

Towards an Efficient Computation of High Quality primary and backup paths in Multi-Domain IP/MPLS Networks*

Marcelo Yannuzzi¹, Alex Sprintson², Xavier Masip-Bruin¹, Ariel Orda³,
Sergio Sánchez-López¹, René Serral-Gracià¹, Josep Solé-Pareta¹,
and Jordi Domingo-Pascual¹

¹ Department of Computer Architecture, Technical University of Catalonia (UPC),
Barcelona, Spain

{yannuzzi, xmasip, sergio, rserral, pareta, jordid}@ac.upc.edu

² Department of Electrical and Computer Engineering, Texas A&M University, USA
spalex@ece.tamu.edu

³ Department of Electrical Engineering, Technion, Israel Institute of Technology,
Haifa, Israel
ariel@ee.technion.ac.il

Abstract. Multi-Protocol Label Switching (MPLS) has become the core switching infrastructure at the intra-domain level. However, little progress has been made to extend the reach of MPLS Label Switched Paths (LSPs) across domains. Among the problems that remain unsolved is how to efficiently find and establish primary and protection inter-domain LSPs for mission-critical services subject to QoS constraints. In this paper we first review the major limitations impeding the deployment of these kinds of LSPs across multiple domains. Next, we discuss about the advantages of the recently proposed Path Computation Element (PCE)-based architecture, and we overview how this architecture can be exploited to tackle some of the main limitations exposed. Finally, we report the key features of a distributed routing scheme that we have recently proposed, which runs between the PCEs. This distributed routing approach allows a source PCE to find and establish *optimal* disjoint QoS paths across multiple domains in an efficient way.

Keywords. Inter-domain routing, PCE, disjoint LSPs, QoS.

1. Introduction

Multi-Protocol Label Switching (MPLS) is one of the key technologies responsible for the changing landscape of switching and routing inside domains. It leverages the integration of legacy layer 2 switching technologies such as ATM or Frame-Relay, and novel switching technologies such as Gigabit Ethernet, into a single converged network infrastructure. In addition to this integration capability, the label stacking and switching features of MPLS turn it into a flexible and efficient technology to transport IP-based services [1]. The advantages of MPLS are not confined only to integration and efficient

* This work was partially funded by the MCyT (Spanish Ministry of Science and Technology) under contract TEC2005-08051-C03-01, the European Commission through the projects IST-NOBEL under contract FP6-506760 and IST-ePhoton/One under contract FP6-001933, for UPC authors.

switching of IP packets. Indeed, the available implementations of MPLS offer a suite of value-added features, which have been key drivers for the success and massive deployment of MPLS at the intra-domain level. Among these features are:

Virtual Private Services–MPLS offers a flexible and scalable way to provision and manage Layer 3 Virtual Private Networks (VPNs) [2], together with Layer 2 Virtual Private LAN Services (VPLSs) [3].

Traffic Engineering (TE)–This feature is supported by extensions of the two most widely deployed Interior Gateway Protocols (IGPs), namely OSPF-TE [4] and IS-IS-TE [5], as well as TE extensions of the Resource Reservation Protocol (RSVP) [6]. While the IGP extensions are used to carry information about the state and availability of network resources, such as link bandwidth utilization or link delays, the RSVP-TE extensions are widely used for signaling purposes. RSVP-TE supplies MPLS networks with valuable TE functionalities including resource reservations, explicit path routing, reroute around a node or a link failure, and preemption of resources along Label Switched Paths (LSPs).

QoS delivery–IP/MPLS networks support resource reservations, which are essential to provide performance guarantees for services subject to QoS constraints. During the establishment of an LSP, a routing process is invoked, which needs to take into account the QoS constraints required for the path setup. Such a routing approach is usually referred to as QoS Routing (QoSR), and it is critical to support QoS paths in MPLS networks [7, 8].

Restoration Features–MPLS offers a fast restoration mechanism to recover from network failures based on pre-established backup LSPs [8]. This approach allows the network to promptly switch to the backup LSP as soon as a failure is detected. The key issues regarding this kind of restoration method are: i) both the primary and the backup paths need to meet the QoS constraints; ii) the primary and the backup paths should preferably be link (node) disjoint.

