Distributed Schemes for Diverse Path Computation in Multidomain MPLS Networks

Fabio Ricciato, Forschungszentrum Telekommunikation Wien
Ugo Monaco, University of Rome La Sapienza
Daniele Alì, CoRiTeL, Roma

ABSTRACT

MPLS is currently used by several ISPs to carry some high-value traffic components, such as telephony over IP trunks and VPNs. For this type of traffic, service availability is a critical QoS dimension that needs to be protected from network failures. With MPLS-TE, this can be achieved by means of path protection schemes, where active and backup LSPs are routed along diverse paths. Besides protection, path diversity can be exploited for load balancing, another common means of QoS improvement. In order to preserve other QoS requirements, the paths must meet certain constraints (e.g., bandwidth availability, low load) and/or minimize some metric (e.g., hop count). This requires the ability to establish path diversity in an optimal way. In many cases of practical interest, the QoS traffic has an interdomain scope. This is the case for ToIP and VPN traffic between different carriers, or between different ASs owned by the same carrier, as found, for example, after corporate acquisitions or mergers. Therefore, path diversity is a requirement for interdomain traffic engineering. In this work we address path diversity in a multidomain network, where individual domains are capable of connection-oriented forwarding and endowed with an MPLS-TE control plane. For administrative and/or scalability reasons intradomain routing information is not disseminated externally, so dynamic path computation must be achieved by a distributed scheme based on interdomain collaboration. We briefly describe three alternative schemes recently proposed for interdomain diverse path computation, and quantitatively assess their performance with simulations over real ISP topologies.

INTRODUCTION

The recent development of MPLS has refocused attention on the connection-oriented paradigm in backbone networks. In this framework, traffic engineering (TE) techniques are used to dynamically establish connections along optimized paths for some objective function and/or set of constraints. This requires that the network be endowed with a dynamic TE-oriented control plane like multiprotocol label switching with TE (MPLS-TE) or generalized MPLS (GMPLS).

Some operators are already using MPLS to enable service convergence (e.g., [1]). A possible approach is suggested by [2]: some high-value traffic components, typically telephony over IP (ToIP) trunks and premium virtual private networks (VPNs), are carried in connection-oriented mode by MPLS-TE label switched paths (LSPs) while other traffic is forwarded in connectionless mode.

In IP/MPLS networks, quality of service (QoS) parameters such as delay and jitter are accommodated by packet prioritization and careful network dimensioning. Additionally, some selected high-value traffic components (e.g., ToIP) require a very high level of service availability, and for that robustness to failures is a key QoS target. With MPLS-TE this can be achieved through path protection, where a pair of active and backup connections[1] is installed between the given endpoints. Path protection requires path diversity, since the two LSPs must be installed along diverse (or disjoint) paths (i.e., without any common network element). Different kinds of diversity (or disjointedness) can be defined depending on the type of network elements considered — typically links, nodes, or SRLGs.[2] In most cases, in order to preserve the QoS requirements, the connections have to meet certain constraints (e.g., bandwidth availability) and/or minimize some metric (e.g., hop count). These aspects require the ability to establish path diversity in an optimal way.

In several cases of practical interest the traffic to be protected has an intrinsic interdomain scope. This includes ToIP traffic and large VPNs between sites hosted by different carriers. Also, some Internet service providers (ISPs) maintain different legacy autonomous systems (ASs) after corporate acquisitions or mergers (as in [2]). In such cases interdomain path diversity is required

---

[1] *We use the terms* connection *and* LSP *interchangeably.*

[2] *A shared risk link group (SRLG) defines a group of links prone to contemporary interruption upon a single failure event (e.g., fibers sharing a single physical conduit, virtual links running over the same physical circuit).*

0163-6804/05/$20.00 © 2005 IEEE

to support protection. Path diversity can also be used for load balancing, or to split a large flow into smaller parallel connections if there is no path with enough bandwidth to support the whole flow [3]. Because of its importance for such applications path diversity is a requirement for interdomain TE [3].

Traditional schemes for disjoint path computation assume full knowledge of the network topology at the computation point. However, this assumption does not hold in the interdomain scenario. Because of administrative policies and/or scalability constraints, intradomain routing information cannot be disseminated externally, so gathering complete routing information at a single computation element is not possible. This calls for distributed mechanisms involving cooperation between the domains.

In this article we survey some recently proposed distributed mechanisms for dynamic computation of diverse paths in a multidomain environment. We summarize their pros and cons, and provide simulation results to quantify their performance in realistic scenarios. Where necessary, we include some general architectural considerations and references to ongoing standardization.

The rest of the article is organized as follows. We clarify the underlying network model and problem statements. We describe three options for interdomain diverse path computation. We present simulation results and draw conclusions.

PROBLEM STATEMENT

NETWORK MODEL

We consider a network modeled as an ensemble of interconnected domains [4]. We distinguish between border nodes (BNs), placed at the boundary between two or more domains, and internal nodes (INs), which are completely embedded within a single domain. Note that a domain might include BNs only and no INs.

All the information required for path computation within a single domain is called *routing information*: it includes the topology graph plus any other attribute relevant to the computation process (link costs, bandwidth, SRLGs, etc.). We assume that intradomain routing information is fully shared within the domains, and that one or more elements exist within each domain that:

- Maintain a complete view of all intradomain routing information
- Are capable of performing local path computations

Following the terminology in [5], we call these elements path computation elements (PCEs). The organization of intradomain PCEs depends on the architecture under study. The basic options are:

- A single centralized PCE for the whole domain
- PCEs duplicated on each BN

This choice exposes several trade-off dimensions (e.g., communication overhead, duplication of resources, need for synchronization), as is typical whenever centralized vs. distributed designs are compared.

The organization of intradomain PCEs must be distinguished from the issue of coordination between PCEs of different domains. As in [5] we consider the two issues largely independent, and the focus of this work is on the latter. For the sake of simplicity we assume henceforth that PCEs are duplicated within each BN.

SETUP STRATEGIES

For scalability and/or administrative reasons we exclude the possibility of resorting to a centralized super-PCE, with a complete view of routing information for all domains. Therefore, we focus on distributed mechanisms involving some dynamic coordination between the domains.

