A Combined Intra-Domain and Inter-Domain QoS Routing Model for Optical Networks

Marcelo Yannuzzi, Sergi Sánchez-López, Xavi Masip-Bruin,
Josep Solé-Pareta, Jordi-Domingo-Pascual

***Abstract*—Inter-Domain Quality of Service (QoS) Routing has become a strong requirement in the present Internet, and this requirement will also be present in the Next Generation Optical based worldwide network. At present end-to-end QoS Routing (QoSR) represents a complex problem mainly because the de-facto standard Inter-domain routing protocol, namely the Border Gateway Protocol (BGP) has not inbuilt QoSR capabilities. Moreover, BGP entirely obscures the availability of Intra-Domain resources in any transit domain within an end-to-end Inter-Domain path, which shifts any tentative proposal to cope with the issue of Inter-Domain QoSR even farther from optimality. Given that Inter-Domain routing in Optical Networks is an active research area in this moment, it seems wise to address the issue of QoSR provisioning from its very foundations. Thus, in this paper we introduce a Combined Intra-Domain and Inter-Domain QoSR Model for Optical Networks. Our goal is to provide a highly efficient coupling between both routing schemes with the aim that the combined QoSR model could be able to supply multi-constrained end-to-end optical paths closer to optimality.**

***Index Terms*—Inter-Domain, Intra-Domain, QoS Routing, Optical Networks**

I. INTRODUCTION

THE Border Gateway Protocol (BGP) is currently the de-facto standard Inter-Domain Routing Protocol in the Internet. Its current release is BGP-4, which was specified in RFC 1771 on March of 1995 [1]. Throughout these nine years, the number of Autonomous Systems connected to the Internet has augmented enormously, which accordingly increased the demands on the scale of the network. In spite of this burden, BGP has proven to be a resilient routing protocol. Among the strengths that made BGP become so popular are firstly that it was designed to address the issue of scalability and connectivity demands at an Internet scale. Secondly, it provides stability to the biggest network ever deployed, and thirdly, BGP was endowed with policy based routing features, allowing each administrative domain at the edge of a BGP connection to manage its inbound and outbound traffic according to its specific preferences and needs.

Although these significant strengths, BGP also presents several weaknesses. For instance, in many cases BGP requires tens of minutes to recover from a route or a link failure [2]. Moreover, even though BGP allows an Autonomous System (AS) to flexibly manage its outbound traffic, it exhibits a scarce degree of control in order to manage and balance how traffic enters the AS across multiple possible paths. In addition, each BGP router only advertises the best route it knows to any given destination prefix. This implies that many alternative paths that could have been potentially used by any source of traffic will be unknown because of this pruning behavior inherent to BGP. Finally, BGP completely hides the state of intra-domain resources within every AS.

In sum, the current release of BGP supplies a slow reacting and limited routing protocol, which results inadequate to handle most of the emerging demands for inter-domain functionalities. The justification for this behavior is that BGP was intrinsically designed to address overall stability and scalability instead of concerning about issues like fast recovering from a particular link failure, nor bounding delay or the packet loss ratio across the Internet for a given block of prefixes, just to name a few.

At present, several efforts are being carried to address the issue of QoS Routing (QoSR) at an inter-domain level in IP networks. On the one hand, many researchers and manufacturers are trying to enhance BGP with new capabilities such as Traffic Engineering (TE), and QoS extensions mainly because of the ubiquity and success that BGP presents at this moment. [3-6] are very good examples of

Manuscript received on December 13, 2004. This work was partially funded by the MCyT (Spanish Ministry of Science and Technology) under contract FEDER-TIC2002-04344-C02-02 and the CIRIT (Catalan Research Council) under contract 2001-SGR00226.

M. Yannuzzi, S. Sánchez-López, X. Masip-Bruin, J. Solé-Pareta, and J. Domingo-Pascual are with the Department of Computer Architecture and the Advanced Broadband Communications Laboratory (CCABA), Universitat Politècnica de Catalunya (UPC), 08034 Barcelona, Catalonia, Spain (email addresses: yannuzzi@ac.upc.edu; xmasip@ac.upc.edu; sergio@ac.upc.edu; pareta@ac.upc.edu; jordid@ac.upc.edu).

this kind of approach. It is important to notice that even though the non-extended version of BGP presents limited functionality, it is indeed a complex routing protocol, where mistakes and misconfigurations are not infrequent. Moreover, some research groups have extended BGP with layer 2 and layer 3 Virtual Private Network (VPN) discovery and signaling capabilities, within the new Multi-Protocol BGP (MP-BGP). Consequently, these proposals for enhancing BGP not only tend to turn it into a much more complex protocol, but also it remains to be seen if the addition of all these enhancements in a real environment could not overwhelm the protocol.