Many research efforts have been and are still devoted to analyze and improve different facets of these MPLS features in the context of a single domain. At present, most of these efforts should move into the inter-domain area. The reason for this is twofold. First, customers are requiring from their Internet Service Providers (ISPs) the capability to extend the reach of VPLS and VPN services across domains. These kinds of services typically involve some mission-critical applications, demanding hard QoS-guarantees and fast restoration capabilities from the network, making MPLS the natural option to comply with such demands. Second, ISPs are eager to offer these services, so the research community is facing the challenge of devising the most suitable way of extending the reach of QoS-constrained MPLS LSPs beyond domain boundaries.

A recently chartered working group (WG) by the IETF is starting to address the issue, and their first contribution is the introduction of a new network component inside each domain called the Path Computation Element (PCE) [9]. The WG is expected to draft solutions and provide guidelines for a wide range of unsolved problems, including:

- i) the extension of MPLS Traffic Engineering (MPLS-TE) capabilities across domains;
- ii) the design of novel communication protocols to handle requests for the computation of paths subject to multiple constraints within and between domains; iii) and the definition of the extensions needed for some of the existing routing and signaling

protocols.

From this range of open problems, we focus on exploring the major limitations impeding the deployment of primary and protection inter-domain LSPs for mission-critical services subject to given QoS constraints. Our interest is in advance path protection strategies, i.e., backup paths need to be established jointly with the primary LSPs. The rationale for this approach is that in many practical settings, it might not be possible to restore all QoS protected paths after a failure. This typically depends on the type of failure, and the amount of traffic that needs to be restored. Furthermore, restoring inter-domain QoS LSPs after a failure might take an unacceptably long time for a number of mission-critical applications. Thus, for this kind of applications, switching promptly from a primary to a backup path in the event of a failure can be guaranteed by provisioning two disjoint QoS paths between the source and destination nodes. We hope that some of the discussions presented here encourage researchers to explore novel Internet routing models, supplying solutions supporting the establishment of high quality (primary and protection) disjoint QoS paths at a multi-domain level.

The remainder of this paper is structured as follows. In Section 2 we analyze the main limitations imposed by the current multi-domain routing model. In Section 3 we succinctly describe the main features of the PCE-based proposals coming from the IETF, including the Explicit Route Object (ERO) expansion [10]. In Section 4 we outline the key features of a distributed routing algorithm running between the PCEs that we have recently proposed, which overcomes the fundamental problems exposed. Finally, a discussion is presented in Section 5.

2. Limitations imposed by the existing multi-domain routing model

In this section we review the main limitations imposed by the existing routing model in order to find and establish primary and backup inter-domain LSPs subject to given QoS constraints. These limitations can be grouped into the following three categories: 1) lack of state information exchange between domains; 2) export policies and autonomic policy-based routing; and 3) scarce path diversity.

2.1 Lack of state information exchange between domains

The Border Gateway Protocol (BGP) is the inter-domain routing protocol used in the global Internet [11]. BGP is a path-vector routing protocol, which only handles and exchanges reachability information between Autonomous Systems (ASs). In other words, BGP routers do not exchange network “state” information, such as link or path bandwidth utilization, or link or path delays. Furthermore, BGP routers are completely unaware of the topology of the Internet. A BGP router only handles destination prefixes, and the Next-Hop to reach each destination. The approach of handling solely reachability information in BGP has proven to supply a highly scalable inter-domain routing framework, but unfortunately, it hinders the deployment of QoS across domains [12].

2.2 Export policies and autonomic policy-based routing

In order to understand the way inter-domain routing information flows in the Internet,

as well as the basic content of this information, it is mandatory to introduce first the business relationships between ASs. There are two major types of business relationships, i.e., customer-provider and peer-to-peer, which correspond to the two different traffic exchange agreements between neighboring domains. The former applies when a domain buys Internet connectivity from an Internet Service Provider (ISP). The latter typically applies when two providers that exchange a significant amount of traffic, agree to connect directly to each other to avoid transiting through a third-party provider. Peering domains share the costs of the connection between them, so there is no customer-provider relationship in this case. These two types of relationships imply the following usual export policies of the ASs [13].