For scalability and/or administrative reasons we exclude the possibility of resorting to a centralized super-PCE, with a complete view of routing information for all domains. Therefore, we focus on distributed mechanisms involving some dynamic coordination between the domains.

It is useful to distinguish two functions involved in the connection setup:

- *Path computation*: the logical function of determining the connection path
- *LSP installation*: the physical function of installing the LSP (i.e., label allocation)

Both require some communication between the network elements and domains. They can be accomplished in separate phases (first compute then install, or “think-before-make”) or combined in a single signaling phase (“combined approach”).

Think-before-make and the combined approach each has advantages and disadvantages. The main trade-off dimensions are:

- Delay of the overall setup procedure
- Signaling overhead
- Optimality of the solution (path cost)

It is not in our scope to examine these dimensions, nor to say the last word on the underlying architectural framework. This process is currently being conducted by the IETF.[3] The contribution of this article is to provide insight into the performance achievable by the different computation schemes and to assess them quantitatively in realistic scenarios.

Generally, it can be expected that the think-before-make strategy holds potential for better performance than the combined approach; in fact, the latter is inherently prone to missing some possible paths, while the former allows an exhaustive search that always considers all possible solutions. However, the potential for higher performance might come at the cost of certainty of higher complexity and/or overhead, which are particularly critical in a distributed environment. Therefore, a global choice between the two strategies cannot be made a priori, but must take into account a quantitative assessment of the *actual* achievable performances of both options.

TOPOLOGY AND ROUTING ASSUMPTIONS

In the rest of this article we adopt the following assumptions:

- Complete node disjointedness is considered as the relevant diversity constraint. However, the computation schemes presented below can be extended to handle more sophisticated definitions (partial disjointedness, SRLG disjointedness, etc.).
- Interdomain linkage involves at least two different BNs on each side. This assumption is required for node disjointedness, and corresponds to the usual way inter-AS connections are arranged in practice.
- The sequence of crossed domains (i.e., the domain-level path) has been computed

---

[3] *The CCAMP Working Group (http://www.ietf.org/html.charters/ccamp-charter.html).*

> *The ERO can be used by any node along the path to enforce explicit routing of a downstream LSP section. In an interdomain environment, this feature can be used by any ingress BN to explicitly control the LSP section within its own domain to the egress BN.*

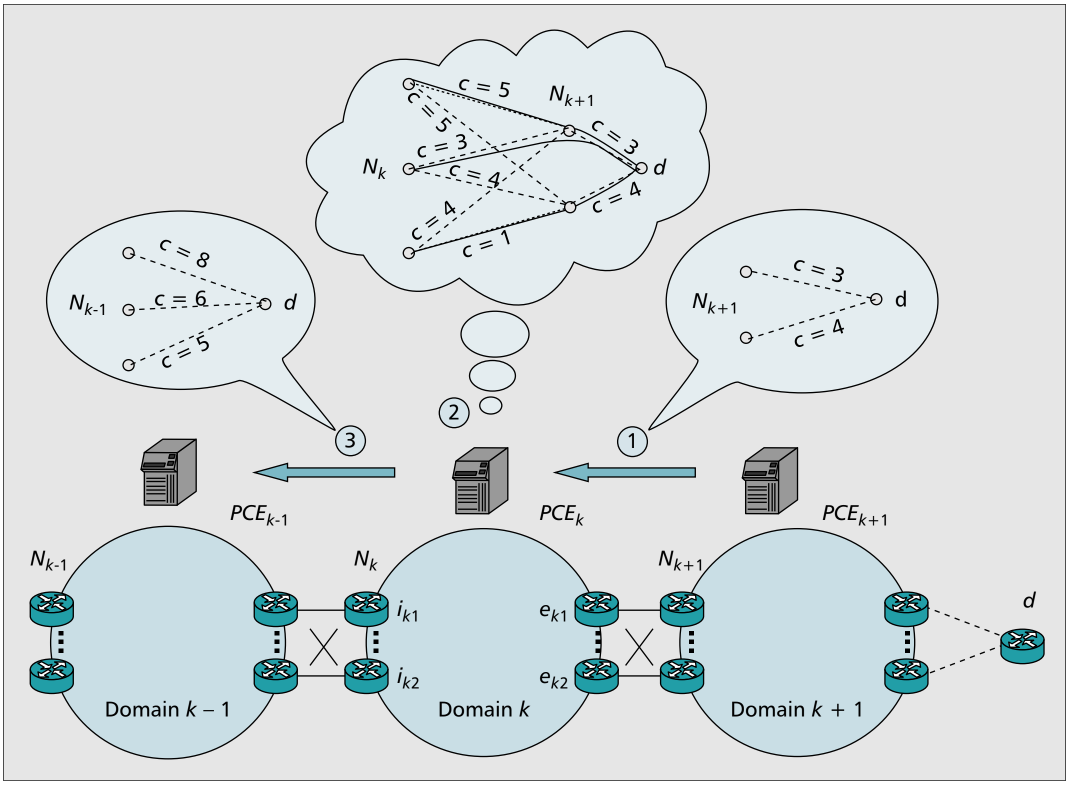

Figure 1. *Reference interdomain scenario and PCE-based computation scheme.*

offline or in any case is known in advance, so dynamic computation only addresses the detailed node-level path (INs and BNs).
- Diversity requirements apply at the node level but not at the domain level (i.e., both paths cross exactly the same set of domains).

These assumptions are helpful for a simpler presentation of the computation mechanisms but do not restrict the generality of the results presented below.

Additionally, Resource Reservation Protocol with TE (RSVP-TE) [6] is the signaling protocol assumed. Recall that a typical RSVP-TE signaling procedure starts with a downstream flow of PATH messages (from source to destination), and follows with an upstream flow of RESV messages. The head-end node (source) can enforce an explicit route by including the list of network elements to be traversed into a specific object within the PATH messages, an *explicit route object* (ERO). The ERO can be strict or loose, depending on whether or not the list is complete. More generally, the ERO can be used by any node along the path to enforce explicit routing of a downstream LSP section. In an interdomain environment, this feature can be used by any ingress BN to explicitly control the LSP section within its own domain to the egress BN.