Despite the well-known limitations of BGP in the areas of QoSR and TE, during the past few years some researchers have proposed to adopt an Optical Border Gateway Protocol (OBGP) as the future inter-domain routing protocol for optical networks [7, 8]. These proposals basically attempt to enhance the legacy and limited release of BGP so that it can carry and update optical path information among OBGP neighbors. Latterly, the proposals were not appealing to the Internet Engineering Task Force (IETF) community, but there are still ongoing research efforts in this area.

Rather than enhancing BGP, an alternative approach to inter-domain routing appeared in 2002 when the Optical Internetworking Forum (OIF) proposed the Domain-to-Domain Routing Protocol (DDRP) [9]. This routing protocol is based on a link state protocol, to be precise, OSPF-TE. However, link state protocols are known to be inadequate for inter-domain routing due to their scalability limitations. Thus, the OIF proposed to modify OSPF-TE and turn it into a hierarchical link state protocol, so DDRP actually uses a modified Dijkstra algorithm. DDRP has mainly two drawbacks. Firstly, it represents a major change in terms of routing and service provisioning when compared with the current IP based Internet, since it proposes to move toward a fully hierarchical model. Therefore, these kinds of proposals are not precisely aligned with several of the premises within the IETF community. Secondly, the modified Dijkstra algorithm still presents very limited flexibility and functionality in the areas of inter-domain QoSR and TE. For instance, this modified algorithm returns a single optimal path at a time so complementary algorithms should be adopted for diverse routing computation and path protection.

An appealing option given its remarkable flexibility in terms of inter-domain QoSR and TE is the overlay approach, which has become a strong candidate to address these issues [10-14]. The main idea behind the overlay is to decouple part of the policy control portion of the routing process from the BGP devices. In any case, it is important to keep in mind that at present the only way to engineer inter-domain traffic in IP networks is by means of smartly configuring BGP. In this sense, the BGP approach differs from the overlay approach in how policies are controlled and signaled. While the BGP enhancements tend to provide in-band signaling, the overlay approach provides out-of-band signaling. However, at the end both mechanisms rely on appropriately tuning BGP to comply with their respective traffic policies.

Given the present limitations of inter-domain routing in IP networks it seems sound to address the issues of QoSR and TE provisioning in the future Optical Networks from its very foundations. Neither OBGP nor DDRP timely address these issues. Therefore, in this paper we introduce a Combined Intra-Domain and Inter-Domain QoSR model for Optical Networks. Our goal is to provide a highly efficient coupling between both routing protocols with the aim that the combined QoSR model could be able to supply multi-constrained end-to-end optical paths closer to optimality. Furthermore, the combined model will not tightly depend on either of these routing protocols. In this sense, enhancements or even a complete replacement of either of the protocols will not tear down the combined routing model. The key to accomplish these goals is a special device that we introduce in this paper called Inter-Domain Routing Agent (IDRA).

The rest of the paper is organized as follows. Section II presents an up-to-date comparison between the most compelling models proposed for optical networks. Section III introduces our novel combined routing model, while Section IV highlights our future research work and concludes the paper.

II. OPTICAL NETWORK MODELS

A. *Overlay versus Peer-to-Peer Model*

What has changed in terms of the network business model? Perhaps the answer to this open question is that many customers are willing to no longer purchase monthly or yearly contracts for one-time capacity, instead Network Service Providers (NSPs) should be able to sell end-to-end lightpaths with the capacity required for periods of hours, days or even months. Those lightpaths could be configured similarly as Frame-Relay or ATM circuits, in which the customer pays for a certain amount of basic transport, but then it can get higher/lower capacities as necessary. A major requirement is that customers should be able to acquire or release those end-to-end lightpaths on demand and in real-time, so the setup/tear-down of those lightpaths should be solved dynamically. This requirement imposes several demands on the future optical networks because establishing/releasing this kind of paths requires not only that a customer could be able to purchase capacity from its NSP, but also that a NSP could be able to request/purchase on demand the additional capacity needed to accommodate large traffic flows from other NSPs. Thus, one of the most important open issues nowadays is the following; what should be the architecture for the optical Internet?