Customer-Provider Advertisements:

- Each AS advertises to its providers all its allocated IP prefixes and those learned from its own customers, but never those learned from its peers or from other providers.
- Each AS advertises to its customers all the reachable IP prefixes it knows (or sometimes only a default route).

Peer-to-Peer Advertisements:

- Each AS advertises to its peers its own IP prefixes as well as those learned from its customers, but never those learned from its providers or other peers.

These export policies determine the inter-domain routing preferences of a provider as follows. A provider prefers customer routes over peer routes or higher hierarchy provider routes, independently of the AS-path length. Moreover, a provider always prefers peer routes over higher hierarchy provider routes. The overall effect of these routing preferences is that inter-domain routes cannot be inferred from the topology. These set of rules turn inter-domain routing into policy driven rather than topology driven or network state driven, so finding multiple paths across domains is strictly limited according to these rules. This can be easily shown by means of the example in Fig. 1.

The top figure shows six ASs and a set of inter-domain links connecting the ASs. Let us suppose that AS1 is a customer of AS2 and AS3, which are in turn peers of AS4. Let us also suppose that AS2 and AS3 have a peer-to-peer relationship. In addition, AS5 and AS6 are customers of AS4. The arrows in the figure represent the flow of BGP routing advertisements for the set of prefixes owned by AS4, according to the usual export policies. At a pure AS-graph level, AS3 has three possible paths to reach AS4, i.e., one through AS1, one through AS2, and the one directly linked to AS4. However, the export policies determine that the path directly connecting AS3 and AS4 is actually the only one available for AS3. This very simple example captures the essence of the hierarchical and policy-based routing structure of the Internet, and helps to understand how to determine the number of available paths between ASs.

The export-policies have two major side-effects. First, algorithms for finding optimal disjoint QoS paths typically rely on a directed graph that abstracts the network topology. However, in [14] we show that a multi-domain network cannot be abstracted as a directed graph in the presence of the export policies. Thus, efficient intra-domain algorithms such as the ones proposed in [15] cannot be simply extended for AS-diverse

routing. The second side-effect is that policy-based routing is one of the factors responsible for the next limitation, i.e., scarce path diversity.

2.3 Scarce path diversity

Besides the limitations due to the export policies, other factors contribute to the problem of the scarce path diversity between nodes located in distant ASs. The power-law relationship of the Internet topology, which was first reported in [16], is in fact one of the main contributors to the problem. It reveals the hierarchical nature of the Internet and exposes the issue that only a very few highly connected ASs keep the Internet as a whole [17]. At present, around only twenty of these ASs exist, which means that, at the AS-graph level, the core of the Internet is very small. It also means that the ASs located at the edge of the Internet, which are in fact the ones exchanging the bulk of the traffic, tend to connect to this highly connected group of ASs, which translates into very few paths between distant ASs. In the top of Fig.1, AS1, AS5 and AS6 represent the ASs located at the edge, while AS2, AS3, and AS4 represent the highly connected core. Most of the traffic is generated by the edge ASs, and clearly the path diversity between any pair of them is scarce. For instance, there is only one possible path between AS5 and AS6. Moreover, due to the export policies, there are only two paths from AS1 to AS5, as well as from AS1 to AS6. These paths are {AS1, AS2, AS4, AS5}, {AS1, AS3, AS4, AS5}, {AS1, AS2, AS4, AS6}, and {AS1, AS3, AS4, AS6}. These sequenced lists of ASs are referred to as *AS-paths*.