CURRENTLY PROPOSED SCHEMES

Here we present three schemes that have been recently proposed for interdomain diverse path computation. The first scheme follows the think-before-make approach, and is based on the so-called PCE architecture currently under consideration by the IETF[4] [5]. The other two mechanisms can be implemented with small extensions to the current RSVP-TE protocol, and do not require architectural modification. All three schemes are being discussed in the IETF.

PCE-BASED SCHEME (THINK-BEFORE-MAKE)

In the scheme proposed by [7] path computation and LSP installation are performed in separate phases. The first phase involves collaborative communication between the PCEs along the chain of domains in order to collectively determine the detailed node-level path. In the case of multiple path computation, all paths can be computed jointly as described below. At the end of this phase, the detailed path list is returned to the head-end node so that it can start the installation phase by including it in the ERO. The computation mechanism proposed in [7] ensures an optimal end-to-end solution (i.e., a path pair with minimal total cost). It achieves this through distributed coordination of the PCEs along the domains involved. Figure 1 depicts a generic section of the domain-level path. Index k denotes the order of the crossed domain ($k = 1$ is the source domain), and N_k the number of BNs between domains $k - 1$ and k. Let us first consider the case of single path computation to a remote destination d, with some constraints (e.g., available bandwidth) and the objective of minimizing some additive link cost. The basic idea is that PCE($k + 1$) advertises a set of N_{k+1} candidate paths to PCE(k) along with associated costs (arrow 1, Fig. 1). This set includes the optimal path from each of the N_{k+1} entry points (from domain k) toward d. Then PCE(k) can build a virtual topology where the N_{k+1} border nodes are directly connected to d by virtual links

---

[4] PCE Working Group (http://www.ietf.org/html.charters/pce-charter.html).

with the advertised costs. On top of this virtual topology, PCE(k) can compute a set of N_k optimal paths, from each of the N_k ingress border nodes (from domain $k - 1$) towards d. These N_k candidate paths are then advertised back to PCE($k - 1$) and the procedure is repeated. It can be easily shown that the number of paths advertised at each stage from PCE(k) to PCE($k - 1$) is equal to the number N_k of ingress candidate BNs. The optimality of the final solution is guaranteed since each stage advertises a candidate optimal path for each entry point, which leads to an exhaustive search.

The extension of this mechanism for diverse path computation is straightforward. In order to ensure a globally optimal solution, PCE(k) has to advertise an optimal candidate path for each possible *pair of entry points*, for a total of $N_k(N_k - 1)/2$ paths. By extension, optimal computation of m diverse paths requires that at each stage

$$\binom{N_k}{m}$$

paths are advertised by PCE(k).

The implementation of this mechanism relies on the PCE architecture proposed in [6]. Therein no assumption is made about the communication protocol between the PCEs during the computation phase. At a more abstract level, the PCE architecture introduces a kind of interdomain path computation plane, composed of the PCEs and associated PCE–PCE communication processes, that is logically separated by the signaling plane for LSP installation.

PPRO-BASED SCHEME (COMBINED APPROACH)

A relatively simple scheme for the setup of disjoint LSPs foresees the subsequent installation of the two paths in separate phases and is based on the primary path route object (PPRO), a new RSVP-TE object proposed in [8]. With this scheme, the first LSP (primary) is signaled with the standard RSVP-TE procedure: in each domain, the ingress BN computes the local path section toward the destination until the next egress BNs (a simple run of Dijkstra suffices) and uses the ERO in the PATH messages to drive the label allocation along the computed path section. During the primary LSP signaling, the record route object (RRO) is activated in the RESV messages with the role of collecting the detailed node-level path of the established LSP and reporting it back to the head-end node. Subsequently, the head-end node starts the setup phase of the secondary LSP with a new downstream PATH flow.

The basic idea of this scheme is to include an additional object in the PATH messages of the secondary LSP advertising the node-level network elements included in the primary LSP so that they can be explicitly excluded by the path computation of the secondary LSP. A new object, the PPRO, was proposed in [8] specifically for this purpose; hence, we denote this scheme *PPRO-based*. Each BN computes the downstream (secondary) path section over a reduced graph where the elements listed in the PPRO — and all other elements with common SRLGs — are pruned. Since the PPRO content is copied by the head-end node from the RRO received from the primary path, intersections between the two paths are avoided. With this scheme the computation and installation functions for each LSP (primary, secondary) are combined in a single PATH/RESV cycle. However, the secondary path is set subsequent to the primary path.

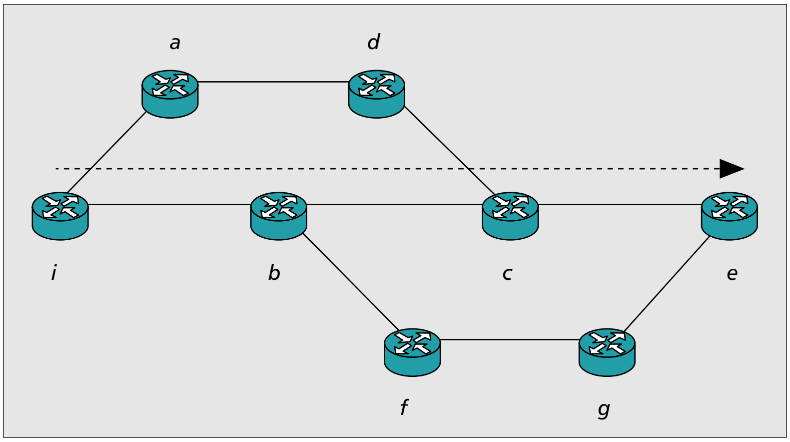

Figure 2. *Trap topology: The shortest path from i to e across b–c leaves the residual graph disconnected. Therefore, sequential computation fails to compute the diverse pair (i-a-d-c-e and i-b-f-g-e).*

Note that also the exclude route object (XRO) proposed in [9] can be used in place of the PPRO. In general, the semantics of XRO and ERO are complementary: while the latter explicitly *includes*, the former explicitly *excludes* specific elements from path computation. Interdomain path diversity is listed among the applications of XRO in [9]. Within the scope of this work the role of XRO is equivalent to PPRO; therefore, we maintain the PPRO-based terminology to address both options.

