ROUTING OF MPLS FLOWS OVER AN AGILE ALL-PHOTONIC STAR NETWORK

Peng He and Gregor v. Bochmann
School of Information Technology and Engineering (SITE)
University of Ottawa, Ottawa, ON, K1N 6N5, Canada
{penghe, bochmann}@site.uottawa.ca

ABSTRACT
In this paper, we study how MPLS flows are routed in an Internet that contains a centrally-controlled agile all-photonic star WDM network (AAPN). Two scenarios are considered, namely deploying AAPN within one OSPF (Open Shortest Path First) area and within several OSPF areas. Since the AAPN provides an NxN interconnection structure for the N edge nodes of the AAPN architecture, the straightforward usage of a routing protocol like OSPF leads to scalability problems. In the first scenario, we have identified several schemes by which this scalability problem can be reduced. The idea is to introduce "virtual routers" that represent a collection of edge nodes (and possibly also the core node), thus reducing the number of paths between the "routers". In the second scenario, we focus on inter-area routing in large-scale IP/MPLS networks. This paper proposes a novel framework for inter-area MPLS Traffic Engineering. The key to our proposal lies in deploying the AAPN architecture as the OSPF backbone area and introducing the concept of "virtual area border routers" (v-ABRs). Compared with other proposals, our proposal can provide globally-optimized inter-area routing and has very good compatibility to existing traditional IP/MPLS routers.

KEY WORDS
MPLS, OSPF, Inter-area Traffic Engineering, AAPN

1. Introduction

Most carriers, including Telus, Bell Canada, AT&T, MCI and British Telecom, are migrating to an IP based converged network for provisioning multi-services (data, voice, video, etc.). In such IP networks, Multi-Protocol Label Switching (MPLS) is adopted to enable Traffic Engineering (TE) and support Virtual Private Networks (VPN). Together with Diffserv, MPLS can also provide Quality of Service (QoS) support.

An Agile All-Photonic Network (AAPN) [1,2,8], with a composite star topology, can potentially provide an efficient high bandwidth/high performance core transport network solution for carriers. Hence, it is very important to design and position AAPN to support IP/MPLS architecture and protocols. Deploying AAPN in an IP/MPLS network environment needs signalling and routing information exchange between them. The routing information exchange and associated signalling for this inter-working is the focus of this paper. Particularly, we study the Open Shortest Path First (OSPF) [4] IP routing protocol, which is commonly used for routing within a single administrative domain.

1.1 OSPF/OSPF-TE Review

OSPF [3] is a link-state routing protocol that is used by MPLS and GMPLS (Generalized MPLS) (with extensions). Each OSPF-running router exchanges LSAs (link state advertisements) through a reliable flooding mechanism to build up and synchronize its link state database (LSDB) with the database of other nodes in the network. The LSDB thus becomes a complete representation of the network topology and resource information (OSFP with TE extensions, OSPF-TE [4,5]). Based on it, each router can run the shortest-path-first (SPF) algorithm to compute its routing table, or run constraint-shortest-path-first (CSPF) algorithm to perform source routing. OSPF-TE and RSVP-TE (Resource Reservation protocol with TE extensions) [6] are fundamental of MPLS TE that can hence compute and establish explicitly routed LSPs (label-switched paths) whose paths follow a set of TE constraints.

OSPF is a hierarchical routing protocol that supports large networks through multiple OSPF areas: one backbone area (Area #0) surrounded by non-backbone areas. Area border routers (ABR) are located at the border between the backbone and the non-backbone areas, and distribute summarized information among the areas.