Another main contributor to the scarcity of available paths is the inter-domain routing

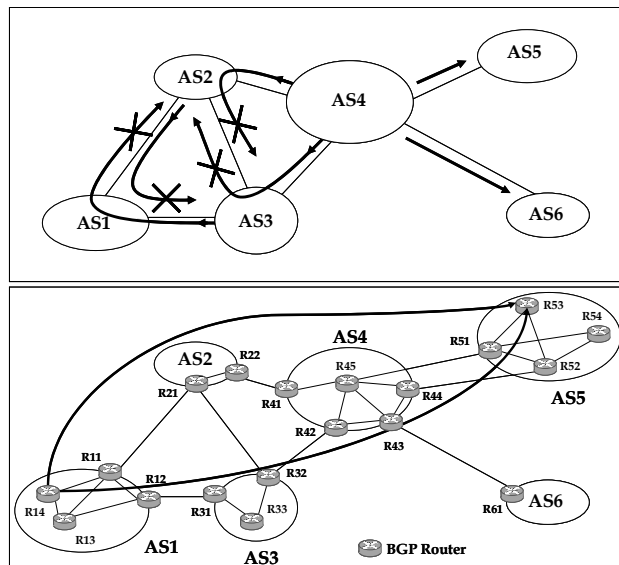


Fig. 1. The multi-domain routing model makes AS-graphs inadequate to find disjoint paths across domains.

protocol. Indeed, BGP introduces two major limitations, which we review in the sequel. First, a BGP routing table typically contains more than one candidate route toward a destination prefix, but BGP routers allocate only one route (the best route) in the forwarding table. Furthermore, BGP routers only advertise the best route they know.¹ This reduces the number of routes handled by upstream domains, supplying a highly scalable routing approach, but unfortunately it drastically prunes the path availability information while flowing upstream. In other words, several paths that could be potentially used for backup purposes are simply unknown by upstream routers due to BGP.

The second limitation introduced by BGP in terms of path diversity is that, for the sake of scalability, BGP handles and advertises highly aggregated information. To be precise, the reachability information advertised by BGP routers only contains AS-path information, that is, a set of destination prefixes and the list of AS hops that need to be traversed to reach those destinations. Clearly, such a list of AS hops offers highly aggregated information by completely hiding the internal structure of the ASs. The advantage of this lack of internal visibility is that it makes BGP highly scalable. A disadvantage, however, is that although several disjoint paths might be available along an AS-path, they cannot be determined.

In the bottom of Fig.1 we disclose the internal structure of the ASs. At the AS-graph level, there are no disjoint paths between AS1 and AS5 (all available paths traverse AS4). Yet, at the router-level, there are in effect link (node) disjoint paths between any router inside AS1 and any router inside AS5. Indeed, it is a well known fact that the ASs at the top of the hierarchy in the Internet have a very dense link/node structure, exhibiting high path diversity between any pair of nodes inside the AS [18]. Unfortunately, for scalability reasons, this information is not conveyed in BGP.

3. PCE-based architecture

The limitations exposed above have motivated the creation of the PCE WG within the IETF. The aim of this initiative is to standardize a PCE-based model to distribute the computation of TE LSPs among different areas of a single domain or within a small group of domains. This model is not considered to be applicable to the entire Internet, and this stems from the fact that there is no such demand at the moment. Most of the ongoing work at the IETF is still focused on inter-area (single domain) issues. Even though the inter-domain case has begun to be analyzed, the discussions are in an early stage. This section provides an overview of the key aspects of this model, and succinctly explores its possibilities in terms of provisioning primary and backup QoS LSPs across domains. Besides the recent standardization of the architecture [9], most of the work in the WG is in the draft stage. Many issues remain open, so from the alternatives that are being discussed, we describe here the one that we consider supplies the most viable approach.

This approach proposes a decoupled architecture, in which path computation tasks are performed by a device that is detached from the head-end MPLS Label Switching Router (LSR). Such device is referred to as the PCE. Each domain may allocate one or

¹ A BGP router usually chooses the shortest AS-path as the best route, and this is the route it uses to forward IP packets and the only one advertised to other BGP peers.

more PCEs depending on its size. For instance, large transit domains can be split into several areas, and use one PCE to handle the path computations within each area. For the distributed computation of inter-area LSPs, a communication protocol is used between the PCEs of the involved areas [19]. Actually, the same model applies at the inter-domain level, so the set up of LSPs spanning multiple domains involves at least one PCE per domain.