The advantages of the PPRO-based solution are its simplicity and the marginal impact on the existing RSVP-TE protocol, since only one additional new object is required. However, a potential point of concern is the performance of the underlying computation scheme. In fact, in certain cases a pair of diverse paths exists in the network graph but the subsequent computation strategy fails to find them. A simple example is provided in Fig. 2. This and similar cases are called "trap topologies": they are well known in the literature of intradomain protection schemes, and according to [10] are found in typical carrier network backbones. In other cases, subsequent computation might lead to paths with suboptimal total cost (an example is provided in [11, Fig. 6]).

ARO-BASED SCHEME (COMBINED APPROACH)

The mechanism proposed in [12] (see also [11]) aims to overcome the limitations of the PPRO-based scheme, while keeping a comparable degree of simplicity and adherence to the legacy RSVP-TE specifications. Similar to the above, an additional object is required in the RSVP-TE messages, an associated route object (ARO). Again, the setup of a disjoint connection pair involves two subsequent signaling procedures along the primary and secondary paths. But in this case the computation of both paths is concentrated in the first phase. In other words, the two paths are computed jointly. This is accom-

Since our focus is on the node-level path computation we consider a simple domain-level topology: a linear chain of K domains. In this work we always consider a single intermediate domain between the source and destination domains (K = 3).

plished by using the ARO in conjunction with the ERO during the first PATH flow, as briefly explained here.

The key point of this scheme is that each BN in domain k is able to jointly compute a pair of diverse paths from any pair of BNs toward the remote destination d. As an example, consider Fig. 1, where BN i_{1k} has received a PATH with a request to install an LSP toward d and at the same time compute an associated diverse path from i_{2k} to d. At this point node i_{1k} has to jointly compute the paths P_{1k} (from i_{1k} toward d) and P_{2k} (from i_{2k} toward d). Any algorithm for joint computation of diverse paths can be adopted (e.g., the well-known Suurballe algorithm) [13]. While both ingress nodes (i.e., i_{1k} itself and i_{2k}) are assigned, the selection of the egress nodes (e_{1k} and e_{2k} in Fig. 1) follows from the computation of P_{1k} and P_{2k}. At this point, the BN i_{1k} includes P_{1k} into the ERO and P_{2k} in the ARO of the new PATH message, and then proceeds with the installation of the primary path. The ERO will drive the installation of the primary LSP section from i_{1k} to e_{1k}, while the ARO is left idle. When the PATH message is received by e_{1k}, it computes another section of diverse paths from e_{1k} and e_{2k} — the latter being implicitly advertised in the ARO — toward d, and the process of ARO expansion in the PATH messages is continued until the destination is reached. At this point the primary path has been computed and partially signaled (PATH phase), and the secondary path has been completely computed and collected in the ARO. The secondary path can be returned in the RESV messages to the head node, which will start the subsequent installation of the secondary LSP.

DISCUSSION

Neither the PPRO-based nor the ARO-based schemes guarantee the globally optimal solution. In fact, in both cases the coordination between neighboring domains is too weak for ensuring convergence to the global optimum. The key point is that in both cases the determination of the egress BNs at each stage is not coordinated between the domains. As discussed above, the PCE-based approach described earlier was specifically designed to achieve this coordination and therefore guarantees optimal computation. For this purpose it requires a separation of the computation and installation phases and a certain communication overhead between the domains.

When considering diverse path computation in a single domain, the joint computation strategy — with Suurballe — is always superior to the subsequent computation (i.e., two runs of Dijkstra, the latter with a pruned topology). In fact, it can be demonstrated that in a single domain the former always returns the optimal path pair (i.e., with minimum total cost) and does not suffer trap topologies. Intuitively, it can be expected that the joint computation strategy is also superior to the subsequent computation in the interdomain scenario with distributed (per-domain) computation, despite both schemes being su-optimal due to weak interdomain coordination.

In principle, both schemes can fail to complete the setup at the first try, despite the existence of some suitable path pair in the network (e.g., in case of a trap topology). In this case, one might revert to the so-called crankback procedure [14], where an already established path section is torn down and resignaled along an alternative route. Therefore, the signaling procedure is (partially) repeated and eventually converges to a solution in a finite number of iterations. However, implementation of crankback is rather inconvenient since it incurs additional protocol complexity and overhead, and it will not be considered further here.

SIMULATION RESULTS

GOALS

In this section we contribute to a quantitative performance evaluation of the above computation schemes, particularly PPRO- and ARO-based. The primary performance indicator is the success rate, defined as the conditional probability of finding a diverse path pair on the first attempt (i.e., without crankback) given the existence of a solution between the assigned endpoints. A secondary performance indicator is the total cost of the solution, defined as the sum of the costs of the two paths.

In summary, the considerations made above suggest the following qualitative statements:

- The PPRO- and ARO-based schemes are generally sub-optimal, with success rate < 100 percent and nonminimal total cost.
- The ARO-based scheme is expectedly superior to the PPRO-based one: higher success rate and lower total cost.

In order to empirically verify these claims and assess the performance quantitatively we implemented the PPRO- and ARO-based schemes in an ad hoc simulator. For each experiment we also computed the global optimum path pair, assuming full knowledge of the complete interdomain node-level graph. Recall that the PCE-based scheme is guaranteed to be optimal, with success rate 100 percent and minimal total cost.

SIMULATION SCENARIO

Since our focus is on node-level path computation we consider a simple domain-level topology: a linear chain of K domains. In this work we always consider a single intermediate domain between the source and destination domains ($K = 3$). The internal node-level topology of each domain is derived from real ISP topologies, specifically the ISP maps listed in Table 1. All links are assumed to have the same capacity and cost.

The interconnection between two neighboring domains was designed as follows. On each side we selected a number N of nodes to act as border nodes (we used $N = 2, 3$), then connected the N pairs with 1:1 direct links, as in Fig. 3b. We considered two different policies for BN selection: random selection and max-degree. With the latter, the nodes with the highest number of attached links are selected as BNs. The interdomain connectivity patterns produced in this way are not necessarily representative of real practice. However, without explicit knowledge of how two ISPs would connect to each other, this simple scheme is sufficient for performance evaluation of the path computation scheme.