1.2 Overview of Agile All-Photonic Networks

As shown in Figure 1, a centrally-controlled AAPN consists of a number of hybrid photonic/electronic edge nodes connected together via several load-balancing core nodes and optical fibers to form an overlaid star topology. By introducing concentrating devices, AAPN can support up to 1024 edge nodes [2]. Each core node contains a stack of bufferless transparent photonic space switches (one for each wavelength). A scheduler at each core node is used to dynamically allocate timeslots over the various wavelengths to each edge node. An edge node contains a separate buffer for the traffic destined to each of the other

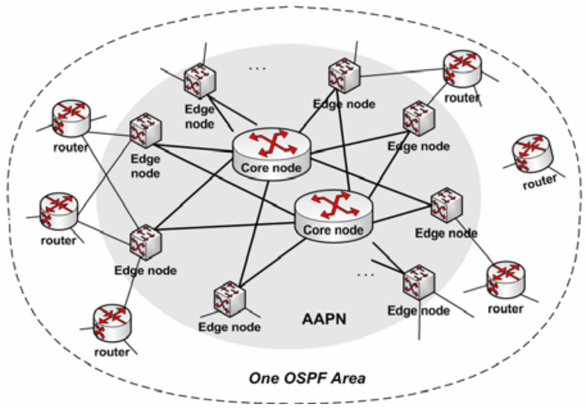

Fig. 1: Agile All-Photonic Network overlaid star topology.

edge nodes. Traffic aggregation is performed in these buffers, where packets are collected together in fixed-size slots (sometimes called bursts) that are then transmitted as single units across the AAPN via optical links. At the destination edge node the slots are partitioned, with reassembly as necessary, into the original packets that are sent to the outside routers. The term “agility” in AAPN describes its ability to deploy bandwidth on demand at fine granularity, which radically increases network efficiency and brings to the user much higher performance at reduced cost.

In this paper, we consider two scenarios to deploy an AAPN in an IP/MPLS network environment, namely within one OSPF area and within several OSPF areas. Since a large portion of the anticipated connections will need to traverse both the backbone area and the non-backbone areas, we focus on the second scenario, in which the proposed inter-networking framework can implement inter-area MPLS Traffic Engineering in an efficient and distributed manner.

2. Solving the Scalability Issue When Deploying AAPN within a Single OSPF area

The first scenario to deploy AAPN in IP/MPLS networks is in a single OSPF/OSPF-TE area (as shown in Fig 1).

2.1 The Problem: Scalability Issue

Since the AAPN provides an *NxN* interconnection structure for the N edge nodes of the AAPN architecture, the straightforward usage of a routing protocol like OSPF leads to scalability problems since the value of N could be very large (e.g., around 1000 and OSPF has to deal with *Nx(N-1)* links). Hence we need to consider what aspect of the AAPN topology should be exported to the IP/MPLS world and how to organize the related routing information exchange. In addition, the exported topology should be:

- as simple as possible (to reduce routing protocol traffic, routing calculation and the size of the link-state database)

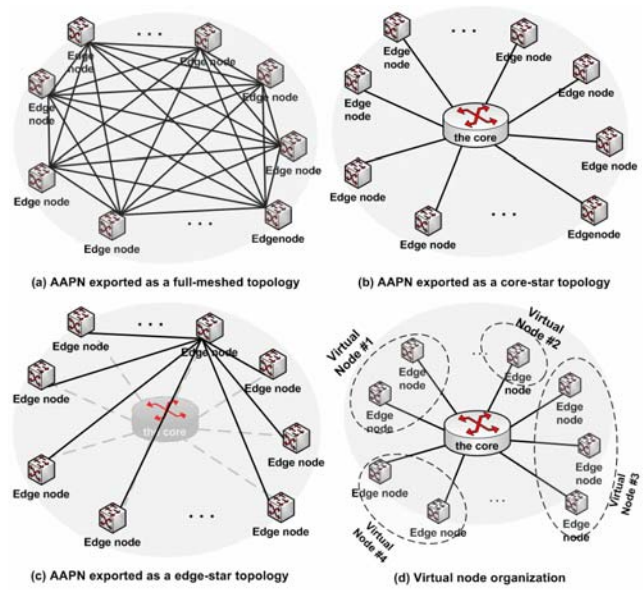

Fig. 2: Exported Topologies of Agile All-Photonic Network

- provide a good match with fault model of the AAPN (e.g., link/edge node/core-node failure)
- meet the traffic engineering requirement (not fully opaque to the outside)

Note: due to the symmetric architecture of AAPN (see Fig. 1), we use the “bundle” [5] concept to further reduce the overhead traffic to the outside. That is, all the links from one edge node to the core nodes are exported as one TE link. Similarly, the overlaid core nodes in AAPN are, if necessary, exported as one core node, named as “the core”.