Each PCE is capable of computing primary and backup QoS paths within a domain or an area of a domain. To accomplish this task, the network state information of the domain (area) is gathered into a Traffic Engineering Database (TED). The TED is fed by the intra-domain routing protocols (e.g. OSPF-TE or IS-IS-TE) and “raw” BGP information, i.e., by the set of BGP routes that are available before BGP chooses the best route. This increases the number of candidate paths inside the TED. The PCE uses the information contained in the local TED to find primary and backup QoS paths by means of heuristics especially designed to tackle the intractability of the path computation problem. By detaching the path computation tasks from the routers, dedicated PCEs can relieve the LSRs from intensive computations such as finding disjoint QoS paths.

The WG has already drafted the communication protocols between the LSRs and the PCEs as well as between cooperating PCEs [19]. In [19] the LSRs are termed Path Computation Clients (PCCs). The protocol specifies both the PCC-PCE communication, and the PCE-PCE communication for the distributed computation of LSPs. The PCC-PCE part of the protocol supports path requests subject to multiple QoS constraints; it is able to return multiple (disjoint) paths, and takes into consideration features such as security and policies. Accordingly, some of the limitations exposed in the previous section are partially addressed by means of this approach.

Fig.2a) illustrates the PCE-based architecture. The LSR1 in AS1 is the head-end of a requested LSP toward a destination node located in a distant AS (not depicted in the figure). When LSR1 receives the LSP request, the following sequence of actions occurs:

- (1) LSR1 requests PCE1 to compute the path.

- (2) PCE1 queries the TED in AS1 and computes the segment of the inter-domain LSP up to the Next-Hop (NH) AS Border Router (ASBR). If more than one candidate path exists, the heuristic algorithm in PCE1 selects the “best” segment towards the destination (we will discuss this selection process in next section). Suppose that PCE1 selects ASBR21, so it responds LSR1 with a set of strict hops toward this node. Notice that the NH ASBR denotes the ingress ASBR to the downstream domain, so the NH ASBR and the PCE computing the local segment of the path belong to different domains.

- (3)-(4) These steps represent the signaling messages, i.e., the resource reservations and explicit path routing performed by a protocol like RSVP-TE.

Once the signaling messages reach ASBR21, the same process occurs inside AS1, which is represented as the actions from (5) to (8), and this process is repeated on a per-domain basis until the destination AS is reached.

The Fig.2b) shows a more detailed description of the sequence of actions and the role of the different protocols involved in the set up of an inter-domain LSP. The distributed path computation approach explained above is referred to as *Explicit Route Object (ERO) expansion* [10]. The name comes from the RSVP-TE ERO, which allows

signaling a mix of strict and loose hops to be used in the path. A hop may be even an “abstract” node such as an entire AS. Abstract and loose hops are expanded inside each transit domain to a set of strict hops between the ingress ASBR and the NH ASBR.

This approach has two practical advantages. First, it supplies a scalable path computation scheme, since the responsibility and “visibility” of each PCE ends up in the corresponding NH ASBR. Second, it supplies an appealing approach to ISPs, since it leverages confidentiality by hiding the internal network topology of downstream domains. The approach is simple since each PCE computes a piece of the LSP based on its knowledge of the state of resources within its AS, and the reachability information obtained from BGP. Unfortunately, the major drawback of computing paths by segments is that the resulting paths are likely to be far from optimal. For instance, it is a well-known fact that high quality paths are frequently uncorrelated with the routing choices made by BGP.

The issue that remains wide open is how to exploit the PCE-based model to compute *high quality* primary and backup LSPs across a small group of domains in a viable way, that is, without adversely affecting the scalability and the confidentiality features of the




Fig.2a) LSP computation based on ERO expansion on each domain.