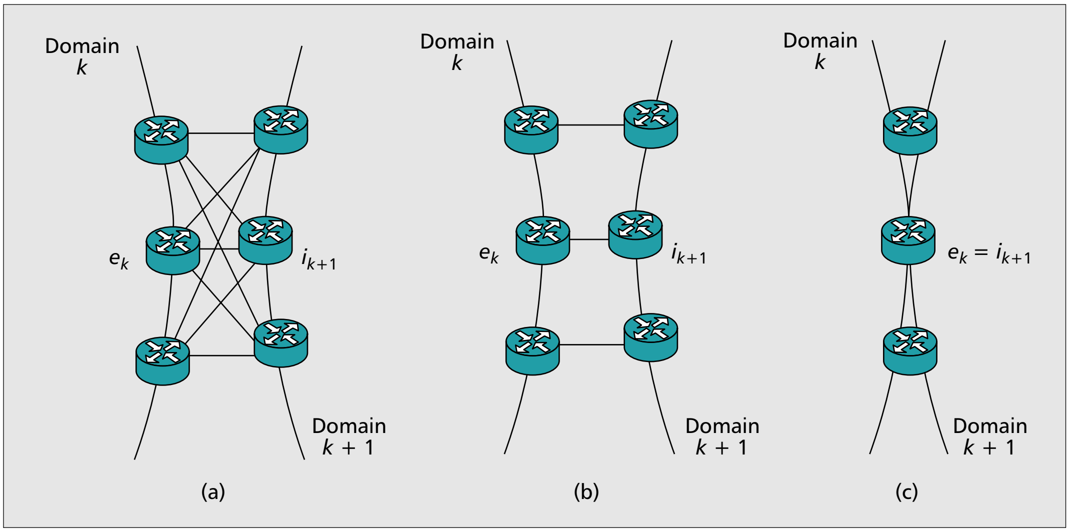

Figure 3. *Interdomain BN-BN linkage patterns: a) full mesh; b) 1:1; c) partnered BN.*

We remark that while BN selection has a critical role in the performance of the combined schemes (PPRO/ARO-based), the exact linkage pattern between the selected BN (ref. to Fig. 3) has a marginal impact. In fact, once the egress BN for domain k has been selected (say e_k), the choice of ingress BN for domain $k + 1$ (denoted i_{k+1}) is restricted to the nodes directly connected to e_k. This choice might take advantage of any information the BNs of domain $k + 1$ have disseminated to BN e_k about the internal connectivity of domain $k + 1$ and/or the availability of paths toward d. Without such information, the choice made by e_k is blind, and stronger linkage (e.g., full mesh, Fig. 3a) provides no improvement over a 1:1 interconnection. On the other hand, if some information were distributed between the two domains, it would also be available to the ingress BN of domain k — recall from an earlier section that we assume full sharing of information within each domain — and therefore would have been exploited at the previous stage. This also applies to the case of a partnered BN, as in Fig. 3c. In summary, from the perspective of the combined schemes (PPRO/ARO-based) the BN-to-BN linkage pattern has marginal impact. On the other hand, since it involves changes to the global network topologies, the reference global optimum will be different in each case.

RESULTS

In our simulations we tested all possible permutations of the four ISP maps in a chain of $K = 3$ domains, for a total of 24 scenarios. For each scenario we ran 500 experiments by randomly picking a pair of source and destination nodes in the first and last domains among the INs with degree greater than one. Each time we computed the globally optimal disjoint path pair with a centralized algorithm (Suurballe) assuming full knowledge of the whole network topology. If a solution was found, we ran the ARO-based and PPRO-based algorithms between the given end-points. For each scheme, if a solution was found (no crankback being implemented) we recorded it as a success along with the total cost of the path pair.

ISP	AS number	No. of nodes	Nodes with degree > 1
Abovenet (U.S.)	6461	368	331
EBONE (Europe)	1755	161	134
Tiscali (Europe)	3257	248	266
Exodus (U.S.)	3967	246	214

Table 1. *ISP topologies used in the simulations. (found at http://www.cs.washington.edu/research/networking/rocketfuel/).*

In Fig. 4 we present the success rates obtained in all the scenarios. The ARO-based scheme succeeds in all cases, while the PPRO-based performs considerably worse: the success rate is mostly around 90 percent, but in some cases dramatically lower. Further inspection displayed some specific topological configuration that repeatedly trapped the PPRO-based algorithm. This occurs, for example, when the ingress and/or egress BNs in domain k lies on the shortest path between the neighboring domains $k - 1$ and $k + 1$, or when trap topologies like Fig. 2 are embedded within one domain.

The most important result here is that the ARO-based scheme seems to be very robust to such specific configurations. We remark that the empirical success rate of 100 percent for ARO-based does not imply that it will be successful in the general case: while it is certainly possible to draw topological configurations that trap the ARO-based scheme, it appears that such configurations were not present in the set of topologies considered.

Besides success rate, a second performance metric is the total cost of the path pair. Since all link costs are set to unity, this is equivalent to the total path length. In Fig. 5 we reported the total cost for each experiment with the ARO- and PPRO-based schemes (y axis) vs. the minimal cost (x axis) for those experiments where both succeed. As expected, in the case where $N = 2$ the ARO always achieves the optimal solu-

