured to implement the method of claim 13.

15. A method of constructing an end-to-end routing path across multiple domains supporting packet based connection oriented technologies and comprising:

- receiving inter-domain data from intra-domain path computation entities, said

inter-domain data including information relating to each border element of each domain and the links between those border elements of each domain and border elements of other domains;

- generating a domains topology identifying all possible domains, the border elements of those domains and the inter-domain links;
- receiving an end-to-end path request from a source entity, said end-to-end path request including details of the desired source and destination of a path and any attributes required for that path;
- removing from the domains topology all domains and inter-domain links that cannot or are unlikely to feature in the path;
- generating for each domain a virtual Domains Topology request including each of the border elements of that domain that may feature in the path and any attributes required by said path;
- sending said virtual Domains Topology requests to intra-domain path computation entities of the source domain, the destination domain and all possible transit domains;
- receiving a virtual Domains Topology reply from each said domain, a virtual Domains Topology reply including logical links representing each candidate intra-domain path and its associated attributes;
- determining the attributes of each of the remaining inter-domain links;
- generating a virtual network topology by combining the logical links and their associated attributes with the remaining inter-domain links and their associated attributes; and
- determining a best end-to-end route by analysing said virtual network topology.

16. Apparatus configured to implement the method of claim 15.

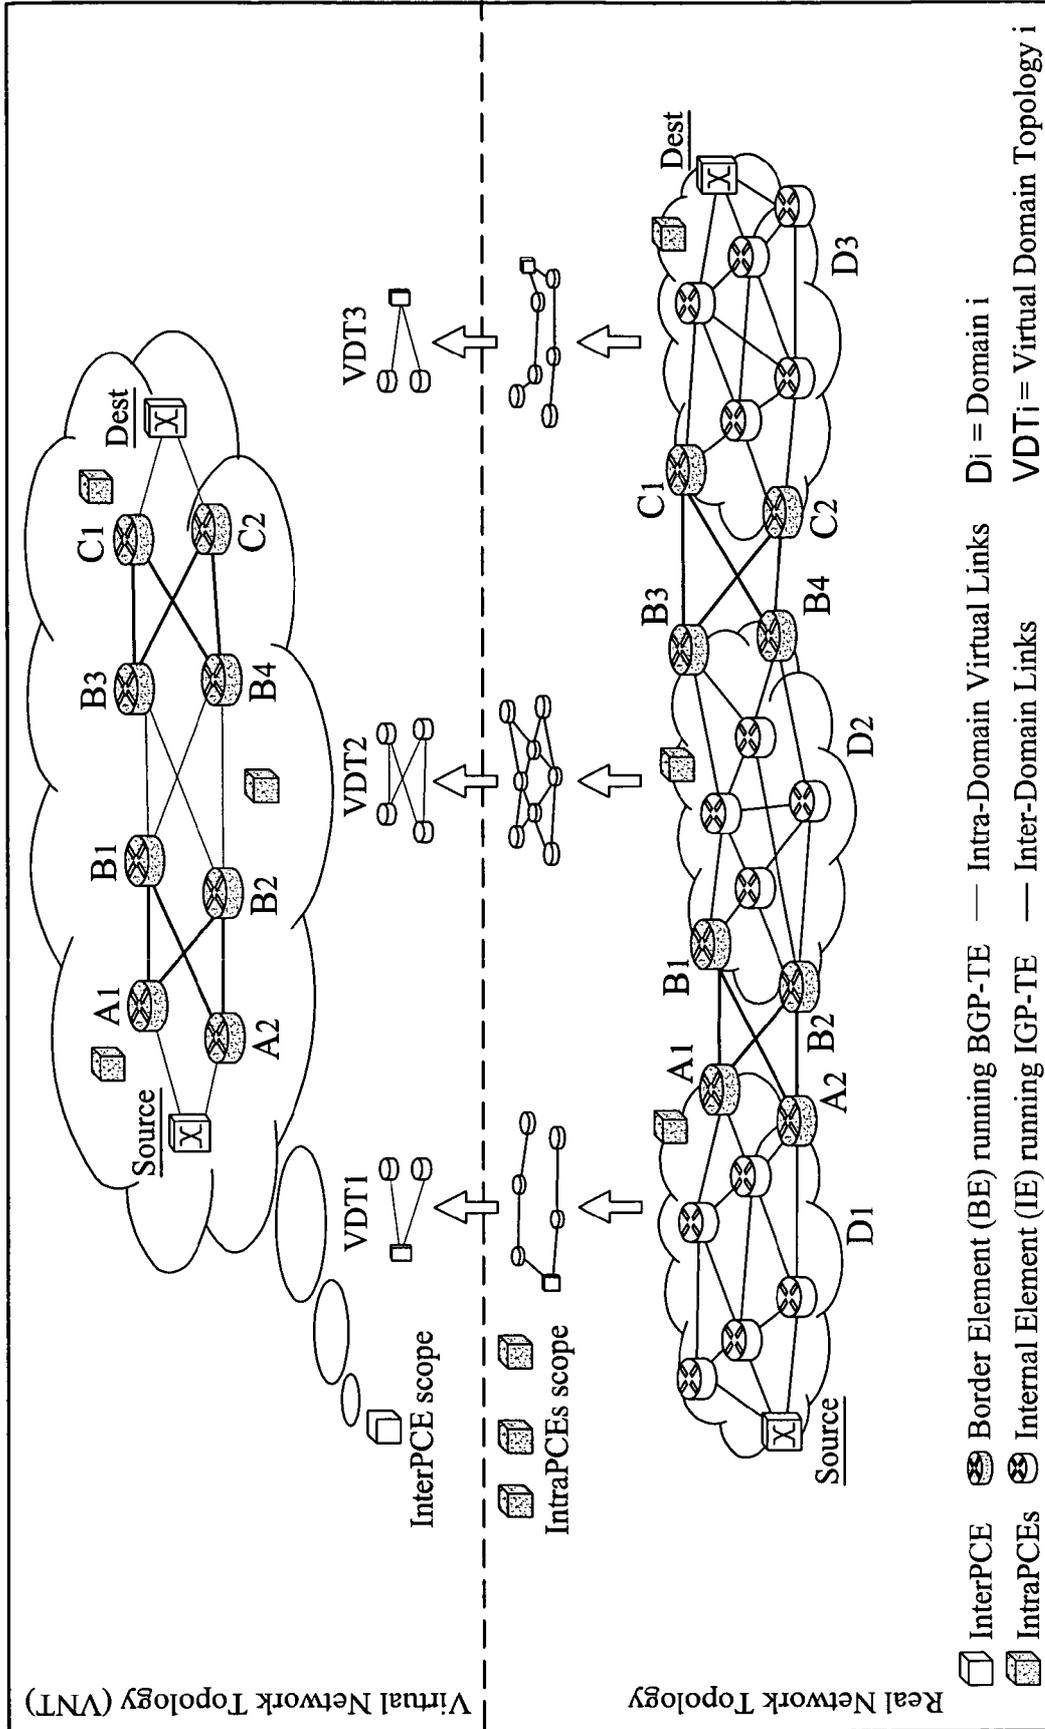


Figure 1



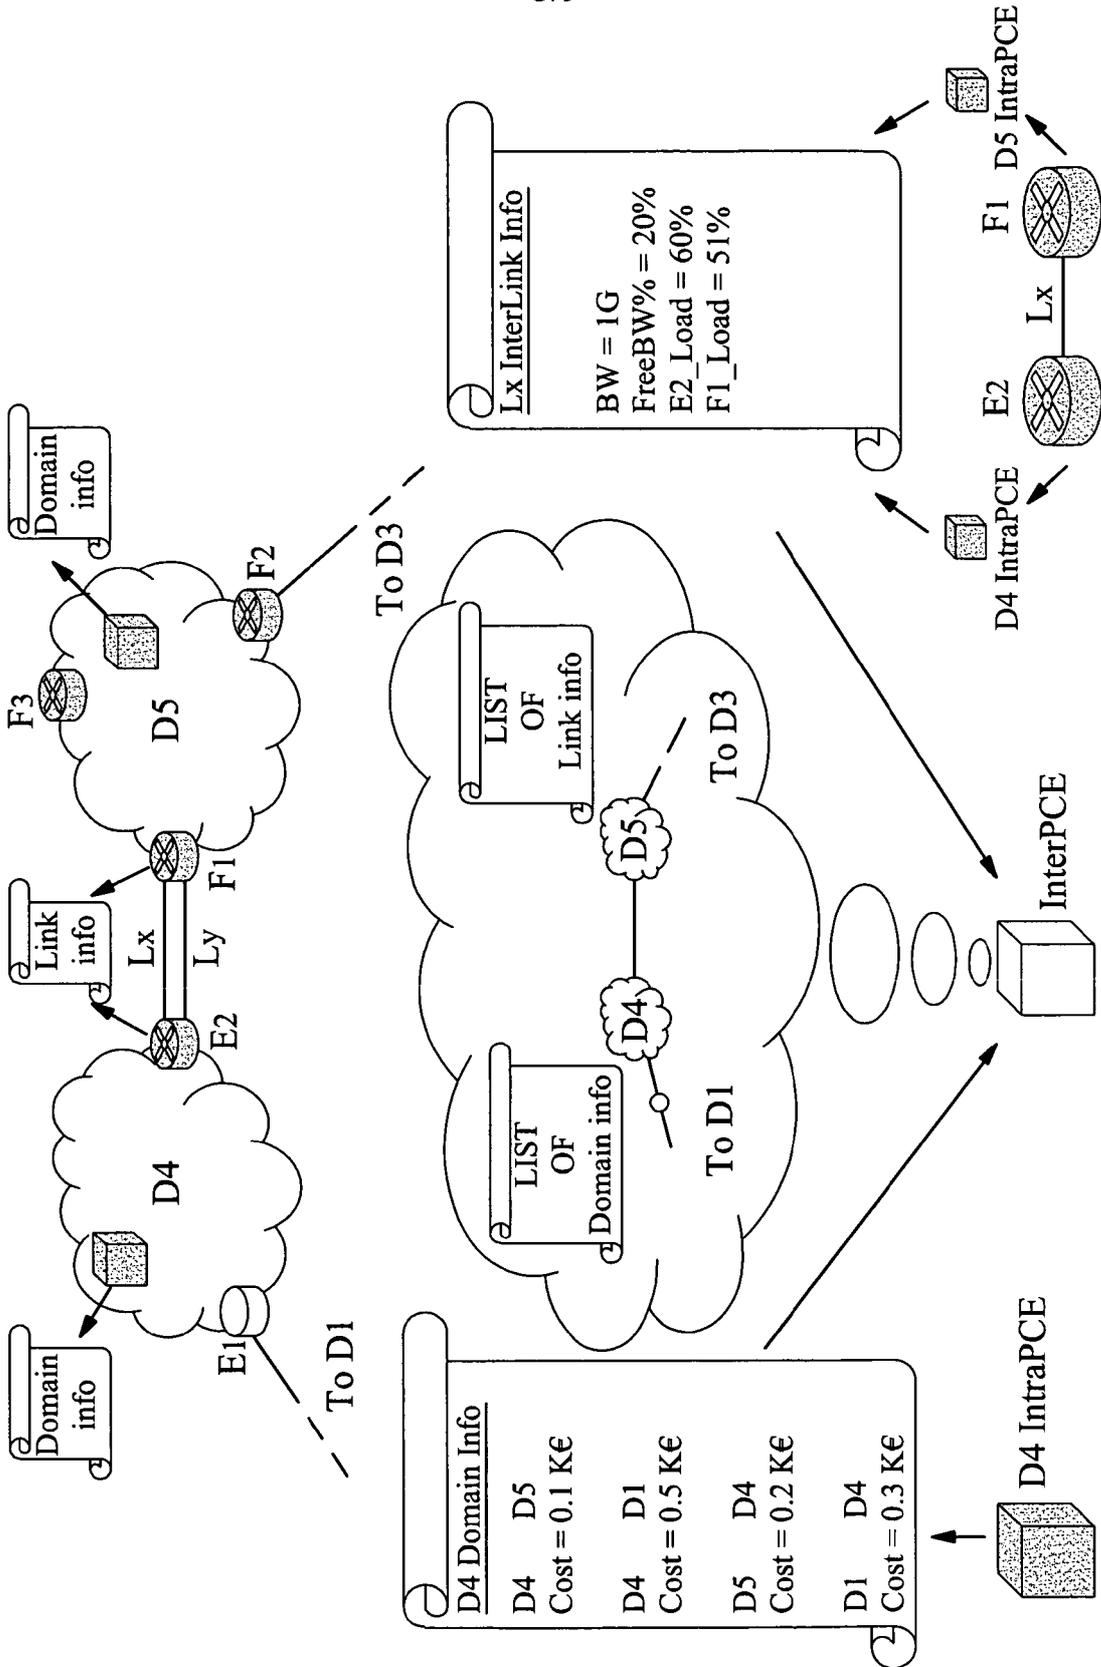


Figure 3

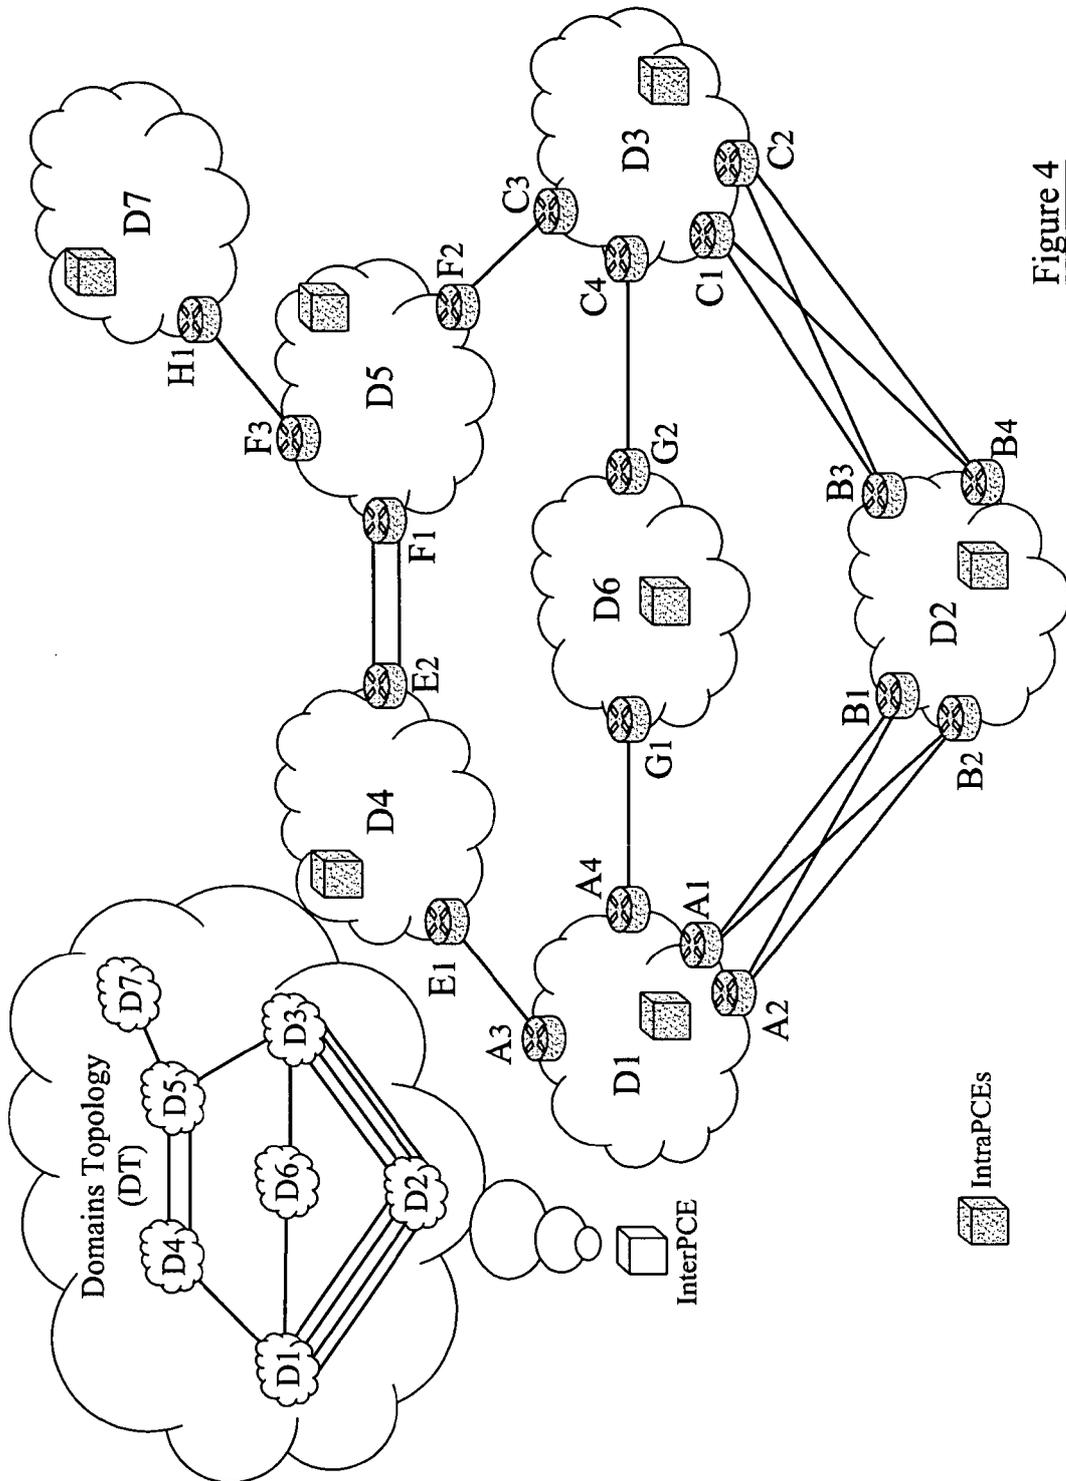


Figure 4

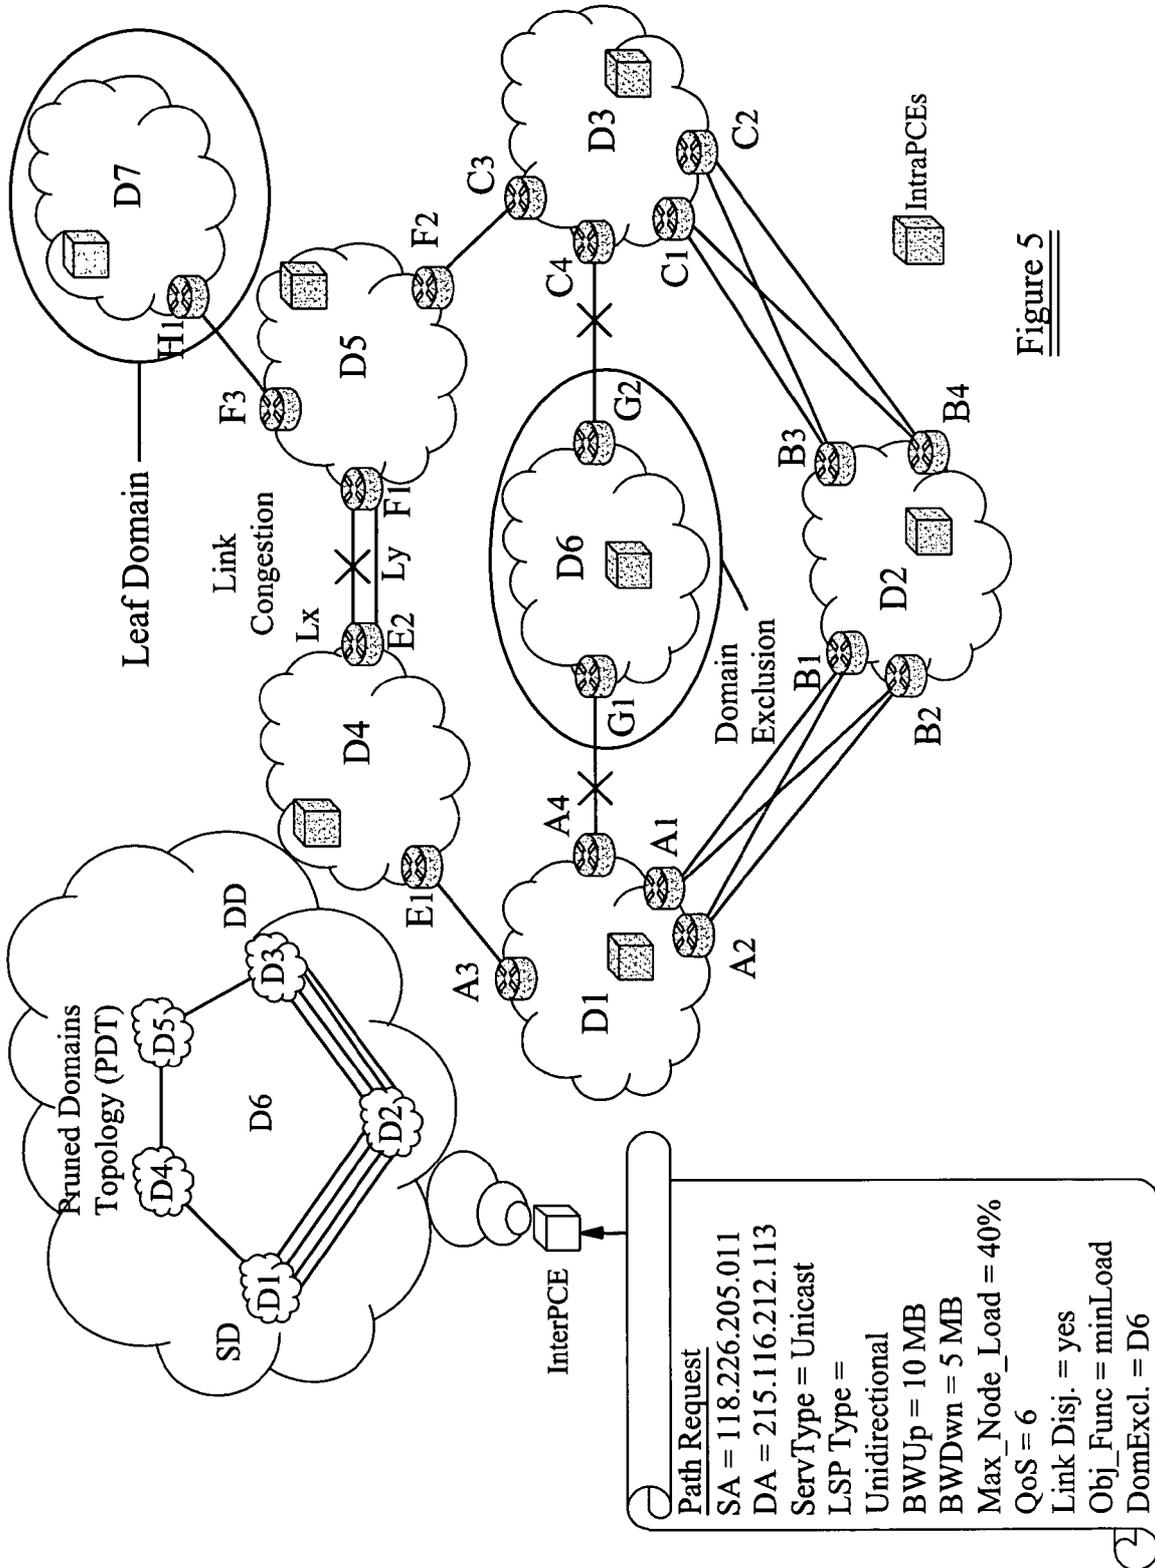


Figure 5

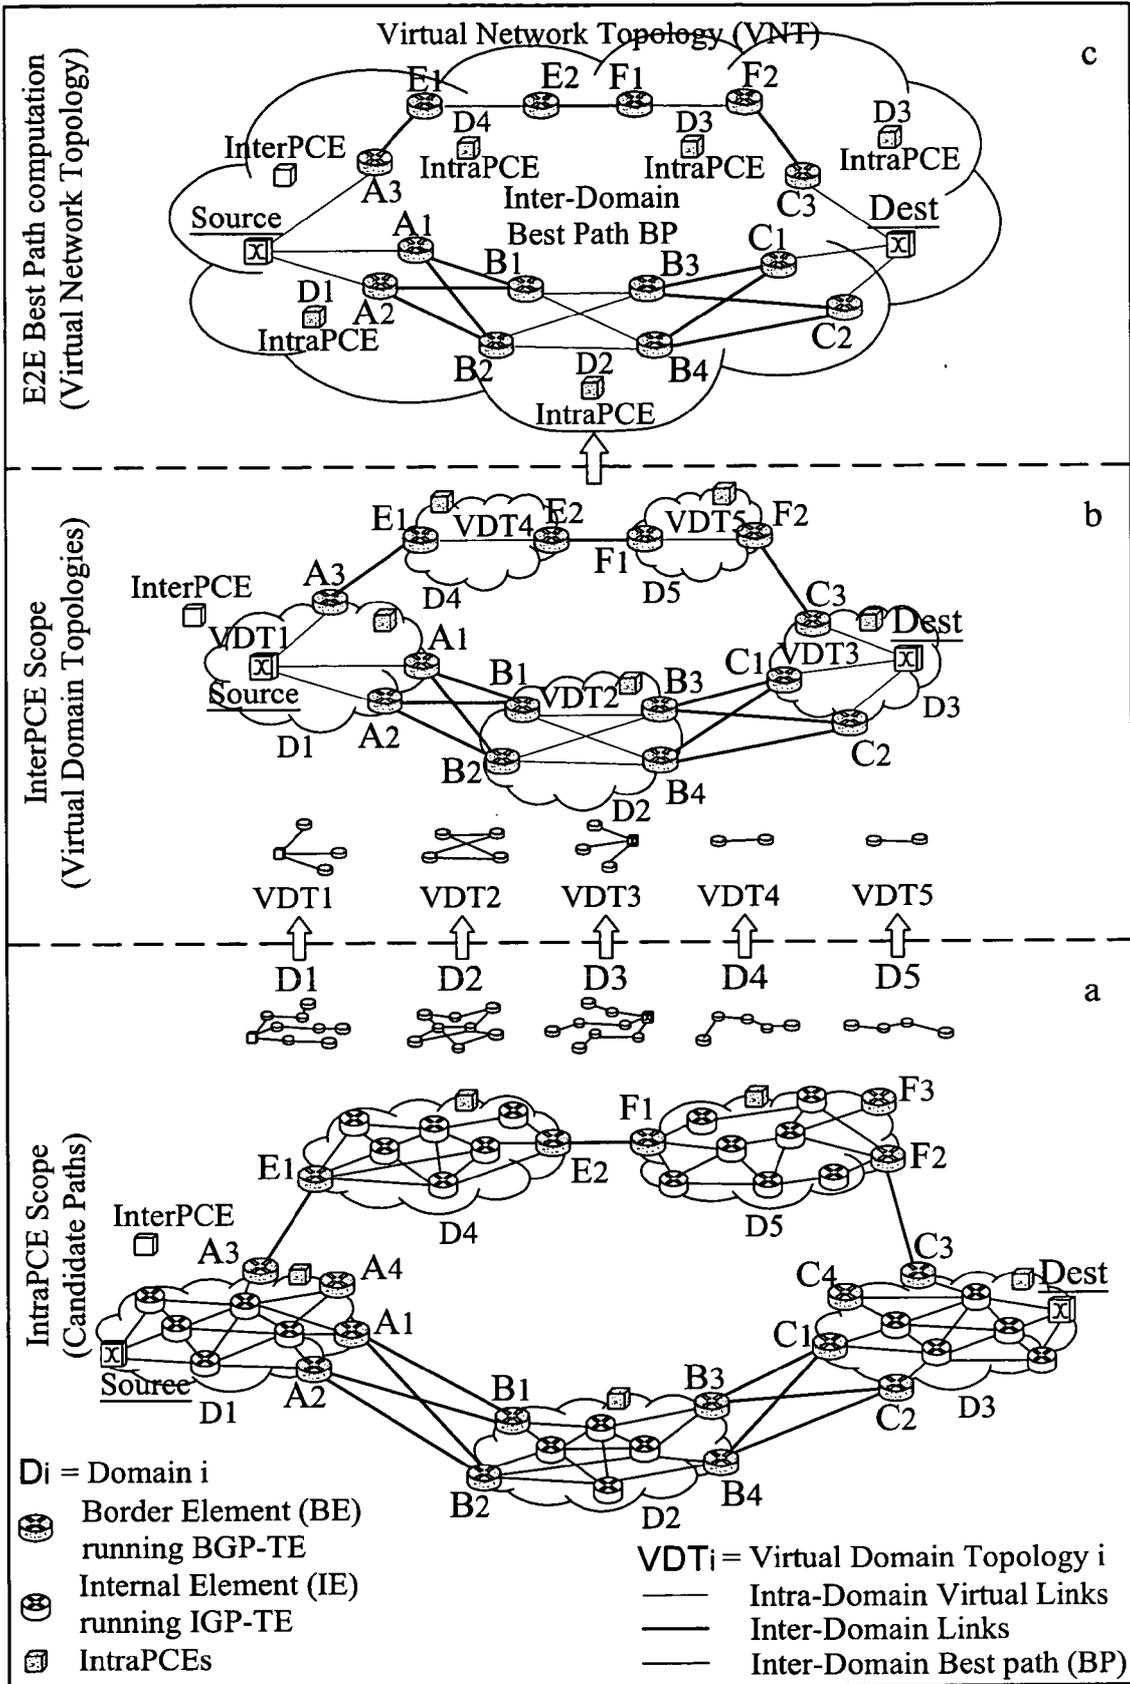


Figure 6

7/9

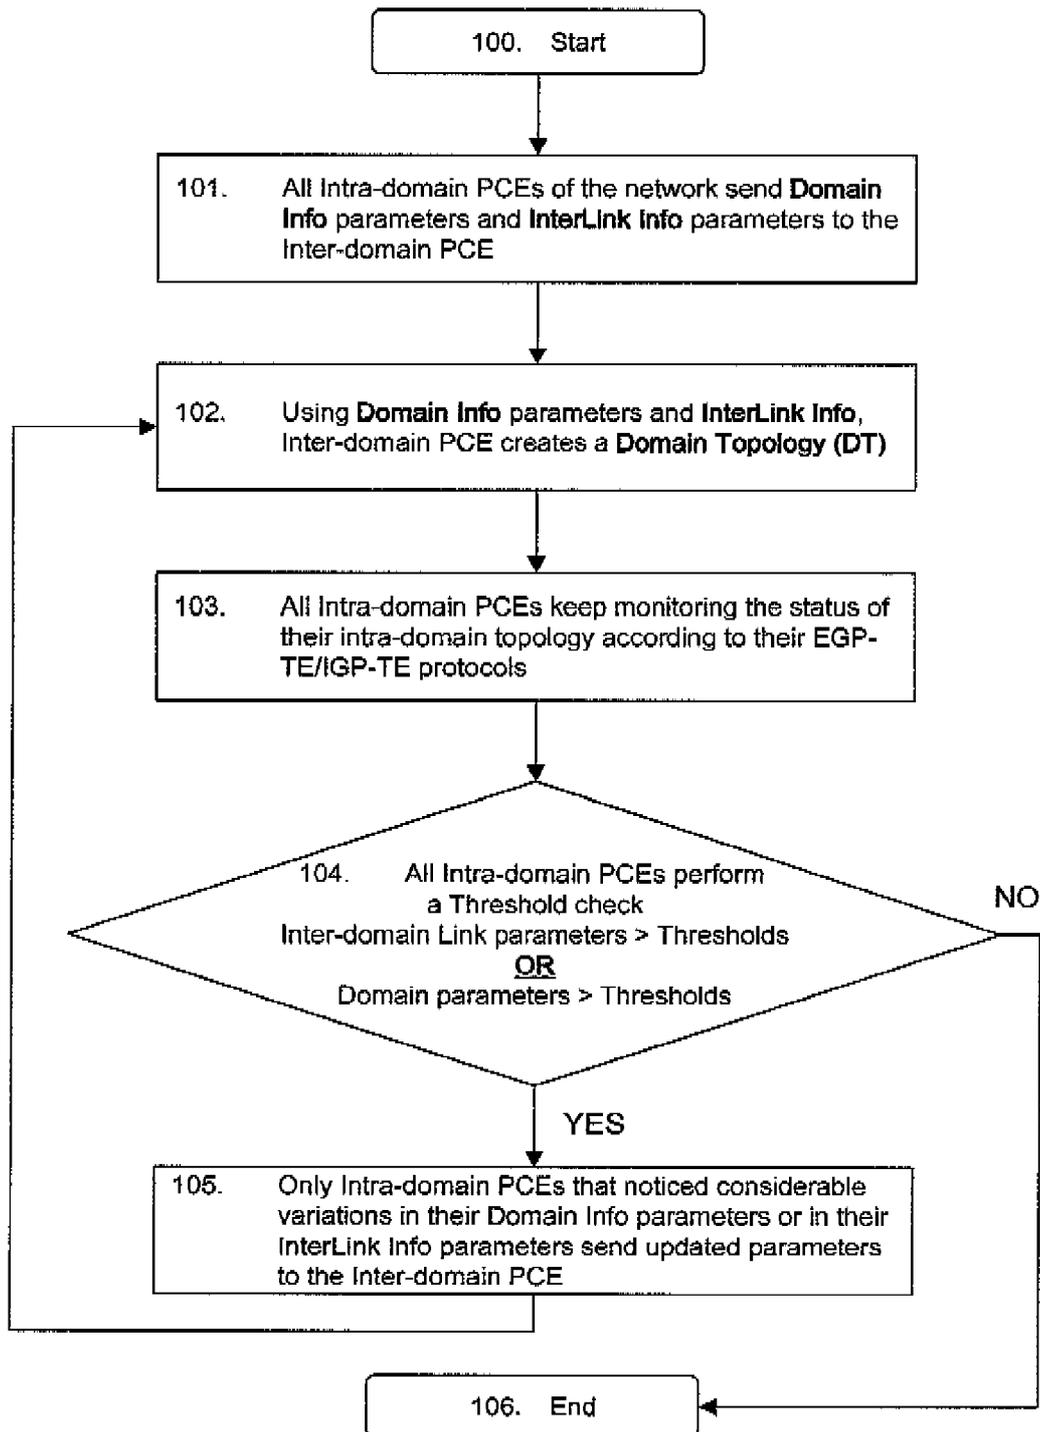


Figure 7

8/9