At present two models are proposed to cope with this problem, namely, the Overlay and the Peer-to-Peer models.

Overlay Model: This model calls for maintaining two discrete networks, namely, an optical network and the customer network. In this model, customer routers or switches request the optical network for a connection, and the optical network either grants or denies it. These requests can be sophisticated, asking for example for a certain circuit size but with a particular grade of restoration. The key in this model is that customer devices cannot "see" inside the NSPs networks. To request capacity from the underlying network, customers access the optical network through a User Network Interface (UNI), which hides the complexity of the NSP network from them. The UNI allows customers of an optical network to establish optical connections dynamically across the optical network, using a neighbor-discovery mechanism and a service-discovery mechanism. On the other hand, devices in the NSP network rely on Network-to-Network Interfaces (NNIs) to access critical network information within their own network, and on External Network-to-Network Interfaces (ENNIs) to access critical network information outside their network. The NNI and ENNI are the control interfaces over which the optical network connections are accomplished, involving basically lightpath routing and signaling. At this point it should become clear that in the Overlay model the business is firmly controlled by the NSPs, so legacy voice operators tend to support this model. Figure 1 depicts the Overlay architecture and routing model for the intra-domain case.

Peer-to-Peer Model: The Peer-to-Peer model argues for a single network in which devices at the edge of the network decide how bandwidth is allocated within the optical core. The separation between the customer and the underlying optical networks becomes blurred. In the Peer-to-Peer model, optical switches and customer routers act as peers, using a uniform and unified control plane to establish lightpaths across the network, and with complete knowledge of network resources. In this model there is little or no distinction among UNI/NNI, or NNI/ENNI. All network elements are direct peers and fully aware of the topology and resources. In the Peer-to-Peer model the complexity is pushed to the edge of the network. However, a significant amount of state and control information flows between the IP and optical layers, making the development of this model more complex. Thus, this model presents a scalability problem because of the amount of information that needs to be handled by any given network element within a NSP. Clearly, in this model the business is no longer tightly controlled by the NSPs, which is one of the main reasons why it remains to be seen if the model will eventually gain enough acceptance among the NSPs. Figure 2 depicts de Peer-to-Peer model and routing architecture for the intra-domain case.