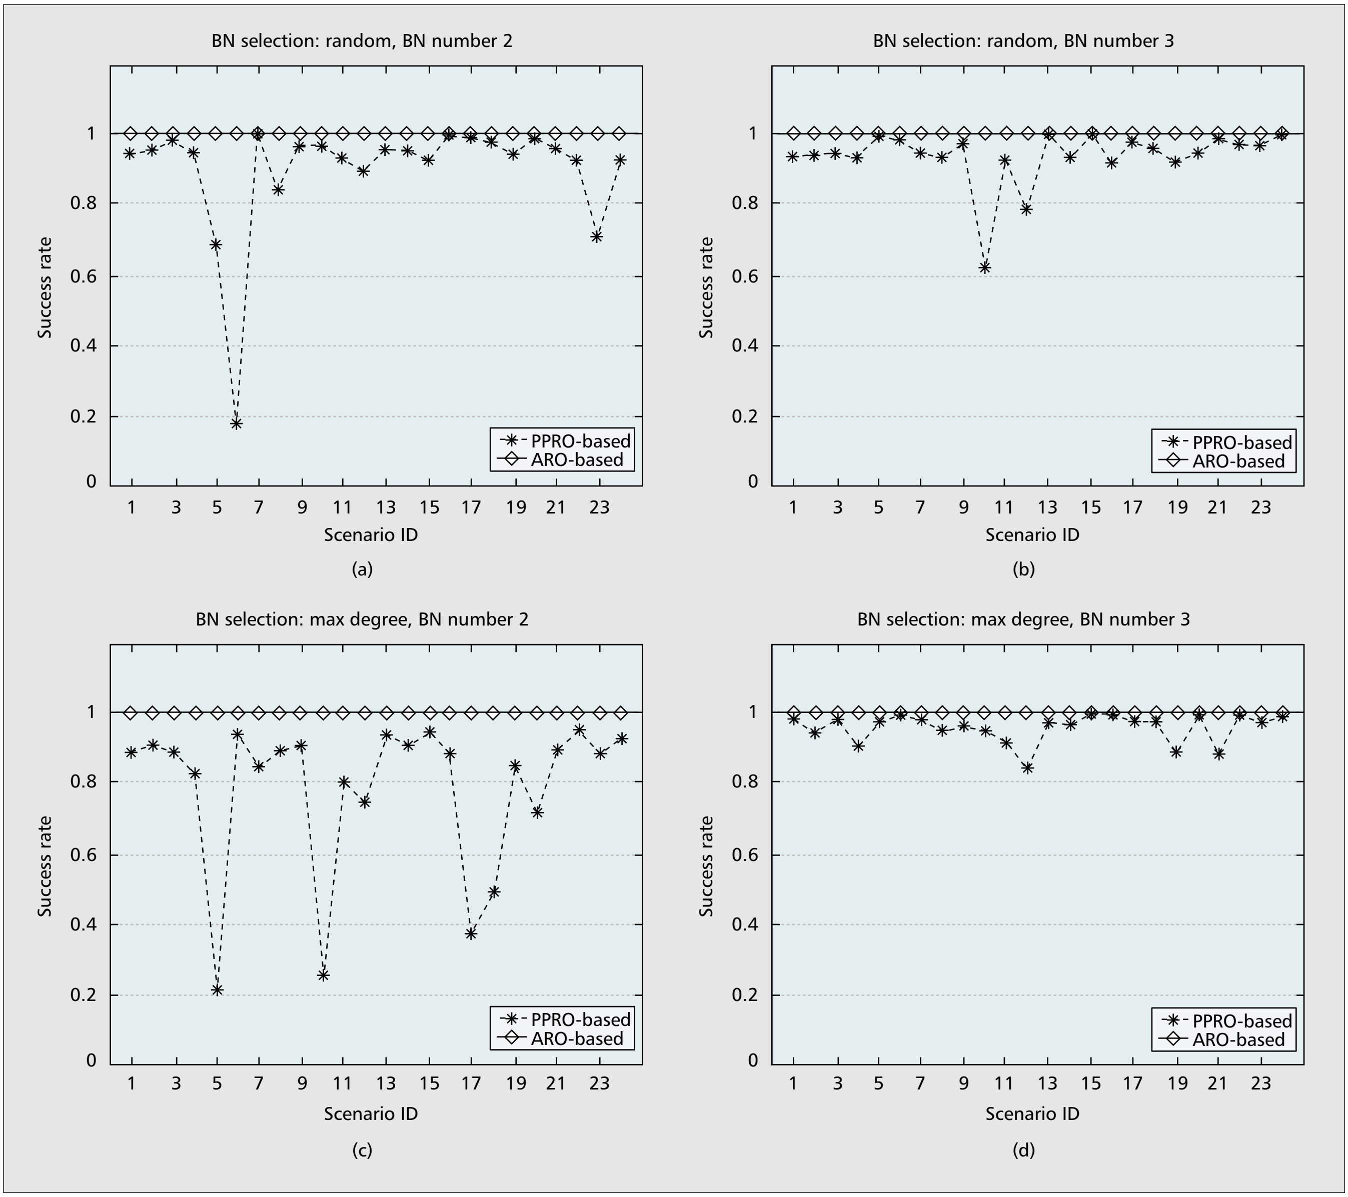

Figure 4. *Simulation result: success rate, ARO-based and PPRO-based. Selection of BN: a, b) random; c, d) max-degree. Number of BNs: a, c)* N = *2; b, d)* N = *3.*

tion. In fact, at each stage the ingress/egress BNs are fixed, and within each domain the ARO-based scheme finds the locally optimal segment pair. Moving to $N = 3$, local optimality is not sufficient for global optimality since there is no guarantee that the ingress/egress BNs were selected optimally (as discussed earlier). However, in most cases the total cost is optimal or within a few hops from the optimum. The comparison between ARO- and PPRO-based (left vs. right column) shows a larger vertical spread of the latter, which indicates slightly worse performance. A second set of simulations with randomly assigned link costs displayed qualitatively similar results.

CONCLUSIONS

Path diversity might contribute to satisfying some QoS requirements since it enables path protection and load balancing. This is particularly important for some high-value traffic components like ToIP and VPNs, which often have an interdomain scope. Therefore, interdomain path diversity could play a role in promoting future interdomain QoS and is among the requirements for interdomain TE.

In this work we have provided an overview of three different schemes for dynamic setup of disjoint LSPs in a multidomain environment. The first is based on the PCE architecture and guarantees a globally optimal solution with separate computation and installation phases. The other two schemes do not require architectural extensions, just one additional object in the RSVP-TE messages. However, these do not guarantee the optimal solution.