2.2 Description of Proposed Solutions

1) Full-Mesh (Fig. 2a)

The whole AAPN (edge nodes, links, and core nodes) is exported as a full-mesh network to the outside (Fig. 2a). Within the AAPN, permanent connections are set up between all edge node pairs for routing information exchange, and may also be used for data exchange; while additional connections for data transmission may be established on demand. Each AAPN edge node behaves as an IP/MPLS router and the core nodes are invisible to the outside.

2) Core-Star (Fig. 2b)

A star topology (the core surrounded by N edge nodes) is exported to outside IP world. Each edge node maintains a two-way permanent connection only with the core for routing information exchange. Data connections will be established on demand. Each AAPN edge node behaves as an IP/MPLS router and the core node is visible from outside as an IP/MPLS router.

3) Edge-Star (Fig. 2c)

As an alternate to the core-star topology, AAPN can also be exported as a star topology where one edge node

TABLE I: COMPLEXITY ANALYSIS OF MESH AND STAR TOPOLOGIES

	FULL-MESH	CORE-STAR	EDGE-STAR
# OF 1-WAY CONNECTIONS TO BE MAINTAINED	$N \times (N-1)$	$2N$	$2(N-1)$
# OF ROUTER-LSAS FLOODED WITHIN AAPN AFTER A SINGLE CONNECTION FAILURE	*O(N^2)*	*O(N)*	*O(N)*
# OF ROUTER-LSAS FLOODED WITHIN AAPN AFTER A SINGLE LINK/EDGE NODE FAILURE	*O(N^3)*	*O(N)*	*O(N)*

is surrounded by the other *N-1* edge nodes. The major differences with the core-star topology are the following: (a) each edge node maintains a two-way permanent connection only with one particular edge node (not the core) for routing information exchange, and (b) each edge node behaves as an MPLS-capable IP router but the core node is invisible from outside.

Table I compares the above three exported topologies. Full-Mesh has a severe scalability problem when *N* is large: there are too many connections to set-up and maintain and hence there is a heavy load of control traffic although it is shared among all the edge nodes. The star topologies are much simpler. However, the load at the center of the star (the core or the central edge node) would become much heavier than at the other edge nodes when *N* is big.

4) Virtual Router Organization (Fig. 2d)

To find a balance between the simplicity of the exported topology and the load-sharing of control traffic, we propose the concept of a virtual router (VR) to organize the routing information exchange in the AAPN in a hierarchical manner. A VR represents a collection of co-located (or near-located) edge nodes and part of the core node switching capability (see dotted line in Fig. 2d). A VR is viewed as one IP/MPLS router, and these VRs, together with the core node, can form a virtual star architecture, thus reducing the number of paths among these "routers" and simplifying exported network topology, as compared with the Core-Star topology.

In an extreme case, a single VR may include all the edge nodes (there is no need for a core node any more), and the whole AAPN can be seen as one big router. In another extreme, a VR may just contain a single edge node. Generally speaking, to reduce the routing protocol traffic, the size of the VR should be big. But the TE requirements may push for smaller VR sizes (fine granularity). Hence the VR size is normally a balance between the above two extreme cases. Meanwhile, the VR-based topology is scalable since adding an edge node in a VR will not affect the whole exported topology.

The VR Organization implies a two-layer organization of the routing information exchange: (1) within each VR domain, and (2) between the VRs (still within the AAPN). The communication between the VRs can adopt the architecture of Full-Mesh, Core-Star or Edge-Star. Among these possibilities, we recommend Edge-Star or Core-Star because it achieves the balance of simple exported topology and load-sharing, and has the smallest number of permanent connections to be set up.