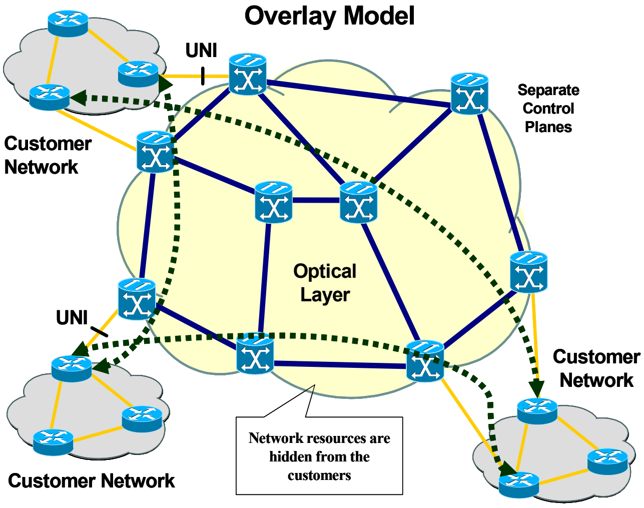

Fig. 1. Overlay Architecture for the intra-domain case

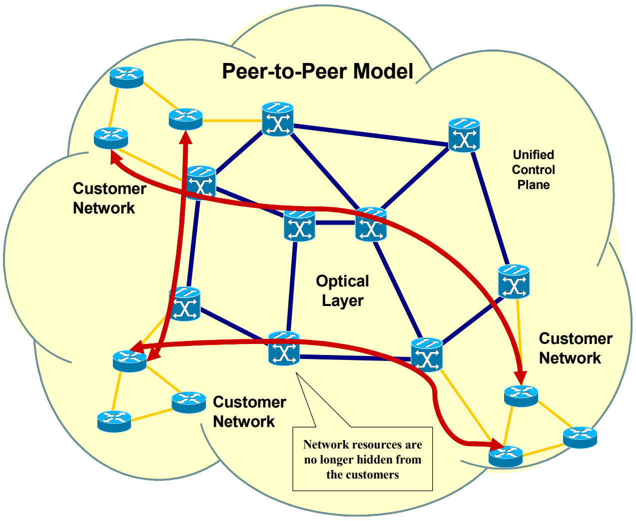

Fig. 2. Peer-to-Peer Architecture for the intra-domain case

B. *An Hybrid Inter-Domain Model*

The NSP-to-NSP relationships could be modeled exactly in the same way as stated before. On the one hand, some NSPs are customers from a higher hierarchy NSP set, so in those cases the Overlay architecture perfectly models the customer/provider relationships. On the other hand, several NSPs share peering circuits with other NSPs at the same hierarchy, thus these peer relationships could be perfectly modeled by the Peer-to-Peer model. It seems quite reasonable to foresee that a hybrid Overlay/Peer-to-Peer model will probably suite for the best the future optical Internet. Obviously, substantial Peer-to-Peer connectivity may potentially provide unparalleled resiliency and traffic load balance, where each domain provides transit and services to its peers [15]. The necessity of Peer-to-Peer structures as part of the inter-domain routing model becomes clear when we consider applications such as GRID projects, in which bulk data files need to be transferred to a number of sites in a bursty way. In these cases, the ability to dynamically change the bandwidth and the topology without re-signaling becomes essential. Furthermore, inter-domain Peer-to-Peer relationships may also become attractive in terms of monetary costs when some domains exchange large amounts of traffic.

The central issue in the Peer-to-Peer model is scalability, so the hybrid model arises as an attractive trade-off between both models. From the Overlay perspective the hybrid model scales supporting a hierarchy of multiple administrative domains. From the Peer-to-Peer perspective the hybrid model supports that multiple optical network domains equally control the circuits among them without any centralized control, and provide transit services to each other in a much more open way, according to its possibilities and requirements.

The Combined routing model we are proposing in this paper is a hybrid routing model, since we believe that this mixed scenario is more likely to become deployed. Our model supports some dense inter-domain Peer-to-Peer cores, while it also allows exploiting the advantages of the Overlay model towards the edge of the optical network where NSPs of a less significant hierarchy are likely to become deployed. This hybrid model has two significant strengths: i) it is indeed a scalable model; ii) the diameter of the network remains bounded by the Peer-to-Peer relationships within the core of the optical network. Clearly, these strengths in our model are of utter importance when trying to address the issues of end-to-end QoS routing and signaling.

III. COMBINED INTRA AND INTER-DOMAIN ROUTING MODEL

The Combined routing model we are purposing in this paper is based on the introduction of specialized devices called Inter-Domain Routing Agents (IDRAs). These devices will act, among other things, as the glue between the inter-domain routing and the intra-domain routing schemes. The IDRAs are able to enrich the network advertisements with combined intra-domain and inter-domain routing information targeting any neighboring NSP, which blurs the current gap between the intra-domain and inter-domain routing protocols. These new kind of combined advertisements will allow any upstream NSP to choose the next downstream NSP for any given destination not only based on the inter-domain state of the network, but also based on the availability of intra-domain network resources of the NSPs within the alternative paths to reach that destination. In order to provide a highly scalable routing model, the NSP-to-NSP information exchanged must be precisely synthesized, which is also one of the tasks carried out by the IDRAs. In this sense, the cost metrics associated with any given destination known are advertised as composite cost metrics by the IDRAs, wherein one part of the cost depends on the state of intra-domain resources, and the other on the state of inter-domain resources. Therefore, the Combined Intra-Domain and Inter-Domain QoS Routing protocol runs between the IDRAs. Each AS participating in our combined routing model must have at least one IDRA, and for resiliency and scalability reasons clusters of IDRAs are allowed per each domain. The IDRAs may be directly connected by physical connections or logically connected, and those connections could be point-to-point or point-to-multipoint.