Our experiments indicate that the ARO-based scheme performs better than the PPRO-based scheme. The latter was trapped with nonnegligible frequency, while the former never failed and in most cases incurred a total cost

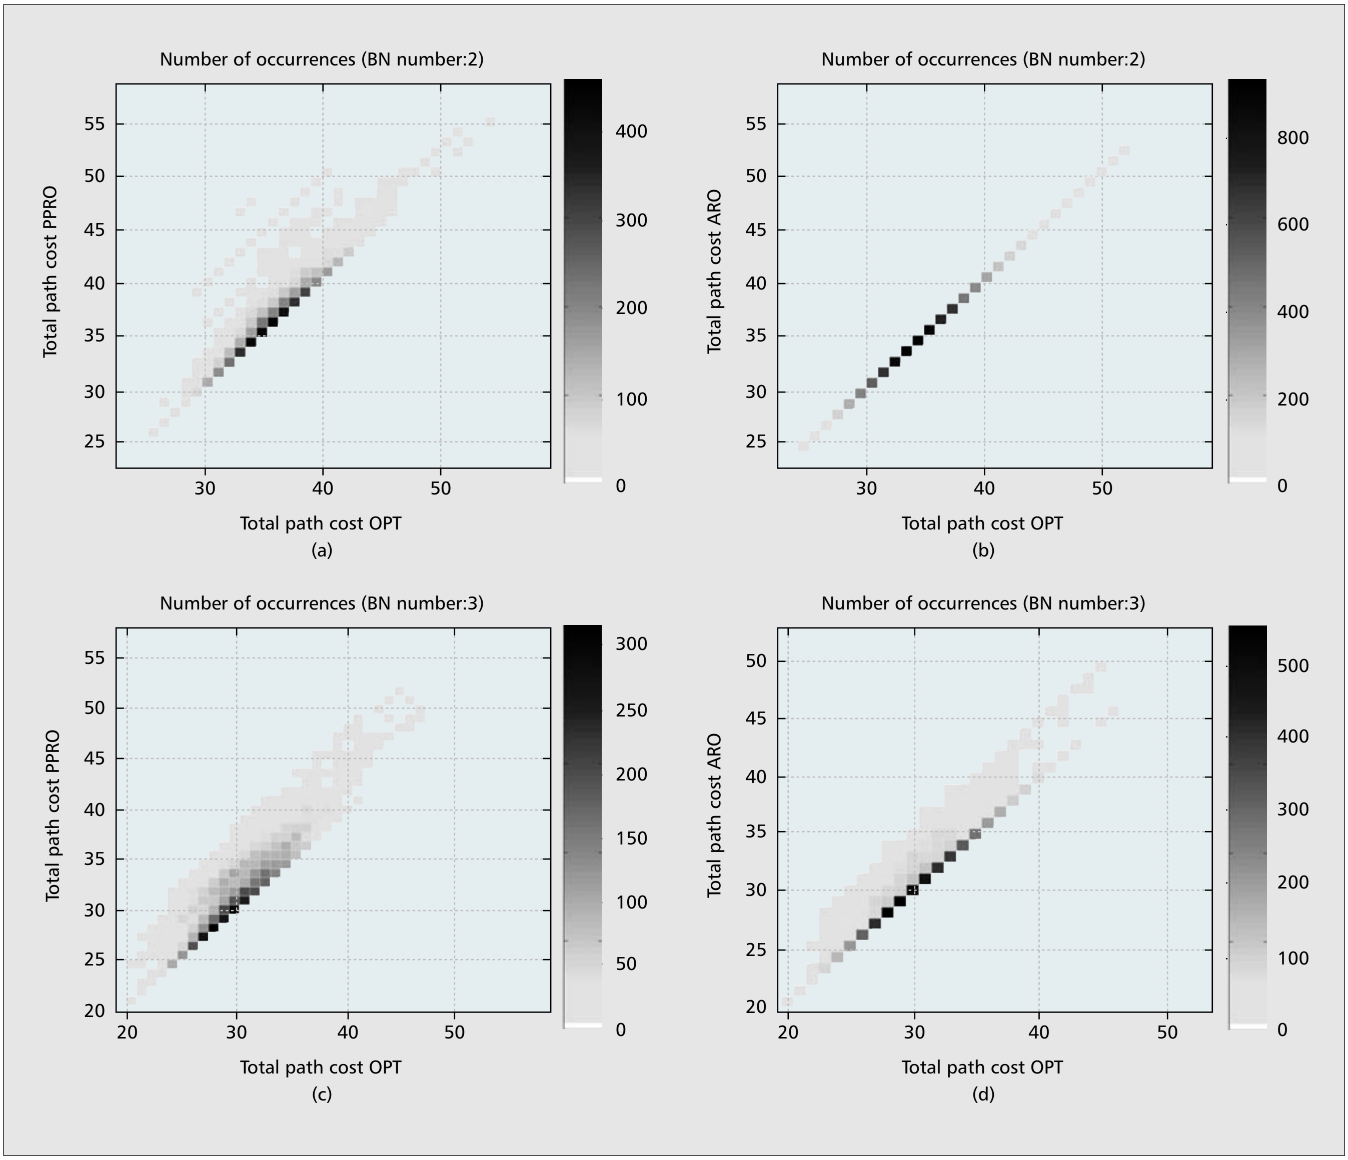

Figure 5. *Simulation results: a, c) total cost for ARO-based; b, d) PPRO-based vs. optimal cost. Selection of BN: random. Number of BNs: a, b)* N = *2; c, d)* N = *3.*

very close to the optimum. But some caution is due before attempting generalizations. While it is certainly possible to design topological configurations that trap the ARO-based scheme or lead to severely suboptimal solutions, it appears that such configurations were not present in the set of topologies we considered. In this respect, we believe that domain size is a critical dimension: the larger the domain, the more degrees of freedom are available in the path computation and the less frequent are "pathological" configurations affecting the performance of the ARO-based scheme. Our future experiments will target a broader set of ISP maps, particularly of smaller size.

These results provide evidence that the ARO-based scheme should be preferred over the PPRO-based. From the perspective of diverse path computation performance, the former was comparable to the PCE-based optimal computation in the scenarios we considered. On the other hand, the "think-before-make" strategy enabled by the latter undoubtedly offers a higher degree of flexibility for implementing any path computation scheme. Such flexibility, which cannot be achieved by a combined computation/installation, might be decisive when considering other applications of distributed path computation besides path diversity.

REFERENCES

[1] D. Fratianni *et al.*, "Il Backbone IP per i Servizi Telefonici," *Notiziario Tecnico Telecom Italia*, anno 13, no. 1, June 2004 (in Italian).

[2] A. M. Langellotti *et al.*, "Il Backbone IP di Telecom Italia Wireline," *Notiziario Tecnico Telecom Italia*, anno 13, no. 2, Dec. 2004 (in Italian).