Each VR has a head (a designated edge node, possibly with a designated backup node). When an edge node finds a routing update from its neighbour router(s), it reports the update to its head node. The head checks the update, aggregates it, if possible, and forwards it to the heads of other virtual routers. Those heads then distribute the update to the edge nodes that are member of their respective VR domain. A simple internal routing cooperation protocol, like the one in [17], can be used for this purpose within VR domain. Note that the forwarding table of each edge node includes both the forwarding information per local external (non AAPN) output port and information for forwarding through the AAPN network.

2.3 Short Summary

Based on all the above analysis, we have the following conclusions when OSPF with a single area (or any other non-hierarchical routing protocol) is to be deployed over a network including an AAPN:

- A very small AAPN can be viewed as single big IP router with MPLS capability.
- A small or medium-sized AAPN (with a few tens of edge nodes) can be viewed as a full-mesh or a star topology where each edge node is viewed as an IP router with MPLS capability.
- A large AAPN should use hierarchical information exchange using the concept of virtual routers (VR) interconnected in a VR-star topology, as explained above.

3. An Inter-Area Traffic Engineering Framework for Deploying Several OSPF Areas Over an AAPN

Currently, several carriers have multi-area networks, and many other carriers that are still using a single IGP area may have to migrate to a multi-area environment as their network grows and approaches the single area scalability limits [18]. Hence, it would be useful and meaningful to extend current MPLS TE capabilities across IGP areas to support inter-area resources optimization. That is why RFC4105 was recently published to define detailed requirements for inter-area MPLS traffic engineering and ask for solutions

3.1 The Problem

An inter-area connection normally starts in a non-backbone area, traverses a backbone area, and terminates in another non-backbone area. MPLS TE mechanisms that have been deployed today by many carriers are limited to a single IGP area and can not be expanded to multi-areas

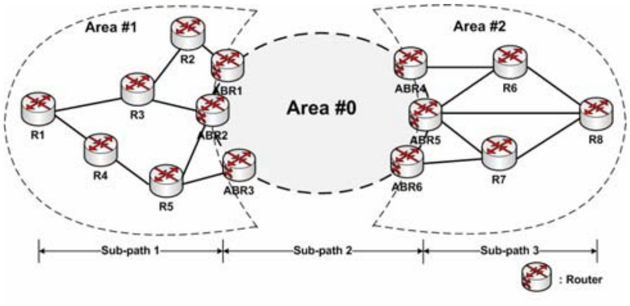

Fig. 3: A common inter-area network.

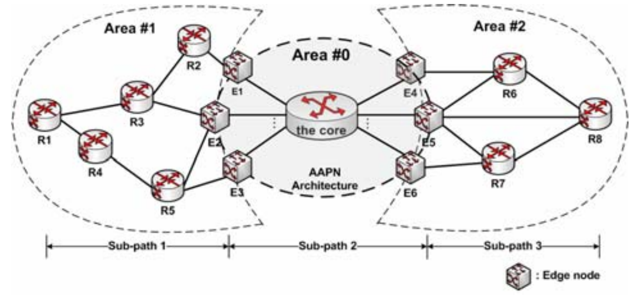

(a) Directly deploying AAPN as backbone area

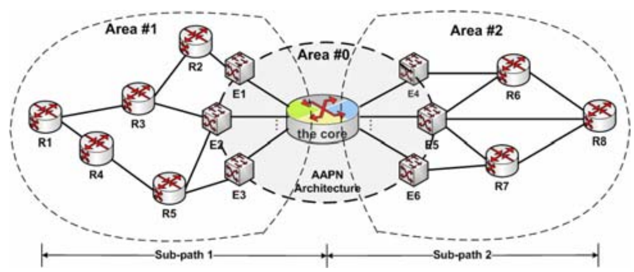

(b) Our proposal

Fig. 4: The framework we proposed deploys AAPN as backbone area #0.

directly. The limitation comes more from the routing and path computation components than from the signalling component. This is basically because the OSPF/OSPF-TE hierarchy limits topology visibility of head-end LSRs (Label Switch Routers) to their area, and consequently head-end LSRs can no longer run a CSPF algorithm to compute the shortest constrained path to the tail-end, as CSPF requires the whole topology information in order to compute an end-to-end shortest constrained path.