Any pair of peering IDRAs could establish two kinds of relationships: i) a customer/provider relationship, which follows the Overlay paradigm, and this is the kind of relationships that we expect to find between domains scattering from the optical cores towards the edge of the network; ii) a peer relationship, which follows the Peer-to-Peer paradigm, and this is the kind of relationships that we expect to find between domains that are part of the optical cores of the network.

The IDRAs are responsible then for carrying inter-domain routing information, and deciding within each NSP which path is the best, among the available paths to reach any known destination. This decision process is affected not only by the state of intra-domain resources within the traversed domains in those available paths, but also by the state of local intra-domain resources within the source domain. Figure 3 shows our Peer-to-Peer inter-domain routing model for a core optical network, and the IDRAs as part of the different network domains. These IDRAs basically exchange the set of offered services, in addition to the combined cost to reach any destination known. In the framework of our model the total number of services foreseen is large, so each IDRA is endowed with flexible input/output policy-based filtering capabilities so that each NSP may use and advertise to its neighboring domains only a sub-set of those services. Those services will be associated with the destinations advertised by each NSP. For instance, a NSP may advertise that it offers strong diversity features for a particular set of destinations, while it offers poor or non diversity at all for the rest of the advertised destinations [16]. Our model is widely flexible in this particular point, given that a NSP may utilize just a single set of services associated with all the destinations advertised, or it may it may advertise different sets of services for different sets of destinations as in the example above.

Furthermore, as aforementioned synthesizing the information exchanged between the IDRAs is mandatory in order to design a highly scalable QoSR model. Our approach is that QoS state information is carefully exchanged by means of a QoS composite cost metric. This QoS cost metric will then act as an input to the routing algorithm running on each IDRA, and will be used by the NSPs to select the best path for any given destination. Figure 4 shows the flow of advertisements between a destination domain D and a source domain S, which are comprised of a filtered set of the offered services, in addition to the combined QoS cost metrics associated with each particular destination advertised. The figure shows that usually the set of destinations, and their corresponding offered services received at any upstream NSP may be a condensed sub-set from those announced by the source of the advertisements due to the filtering process within each intermediate domain in the path. This filtering process allows each NSP to advertise what it wants to advertise, so disclosing its internal network capabilities and availability of resources is subject to its own policies and filtering criteria. In other words, the external visibility of the NSP's network is controlled by the NSP.

In terms of the QoS composite cost metric, each NSP will compute the following cost to reach any destination known:

$$C(i) = \min\{\, w_A(i) C_{IA}(i) + [C_{ID}(i) + C(i+1)]\, w_D(i) \,\} \qquad (1)$$

i: represents a NSP domain in the path from the source domain S to the destination domain D

$C_{IA}(i)$: represents the intra-domain QoS cost, which is locally computed by the IDRA within the NSP domain i

$C_{ID}(i)$: represents the inter-domain QoS cost, which is locally computed by the IDRA within the NSP domain i

$C(i+1)$: corresponds to the inter-domain QoS cost advertised by the IDRA within the NSP domain $(i+1)$

w_A: weight of the intra-domain cost, which is locally assigned by the NSP i

w_D: weight of the inter-domain cost, which is locally assigned by the NSP i

Equation (1) is restricted to the following border condition at the destination D:

$$C(D) = C_{IA}(D) \qquad (2)$$

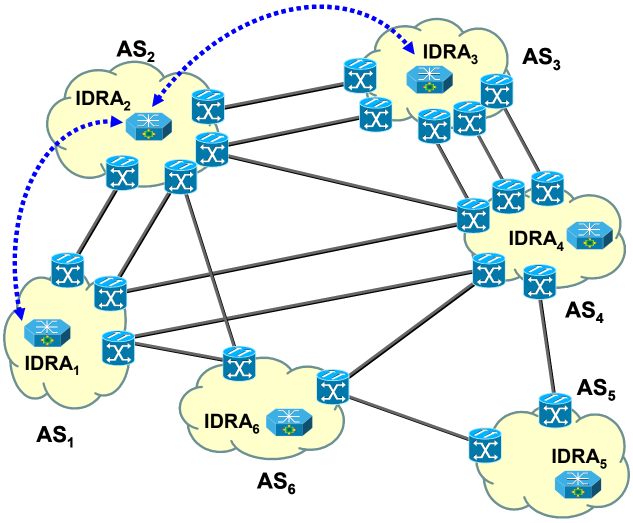

Fig. 3. Peer-to-Peer Inter-Domain Routing Model for a core Optical Network based on Inter-Domain Routing Agents (IDRA) which are able to exchange the set of offered services, and the combined cost to reach any destination known

It is worth noticing that if the adjacent domains i and $(i+1)$ are connected through several links, then the calculus in (1) implies to minimize the cost not only over the available intra-domain paths within i able to reach a given destination, but also over the available egress links connecting both domains.