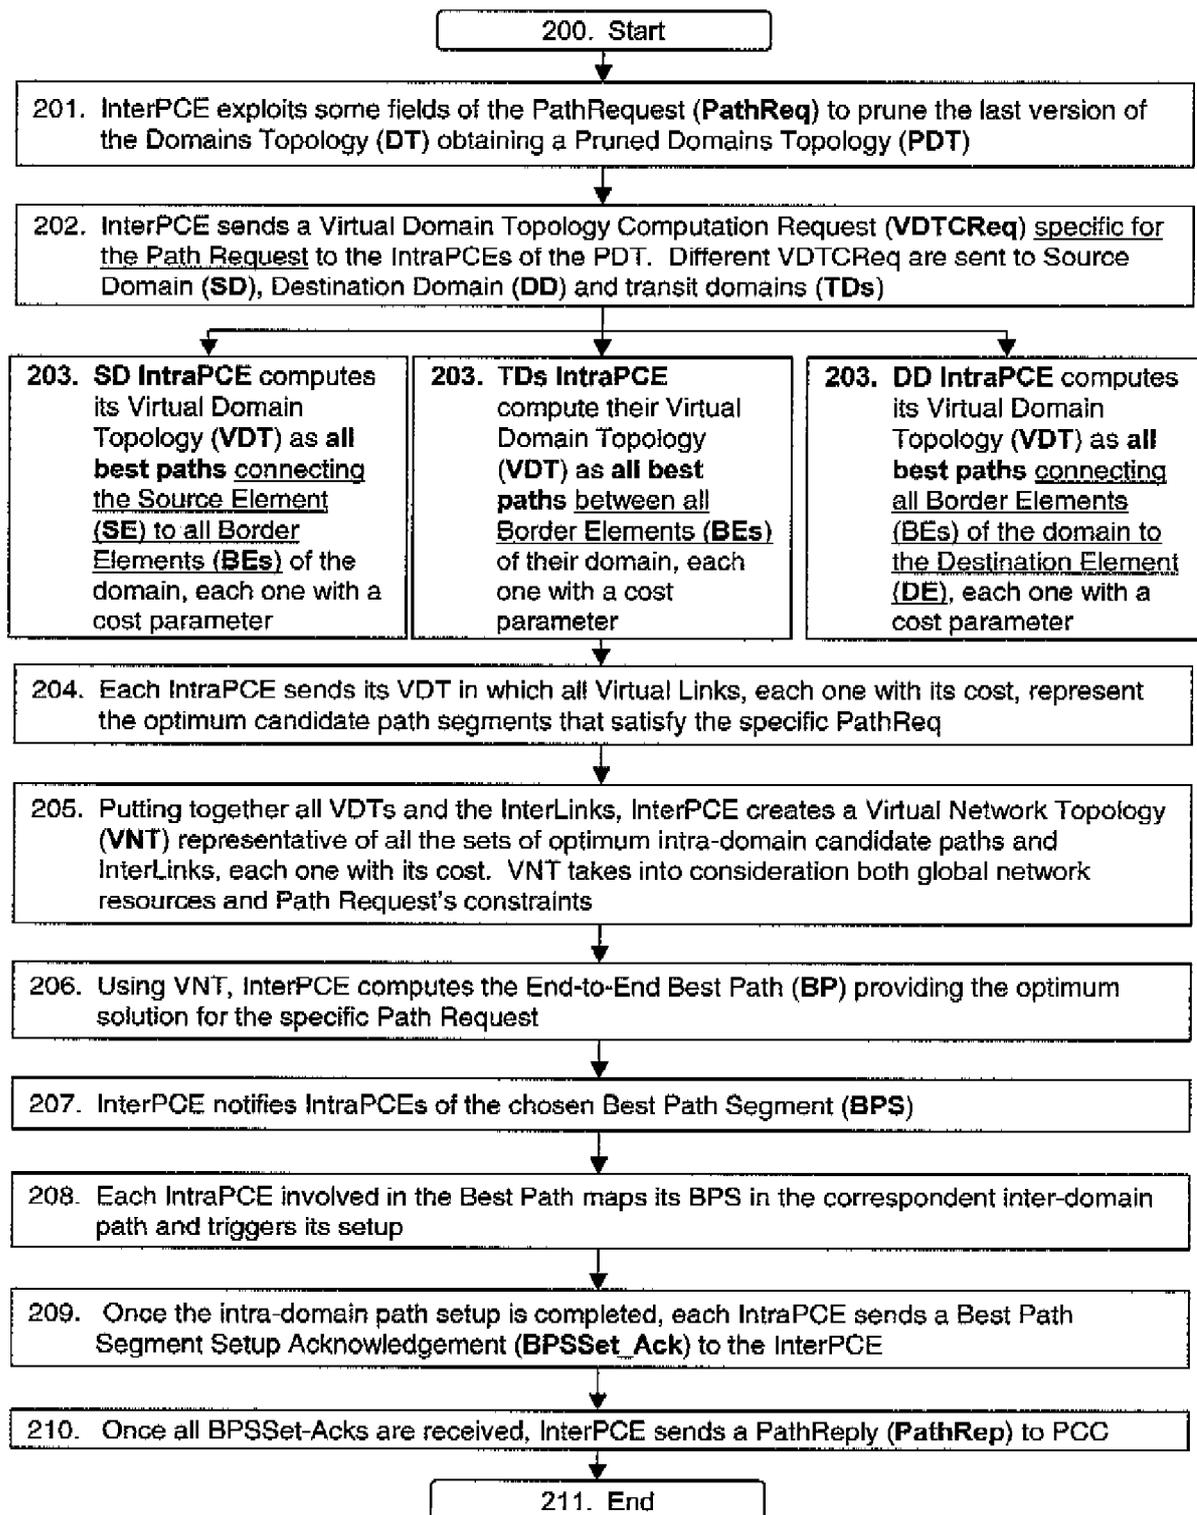


Figure 8

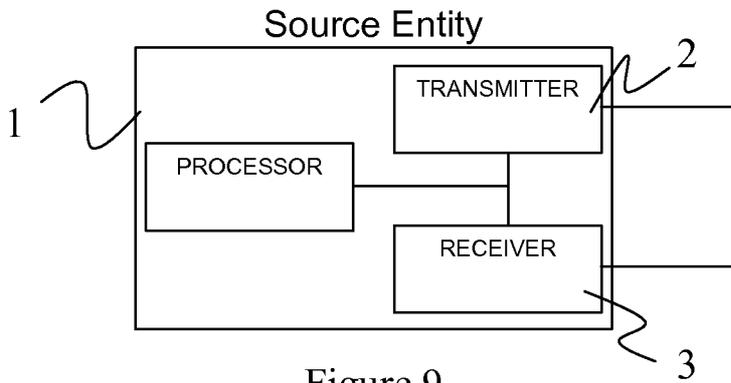


Figure 9

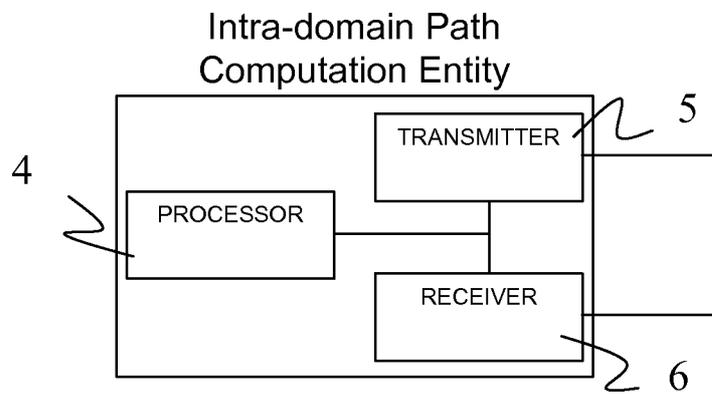


Figure 10

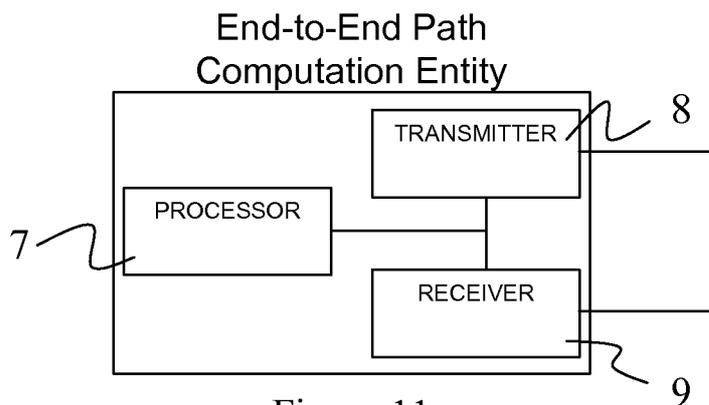


Figure 11

# INTERNATIONAL SEARCH REPORT

International application No  
PCT/EP2008/053772

**A CLASSIFICATION OF SUBJECT MATTER**  
INV H04L12/56

According to International Patent Classification (IPC) or to both national classification and IPC

**B FIELDS SEARCHED**  
Minimum documentation searched (classification system followed by classification symbols)  
H04L

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

Electronic data base consulted during the international search (name of data base and where practical search terms used)  