[3] J. L. Le Roux, J. P. Vasseur, and J. Boyle, "Requirements for Inter-Area MPLS Traffic Engineering," Internet draft, draft-ietf-tewg-interarea-mpls-te-req-03.txt, Nov. 2004, work in progress.

[4] A. Farrel, J. P. Vasseur, and A. Ayyangar, "A Framework for Interdomain MPLS Traffic Engineering," Internet draft, draft-ietf-ccamp-inter-domain-framework-01.txt, Feb. 2005, work in progress.

[5] A. Farrel, J. P. Vasseur, and J. Ash, “Path Computation Element (PCE) Architecture,” Internet draft, draft-ash-pce-architecture-01.txt, Feb. 2005, work in progress.

[6] Awduche *et al.*, “Extensions to RSVP for LSP Tunnels,” RFC 3209, Dec. 2001.

[7] J. P. Vasseur, A. Ayyangar, and R. Zhang, “Interdomain Traffic Engineering LSP Path Computation Methods,” Internet draft, draft-vasseur-ccamp-inter-domain-path-comp-00.txt, July 2004, work in progress.

[8] J. P. Lang, Y. Rekter, and D. Papadimitriou, “RSVP-TE Extensions in Support of End-to-End Generalized Multiprotocol Label Switching (GMPLS)-Based Recovery,” Internet draft, draft-ietf-ccamp-gmpls-recovery-e2e-signaling-02.txt, Oct. 2004, work in progress.

[9] C. Y. Lee, A. Farrel, and S. De Cnodder, “Exclude Routes — Extension to RSVP-TE,” Internet draft, draft-ietf-ccamp-rsvp-te-exclude-route-03.txt, Feb. 2005, work in progress.

[10] B. Doverspike, G. Li, and C. Kalmanek, “Fiber Span Protection in Mesh Optical Networks,” *Opt. Net.*, vol. 3, no. 3, May/June 2002.

[11] F. Ricciato, U. Monaco, and A. D’Achille, “A Novel Scheme for End-to-End Protection in a Multi-Area Network,” IPS ’04, Budapest, Hungary, Mar. 2004.

[12] D’Achille *et al.*, “Diverse Inter-Region Path Setup/Establishment,” Internet draft, draft-dachille-diverse-inter-region-path-setup-01.txt, Oct. 2004, work in progress.

[13] J. W. Suurballe, “Disjoint Paths in a Network,” *Networks*, 1974, pp. 125–45.

[14] A. Farrel *et al.*, “Crankback Signaling Extension for MPLS Signaling,” Internet draft, draft-ietf-ccamp-crankback-04.txt, Feb. 2005, work in progress.

BIOGRAPHIES

FABIO RICCIATO (ricciato@ftw.at) received a Laurea degree in electrical engineering and a Ph.D. in telecommunications from the University of Rome La Sapienza in 1999 and 2003, respectively. He has collaborated with CoRiTeL in several national projects and one EU project (IST-AQUILA). Since 2004 he is a senior researcher at the Forschungszentrum Telekommunikation Wien (FTW), where he leads a research project on traffic measurements in 3G networks. His research interests are in the networking area for wireless and backbone technologies, including QoS, traffic engineering, resilience, traffic measurement and analysis, and network security.

UGO MONACO received a Laurea degree in telecommunications engineering from the University of Rome La Sapienza, Italy, in 2002. Currently he is a Ph.D. student in the Networking Group of the INFOCOM department at the same university. His main research interests are in IP/MPLS intra-/interdomain routing, network optimization techniques, and traffic engineering.

DANIELE ALI received a Laurea degree in telecommunications engineering from the University of Roma La Sapienza. He was involved in the implementation of a MPLS-TE Linux-based testbed and worked on interdomain routing performance evaluation. Currently he is a collaborator of the University of Rome Tor Vergata in the European project Simplicity.

IEEE COMMUNICATIONS MAGAZINE
CALL FOR PAPERS
WIDEBAND SPEECH CODING STANDARDS AND WIRELESS SERVICES

Background

Compared to conventional narrowband telephony, the introduction of a wider audio bandwidth of 50-7000 Hz provides substantially improved speech quality and naturalness and adds also a feeling of transparent communication. Efficient wideband speech coding algorithms have recently emerged and several wideband speech coding standards have been prepared in ITU-T, 3GPP and 3GPP2 for various applications, including wideband telephony, VoIP, video-conferencing, multimedia services (such as messaging and streaming), and Internet applications, among others. The use of wideband speech services has been enabled by the current end-to-end digital networks (e.g., wireless, ATM, xDSL) through the use of Packet Switching (PS) or Tandem/Transcoder Free Operation (TFO/TrFO). Due to the introduction of efficient wideband speech coding algorithms and networks able to support their employment, there is currently an increased interest by operators (both wireless and wireline) for adopting wideband speech services for the above mentioned applications, which are becoming increasingly important. During the next years, the most rapid employment of wideband speech can be expected in wireless applications such as in 3rd generation mobile phone systems, but it is not limited to these.

The goal of the issue is to address the importance and advantages of wideband speech telephony and services, and discuss the recent advances in wideband speech coding technology - in particular to explain the recent wideband coding standards developed in ITU-T, 3GPP, and 3GPP2 and their applications. Further, the issue may address related service issues, such as wideband terminal characteristics, and assessment of wideband speech quality in terminals and networks.

Papers are invited especially in the following topics:
- Wideband speech coding standards and algorithms
- Wideband speech quality and quality assessment
- Wideband speech applications and services
- Employment of wideband speech into terminals and networks (e.g., wideband terminal characteristics)

Manuscripts should be submitted through Manuscript Central at: http://commag-ieee.manuscriptcentral.com/. (Choose your topic from the “Topic or Series” drop down menu.) Authors must follow the IEEE Communications Magazine guidelines regarding the manuscript format. For further details, please find information for authors at the IEEE Communications Magazine website at http://www.comsoc.org/pubs/commag/sub_guidelines.html.

Schedule of Submissions

Manuscripts due:	July 15, 2005
Acceptance notification:	December 31, 2005
Final manuscripts due:	February 28, 2006
Publication:	May 2006

Guest Editors

Kari Järvinen
Nokia Research Center
e-mail: kari.ju.jarvinen@nokia.com

Simão Ferraz de Campos Neto
International Telecommunications Union
e-mail: simao.campos@itu.int