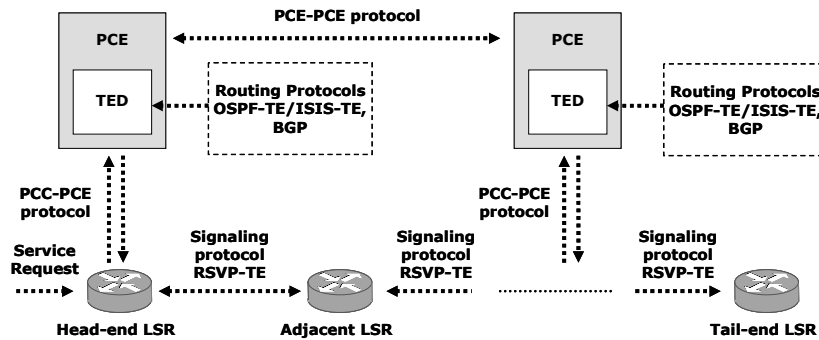


Fig.2b) Messages and protocols involved.

above approach. An analysis about the key challenges raised by this issue can be found in [20].

In the following section we outline the most important features of a distributed routing scheme that we have recently proposed [14]. This distributed routing scheme runs among neighboring PCEs, and supports the computation and establishment of *optimal* disjoint QoS paths across multiple domains in an efficient manner.

4. Towards a PCE-based solution to compute and establish high quality primary and backup LSPs across domains

In [14] we developed a routing model where each PCE is able to compute the *optimal* primary and backup QoS paths to any destination both in a general setting as well as subject to the export policy constraints. One of the major advantages of our approach is that it avoids the well-known *trap topology* problems [21].

To achieve scalability and due to security and administrative considerations, routing domains in our model do not advertise their internal structure, but rather supply an *Aggregated Representation* (AR) to the outside world. Accordingly, a key aspect in the design of distributed routing algorithms is to find an adequate AR that captures the availability of diverse QoS paths across multiple domains. However, there is an inherent trade-off between the accuracy of the representation and the size of the required data structures. In [14], we considered a setting in which a reduced set of neighboring domains are willing to extend the reachability of LSPs across their boundaries. This enables each domain to provide an accurate representation of its traversal characteristics, which, in turn, enables finding optimal disjoint paths across the network. This approach is consistent with that adopted by the IETF PCE WG. The WG has clearly stated that its efforts will focus on the application of the PCE-based model within a single domain or within a small group of neighboring domains, but it is not the intention of the WG to apply this model to the greater Internet.

The distributed routing scheme in [14] is supported by a novel AR for a multi-domain Network, which is small enough to minimize the link-state overhead, and, at the same time, is sufficiently accurate, so that the PCEs can *optimally* find disjoint QoS paths across multiple domains. Our solution guarantees that the confidentiality and administrative limits are respected between domains (e.g., neither the internal topology nor the full IGP state of the domains can be inferred from their ARs).

Another contribution in [14] is that we show that the standard approach of representing a multi-domain network by a graph is inadequate for finding disjoint paths subject to the export policies. However, the export policies can be efficiently represented by employing the concept of the *line graph*. We show that the distributed routing scheme that we developed for the general setting can be easily extended for finding optimal disjoint paths that satisfy the export policy constraints by using the line graph.

4.1 Related work

The problem of finding primary and backup paths subject to QoS constraints in the context of IP/MPLS networks has been widely studied at the intra-domain level. With the advent of the PCE-based architecture, a few recent works have started to extend the study of this problem to LSPs spanning multiple domains. In the current IGP/BGP

routing context, a major issue is that the PCE in the source domain has to compute inter-domain LSPs based on a very limited visibility of the topology and state of the network, yielding solutions that are far from optimal. To cope with this, enriched topological and path state information should be aggregated and available at the PCE in the source domain [20].