For an example, Fig. 3 shows a common multi-area network and we suppose R1 in Area #1 is the source node while R6 in Area #2 is the destination. Generally speaking, a non-backbone area (e.g., area #1 in Fig. 3) often has multiple ABRs (existing points). One ABR might be much closer to the destination of a requested MPLS connection than another. Because the head-end node does not have the entire topology, it does not know which ABR is the best choice. In Fig. 3, how could R1 choose an optimum ABR in Area #1 to the destination R6? Through local optimization, R1 may select ABR2 to be on the path, but how does ABR2 know what the best path is to go to R6? Although local optimization can be done in each of the respective areas along the inter-area path (R1 to R6), the simple summation of the three local optimizations does not necessarily lead to a global optimization. What many carriers want is to optimize their resources as a whole. Therefore, the question of how to implement inter-area routing with global optimization guarantee is a key issue in inter-area traffic engineering.

3.2 Related Work Review

Untill now, most papers [11,12,13,14,15] talking about inter-area routing center on the "how-to" issue, that is, how to find out an inter-area route (not optimal and not dynamic) and how to build up this path through the inter-area signalling process. [14,15] use a two-step approach to compute an inter-area route: find out a "loose inter-area route" first through topology aggregation/abstraction, then resolve the loose route into a strict path, area by area. Actually, this per area approach would always lead to sub-optimal resource utilization. In [12,13], the authors divided an inter-area path into two segments, one in Area #1 and one in Area #0 & #2 (see Fig. 3. for a rough look). Optimal routing of the 1st segment is done first by the head-end LSR; then based on the 1st segment, a far-end ABR (e.g., ABR5) computes the 2nd segment. Obviously, this approach can not achieve global route optimization either. Another category of approaches [16] claim that they can provide global optimization, but it is at the cost of building up an independent external PCE (Path Computation Element) network, which will decrease the efficiency and increase greatly the cost. In this paper, we propose a novel framework to deploy AAPN in a multi-area network so as to implement global resource optimization in an efficient and distributed manner.

3.3 Description of the Proposed Inter-area Traffic Engineering Framework

The direct and natural way to deploy an AAPN in a multi-area network is shown in Fig. 4a, where the core nodes locate in the middle of Area #0 and the edge nodes act as ABRs at the border between Area #0 and non-backbone areas. However, in this scheme, inter-area routing with global optimization still can not be guaranteed. Therefore, we propose a novel approach/framework, shown in Fig. 4b, that can provide such guarantee. Our proposed framework consists of three main components, namely the routing-info, path computation and signalling components:

1) The routing-info component

This component is responsible for the discovery of the TE topology. As seen in Fig 4b, we expand the OSPF non-backbone areas a little step so that there is an overlap between Area #0 and each expanded non-backbone area. Then the AAPN edge nodes located in the overlap, together with their direct TE links to the core and the associated part of the core, belong to both the Area #0 and a non-backbone area. In such a scenario, legacy routers in a non-backbone area see related AAPN edge nodes as normal internal IP/MPLS routers, see the AAPN TE links as normal internal links and see the associated part of the

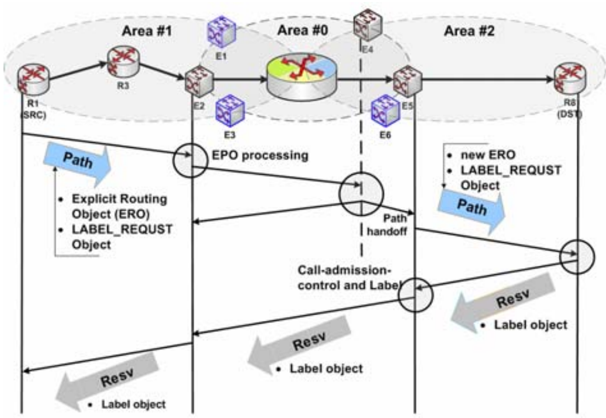

Fig. 5: Inter-area LSP signaling process

core as the (only) ABR of its non-backbone area. In another words, a legacy router sees what it can see in its area about the core as an ABR, which we name as a virtual-ABR (v-ABR). For each legacy router in an expanded non-backbone area, the exchange and distribution of routing/TE information is just like in any other standard OSPF/OSPF-TE area.