EPO-Internal

**C DOCUMENTS CONSIDERED TO BE RELEVANT**

Category	Citation of document with indication where appropriate of the relevant passages	Relevant to claim No
<b>X</b>	WO 03/049342 A (IPSUM NETWORKS INC [US]; GUERIN ROCH [US]; RAJAN RAJENDRAN [US]) 12 June 2003 (2003-06-12) abstract paragraph [0013] - paragraph [0018] paragraph [0034] - paragraph [0052] figures 1-8	1-16
<b>X</b>	US 2006/039391 A1 (VASSEUR JEAN P [US] ET AL) 23 February 2006 (2006-02-23) abstract paragraph [0008] paragraph [0027] - paragraph [0047] figures 1-4	1-16

Further documents are listed in the continuation of Box C       See patent family annex

\* Special categories of cited documents

'A' document defining the general state of the art which is not considered to be of particular relevance 'E' earlier document but published on or after the international filing date 'L' document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified) 'O' document referring to an oral disclosure use exhibition or other means 'P' document published prior to the international filing date but later than the priority date claimed	'T' later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention 'X' document of particular relevance the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone 'Y' document of particular relevance the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents such combination being obvious to a person skilled in the art 'Z' document member of the same patent family
---	---

Date of the actual completion of the international search <b>5 August 2008</b>	Date of mailing of the international search report 12/08/2008
---	--

Name and mailing address of the ISA/ European Patent Office P B 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel (+31-70) 340-2040 Tx 31651 epo nl Fax (+31-70) 340-3016	Authorized officer  <b>Pi ata-Andres, Isabel</b>
---	--

INTERNATIONAL SEARCH REPORT

International application No  
PCT/EP2008/053772

C(Continuatio π). ' DOCUMENTS CONSIDERED TO BE RELEVANT		
Category"	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No
A	<p>US 2003/099203 A1 (RAJAN RAJENDRAN [US] ET AL) 29 May 2003 (2003-05-29)</p> <p>abstract</p> <p>paragraph [0012] - paragraph [0019]</p> <p>abstract</p> <p>paragraph [0005] - paragraph [0017]</p> <p>paragraph [0036] - paragraph [0053]</p> <p>figures 1-6</p> <p style="text-align: center;">-----</p>	1-16

# INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/EP2008/053772

Patent document cited in search report	Publication date	Publication date	Patent family member(s)	Publication date
WO 03049342	A	12-06-2003	AU 2002357001 A1	17-06-2003
			EP 1461884 A2	29-09-2004
US 2006039391	A1	23-02-2006	NONE	
US 2003099203	A1	29-05-2003	AU 2002352874 A1	17-06-2003
			EP 1461883 A1	29-09-2004
			JP 2005512393 T	28-04-2005
			UO 03049341 A1	12-06-2003
			US 2003072485 A1	17-04-2003
			US 2003072270 A1	17-04-2003
			US 2005025059 A1	03-02-2005