In our model each NSP may freely choose which Interior Gateway Protocol (IGP) should be used within its domain. However, while computing $C_{IA}(i)$ each IDRA must have a common framework to determine this value independently of the IGP employed. Thus, our model requires that a standardized mapping exists between the different IGPs and the calculus of $C_{IA}(i)$. The main advantage is that our routing model is able to take into account useful intra-domain QoS state information, but without getting into the details of its implementation. Furthermore, the IGP running within a particular NSP could be enhanced or even completely replaced by a more sophisticated one without affecting our

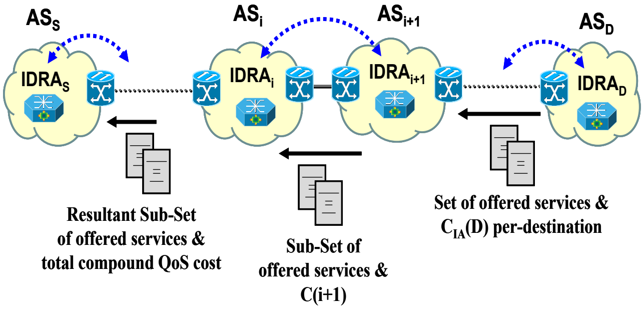

Fig. 4. Inter-Domain Routing advertisements flow between a destination domain D and a source domain S in our combined Intra-Domain and Inter-Domain QoSR model

routing model, as long as the standardized mapping is supported by this new IGP.

From the network advertisements perspective our combined routing model introduces substantial differences when compared with the current inter-domain routing model in the Internet. On the one hand, in our model part of the state of the NSP's internal routing becomes now visible from the outside, which is inherently different from the contemporary inter-domain routing paradigm based on BGP. On the other hand, BGP lacks at present of flexible mechanisms to advertise enriched QoS information such as availability of resources (i.e. availability of wavelength converters), spatial diversity capabilities or monetary costs. These mechanisms are indeed present in our routing model with the aim that the routing algorithms running on the IDRAs could take full advantage of them. The key in supplying these mechanisms is the set of services advertised by the IDRAs, in addition of course to the QoS cost metrics announced. For instance, a QoS routing algorithm running in one particular IDRA could prefer, at some extent, a path with a better end-to-end protection degree rather than simply selecting the best path by using the minimum cost in (1). Therefore, each IDRA in our routing model is capable of determining the best path for any given destination D based on a number of concrete requirements in terms of services, or based on the minimum cost to reach the destination D. In this sense, much more sophisticated heuristics and QoSR algorithms could be empowered by our routing model. Undoubtedly, those heuristics need to be able to efficiently cope with the NP-hard issues present in any Multi-Constrained Path selection problem (MCP).

Among the foremost strengths of our combined routing model are firstly that each NSP may compute (1) and determine that the best end-to-end route to reach a certain destination could be different from the best inter-domain route to get to it. For example, let us assume that the Border Switch Router BSR$_{1S}$ in Figure 5 has a better inter-domain cost than BSR$_{2S}$ to reach a particular destination D. However, if the internal resources in the path between the client network and BSR$_{1S}$ are scarce, while the network resources between the client and BSR$_{2S}$ are quite unloaded, the end-to-end path chosen between the client network and D could be far from optimal if only inter-domain information is taken into account during the decision process. Thus, our combined QoSR model offers stronger end-to-end QoS guarantees than simply relying on an inter-domain QoSR model. Secondly, a main advantage in our routing model is that it is indeed a completely distributed routing architecture based on distributed IDRAs.