In [21] the authors compare the performance of some recently proposed distributed schemes for disjoint path computation of inter-domain LSPs. They assume that the AS-level path was previously computed by BGP at the source domain and that both disjoint paths belong to the same “chain” of domains. This approach has two major limitations. First, solving the two disjoint paths problem restricted to the AS-path selected by BGP will frequently return paths that are far from optimal. This is because BGP does not offer any guarantee about the quality of the chosen AS-path. Second, when several disjoint LSPs need to be established following the same (or part of the same) AS-path, crankback [22] or even blocking might occur, even though the paths could have been established along the alternative AS-paths available at the source domain.

In [14] we studied a PCE-based architecture completely decoupled from the BGP protocol. With this approach, the PCE at the source domain is not compelled to choose both paths along the same chain of domains. This allows the domains to use their multi-homed networks more efficiently. Once the computation of the paths is extended to an expanded AS topology, i.e., not restricting the study to a chain of domains, we need to consider the export policies between domains. As mentioned before, this introduces, however, a major challenge. Whereas the chain of domains can be aggregated and represented as a directed graph, this cannot be done in the presence of the export policies. To solve this problem, we introduced in [14] an AR of the expanded topology using line graphs.

In [23] the authors propose two heuristics so that the PCEs can solve the problem of finding inter-domain LSPs with low end-to-end delay. However, this work addresses the computation of only a single path (without a disjoint counterpart). In addition, the availability of inter-domain paths is inferred directly from the BGP routing information. Accordingly, the authors do not need to address the issue of finding an AR that captures the path diversity and the internal structure of domains.

Another alternative that can be used as an interim solution (e.g. before the deployment of the PCEs) was proposed in [24]. This proposal exploits the multi-connectivity between peering ASs in order to find disjoint LSPs along a chain of domains.

In short, our contributions in [14] can be summarized as follows:

- An accurate AR that captures the path diversity and the internal link state of each domain was proposed.
- A distributed routing algorithm that exploits an AR of the multi-domain network in order to find an optimal pair of link-disjoint paths between the source and the destination in an efficient manner was introduced.
- An efficient method for finding link-disjoint paths subject to the common export policies imposed by customer-provider and peer relationships between routing domains was provided.

5. Discussion and Future Work

The PCE-based model facilitates the provisioning of primary and backup QoS LSPs across domains. The current proposals for finding such paths are based on a coarse selection of the paths by the source domain, and then rely on expansion techniques within the subsequent domains traversed. The strengths of this approach are its scalability and the preservation of the confidentiality of ISPs networks. The main weakness is that the resulting paths are far from optimal.

An important subject is how to exploit the PCE-based model to compute high quality primary and backup LSPs across a small group of domains in a viable way. Approaches tending to endow this model with the capability of aggregating and distributing enriched path state information, allowing a source PCE to compute entire near-optimal LSPs, are worthy of being explored.

We plan to investigate additional approaches that aim at balancing the intrinsic trade-off between the scalability of the aggregated representation of a multi-domain network, and the optimality of the resulting LSPs. We have also plans to address the load balancing and traffic engineering issues related to establishing disjoint QoS paths in a multi-domain environment.