Figure 5 depicts the ligthpath setup steps within a source domain S:

(1) The Client requests a path to a Distribution Layer (DL) device
(2) The DL device interrogates the local IDRA in order to find the best BSR for that specific destination
(3) The IDRA responds with the best BSR based on the state of Inter-Domain resources and also the availability of local resources
(4) The DL device forwards the Explicit Path Setup Request via the correspondent BSR.

Another central issue in our model is how QoS routing updates are managed among peering IDRAs in order to provide a highly scalable QoSR model. Henceforth, we will call these QoS routing updates State Advertisements (SAs). Our approach is to handle different levels of profundity while disseminating these SAs throughout the network. For instance let us assume that a lightpath must be established between two NSPs, namely, NSP$_1$ and the NSP$_3$. This is represented by the dotted curve between the IDRA$_1$ and the IDRA$_3$ in Figure 6. Once the lightpath is established, the IDRAs involved in the lightpath setup process trigger SAs to their Primary Level IDRAs that is to say to the neighboring IDRAs that did not took part in the lightpath setup process. This updating information scheme allows that neighboring IDRAs keep up-to-date QoSR information. However, from the Secondary Level IDRAs perspective, the Tertiary Level IDRAs, and so on, the QoSR information will certainly present some inaccuracy. To tackle this problem we allow triggering Secondary Level SAs, but only when substantial changes have occurred in the state of network. In other words, the Primary Level IDRAs trigger

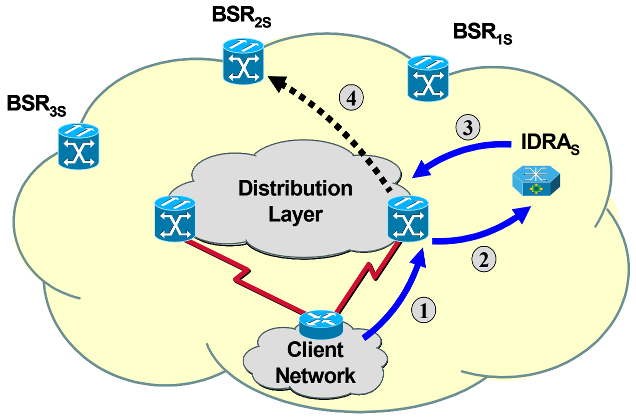

Fig. 5. Ligthpath setup steps within the source domain. (1) Path Request to a DL device; (2) The DL device interrogates the local IDRA in order to find the best BSR for that specific destination; (3) The IDRA responds with the best BSR based on the state of Inter-Domain resources and also the availability of local resources; (4) The DL device forwards the Explicit Path Setup Request via the correspondent BSR.

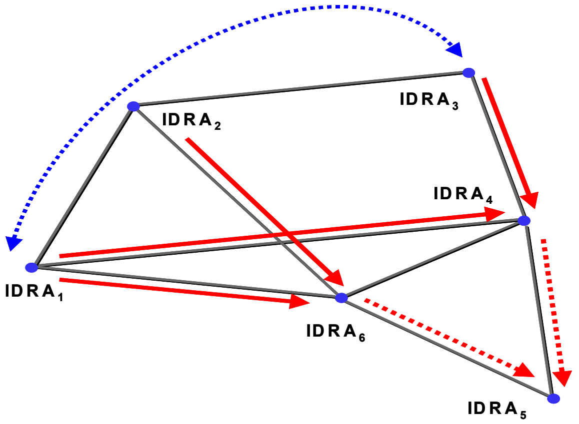

Fig. 6. Lightpath establishment and profundity of QoS SAs. The Solid arrows represent the Primary Level SAs, while the dotted ones represent the Secondary Level SAs which are triggered only when substantial changes have occurred in the state of network

Secondary Level SAs only when a previously defined Secondary Level threshold is met.

Similarly, our model supports managing higher level thresholds in order to accurately control the profundity of the SAs. Figure 6 shows our approach. A justification for this approach is the following. At present, inter-domain traffic characteristics reveal that even though an AS will exchange traffic with most of the Internet, only a small number of ASs is responsible for a large fraction of the existing traffic. Moreover, this traffic is mainly exchanged among ASs that are not directly connected; instead they are generally 2, 3 and 4 hops away [4]. Similarly, in the future optical Internet it is likely that even for money saving reasons, NSPs that share large amounts of traffic try to establish closer client/provider or peer-to-peer relationships in order to avoid the cost of transiting over several NSPs hops. In such cases the number of IDRAs involved in establishing an end-to-end optical path might be diminished. Thus, a profundity of 2, 3 or 4 Levels may be enough to keep most of the relevant remote domains quite up-to-date in terms of QoS routing information. Finally, our model also supplies periodical updates between peering IDRAs.

IV. SUMMARY AND FUTURE RESEARCH WORK

Our goal in this paper was to introduce a combined intra-domain and inter-domain QoS routing model for optical networks. We have shown that our model presents significant advantages empowering QoS based routing in the future optical Internet. Issues like discovery and signaling were left out of the scope of this particular work. These issues are part of ongoing and future research work. Additionally, in the near future we plan to develop and incorporate novel QoS composite cost metrics based on optical QoS parameters as well as novel QoSR algorithms and heuristics using those metrics in order to tackle the MCP problem.

REFERENCES

[1] Y. Rekhter, T. Li, “A Border Gateway Protocol 4 (BGP-4)”, Internet Engineering Task Force, Request for Comments 1771, March 1995.

[2] C. Labovitz, A. Ahuja, A. Bose, and F. Jahanian, "Delayed Internet routing convergence," in Proc. ACM SIGCOMM, 2000, Stockholm, Sweden, September 2000.

[3] L. Xiao, K. Lui, J. Wang, K. Nahrstedt, “QoS Extension to BGP”, IEEE ICNP 2002, Paris, France, November 2002.

[4] B. Quoitin, S. Uhlig, C. Pelsser, L. Swinnen and O. Bonaventure, “Interdomain Traffic Engineering with BGP”, IEEE Communications Magazine, May 2003.

[5] G. Cristallo, C. Jacquenet, “An Approach to Inter-Domain Traffic Engineering”, in Proceedings of XVIII World Telecommunications Congress, Paris, France, September 2002.

[6] B. Quotin, S. Tandel, S. Uhlig, O. Bonaventure, “Interdomain Traffic Engineering with Redistribution Communities”, Computer Communications Journal, March, 2004.

[7] M. Blanchet, F. Parent, B. St-Arnaud, “Optical BGP (OBGP): InterAS Lightpath Provisioning”, Internet Engineering Task Force draft, ietf-draft-parent-obgp-01, March 2001.

[8] L. Wang, H. Zhang, X. Zheng, et al, “A Novel OBGP–based mechanism for Lightpath Establishment in WDM Mesh Networks”, in Proceedings of ECOC 2003, Rimini, Italy, September 2003.

[9] G. Bernstein, D. Cheng, D. Pendarakis, et al, “Domain to Domain Routing using GMPLS, OSPF Extension V1.1(Draft)”, OIF2002.23.06, July 2002.

[10] S. Agarwal, C. Chuah, R. Katz “OPCA: Robust Interdomain Policy Routing and Traffic Control”, IEEE Openarch, San Francisco, USA, April 2003.

[11] L. Subramanian, I. Stoica, H. Balakrishnan, R. Katz, “OverQoS: Offering Internet QoS Using Overlays”, ACM SIGCOMM, Computer Communications Review, January 2003.

[12] M. Yannuzzi, A. Fonte, X. Masip, E. Monteiro, S. Sánchez, M. Curado, J. Domingo, “A proposal for inter-domain QoS routing based on distributed overlay entities and QBGP,” in Proceedings of the First International Workshop on QoSR (WoQoSR), LNCS 3266, Barcelona, Spain, October 2004.

[13] D. G. Andersen, “Resilient Overlay Networks”, MSc thesis, MIT, May 2001.

[14] Z. Li, P. Mohapatra, “QRON: QoS-aware Routing in Overlay Networks”, IEEE Journal on Selected Areas in Communications, June, 2003.

[15] J. Wu, J. M. Savoie, A. Vukovic, H. Hua, “Architecture, Routing, and Signalling for Peer-to-Peer Inter-Domain Optical Networking”, in Proceedings of 2002 National Fiber Optic Engineers Conference (NFOEC), Dallas, Texas, USA, September 2002.

[16] G. M. Bernstein, V. Sharma, L. Ong, “Inter-Domain optical routing”, Journal of Optical Networking, Vol. 1, No. 2, February 2002.