References

- [1] B. S. Davie, and Y. Rekhter, "MPLS: Technology and Applications," Morgan Kaufmann Series in Networking, May 2000.
- [2] E. Rosen, Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)," IETF RFC 4364, February 2006.
- [3] K. Kompella, Y. Rekhter, "Virtual Private LAN Service," Internet draft, draft-ietf-l2vpn-vpls-bgp-06, work in progress, December 2005.
- [4] D. Katz, K. Kompella, D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2," IETF RFC 3630, September 2003.
- [5] K. Kompella, Y. Rekhter, "Intermediate System to Intermediate System (IS-IS) Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)," IETF RFC 4205, October 2005.
- [6] D. Awduche, L. Berger, D. Gan, T. Li, V. Srinivasan, G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels," IETF RFC 3209, December 2001.
- [7] S. Chen and K. Nahrstedt, "An Overview of Quality-of-Service Routing for the Next Generation High-Speed Networks: Problems and Solutions," IEEE Network, Special Issue on Transmission and Distribution of Digital Video, vol. 12, no. 6, pp. 64–79, November/December 1998.
- [8] J. Marzo, E. Calle, C. Scoglio, T. Anjali, "QoS On-Line Routing and MPLS Multilevel Protection: a Survey," IEEE Communications Magazine, October 2003.
- [9] A. Farrel, J. P. Vasseur, J. Ash, "A Path Computation Element (PCE)-Based Architecture," IETF RFC 4655, August 2006.
- [10] J P. Vasseur, A. Ayyangar, and R. Zhang, "A Per-domain path computation method for establishing Inter-domain Traffic Engineering (TE) Label Switched Paths (LSPs)," Internet draft, draft-ietf-ccamp-inter-domain-pd-path-comp-03.txt, work in progress, August 2006.
- [11] Y. Rekhter, T. Li, "A Border Gateway Protocol 4 (BGP-4)," Internet Engineering Task Force, Request for Comments 1771, March 1995.
- [12] M. Yannuzzi, X. Masip-Bruin, and O. Bonaventure, "Open Issues in Inter-domain Routing: a survey," IEEE Network, Vol. 19, No. 6, pp. 49-56, November/December. 2005.

- [13] Y. R. Yang, H. Xie, H. Wang, A. Silberschatz, A. Krishnamurthy, Y. Liu, and L. E. Li, "On Route Selection for Inter-domain Traffic Engineering," *IEEE Network*, Vol. 19, No. 6, pp. 20-27, November/December. 2005.
- [14] A. Sprintson, M. Yannuzzi, A. Orda, and X. Masip-Bruin, "Reliable Routing with QoS Guarantees for Multi-Domain IP/MPLS Networks," accepted for publication in *IEEE INFOCOM 2007*, Anchorage, Alaska, USA, May 2007.
- [15] A. Orda and A. Sprintson, "Efficient Algorithms for Computing Disjoint QoS Paths," in *Proc. of IEEE INFOCOM*, Hong Kong, March 2004.
- [16] M. Faloutsos, P. Faloutsos, C. Faloutsos, "On Power-Law Relationships of the Internet Topology," in *Proceedings of ACM SIGCOMM*, September 1999.
- [17] L. Subramanian, S. Agarwal, J. Rexford, and R. Katz, "Characterizing the Internet hierarchy from multiple vantage points," in *Proceedings of INFOCOM 2002*, June 2002.
- [18] R. Teixeira, K. Marzullo, S. Savage, and G.M. Voelker, "In Search of Path Diversity in ISP Networks," in *proceedings of the USENIX/ACM Internet Measurement Conference*, October 2003.
- [19] JP. Vasseur and JL. Le Roux, "Path Computation Element (PCE) communication Protocol (PCEP)," *Internet draft*, draft-ietf-pce-pcep-06.txt, work in progress, February 2007.
- [20] M. Yannuzzi, X. Masip-Bruin, S. Sánchez-López, Jordi Domingo-Pascual, A. Orda, and A. Sprintson, "On the Challenges of Establishing Disjoint QoS IP/MPLS paths across multiple domains," in *IEEE Communications Magazine*, Vol. 44, No. 12, pp. 60-66, December 2006.
- [21] F. Ricciato, U. Monaco, D. Ali, "Distributed Schemes for Diverse Path Computation in Multi-Domain MPLS Networks," *IEEE Communications Magazine*, vol. 43, n. 6, June 2005.
- [22] A. Farrel, A. Satyanarayana, A. Iwata, N. Fujita, and G. Ash, "Crankback signaling extensions for MPLS and GMPLS RSVP-TE," *Internet draft*, draft-ietf-ccamp-crankback-05.txt, work in progress, May 2005.
- [23] C. Pelsser and O. Bonaventure, "Path Selection Techniques to Establish Constrained Interdomain MPLS LSPs," *Networking 2006*, Coimbra, Portugal, May 2006.
- [24] R. Romeral and D. Larrabeiti, "Combining Border Router Policies for Disjoint LSP computation," in *Proc. of the V Workshop in G/MPLS Networks*, Gerona, Spain, March 2006.