While within the Area #0, the AAPN edge nodes that belong to the same non-backbone area can be considered as a big virtual router and one edge node in each virtual router is selected as the head of this virtual router. Then the area-specific reachability information is exchanged among these heads and distributed to each edge node and then outside routers. This is very similar to the hierarchical routing information exchange described in Section II B. Therefore, it is the head edge node that actually performs the functions of the virtual-ABR. Note that only the reachability (not TE) information, which is enough, is exchanged among virtual routers.

2) The path computation component

In our framework, an inter-area LSP can be considered consisting of two segments (instead of three, as in Fig. 3) as shown in Fig. 4b: one in the head-end (expanded) area and one in the tail-end (expanded) area. The core connects these two segments/sub-LSPs to form a complete inter-area LSP.

The most interesting thing is that local routing optimization with both of these two sub-LSPs can lead naturally to a globally-optimized inter-area LSP. As seen in Fig. 4b, this is due to the particular star topology of the AAPN architecture and the load-sharing core nodes that can be exported as one single core to the outside.

The local routing optimization in the head-end area can be performed by the source LSR, which takes TE topology and LSP constraints as input. While in the tail-end area, local routing optimization is done by one of the edge nodes in the area (we will discuss this later). Obviously, dynamic inter-area routing can be implemented in our proposed framework.

3) The signalling component

This component is responsible for the establishment of the LSP along the computed path. In Fig. 4b, consider the case that a source LSR (R1) wants to set up a LSP to a destination LSR (R8). R1 must first compute an optimized path to the virtual-ABR of Area #1 through CSPF, and then signal this establishment request to the network. Suppose RSVP with TE extension (RSVP-TE) [6] is used as the signalling protocol.

As shown in Fig. 5, R1 starts the signalling process by creating a RSVP Path message with two objects inserted, namely LABEL_REQUEST Object (LRO) to request a label binding for the path, and EXPLICIT_ROUTE object (ERO) to indicate the computed explicit path (with one sub-object per hop). However, R1 has to use the loose ERO sub-objects for the hops outside Area #1. In Fig. 4b, the ERO specifies the explicit path as R1→R3→E2→v-ABR #1→→R8, where R8 is a loose ERO sub-object. Then, R1 sends the Path message to the next hop defined in the ERO, which is R3.

R3 (a non AAPN node) receives the Path message and processes it as follows:

a) Checks the message format to make sure everything is OK,

b) performs admission control to check the required bandwidth,

c) Stores the “path state” from the Path message in its local Path State Block (PSB) to be used by the reverse-routing function, and

d) if successful, deletes the 1st sub-object (itself) in the ERO and forwards the Path message according to the new 1st sub-object (next hop) in the ERO, in our case, E2.

E2, an AAPN edge node, receives the Path message from R3 and checks the contained ERO. If E2 finds that the IP address of the 2nd sub-object in the ERO is a v-ABR #1 and the 3rd sub-object (with the loose attribute) is beyond Area #1, then E2 has the task of resolving the loose sub-object into strict ones. In our case, there is one loose sub-object, R8, which represents the destination of the requested LSP. Although E2 can not find a strict path from v-ABR #1 to R8 by itself, it knows who can. First, by checking the inter-area reachability information and internal parameters, E2 finds out which group of edge nodes (also which associated v-ABR) locates in the same non-backbone area as R8. In Fig. 4b, these are E4, E5 and E6 (v-ABR #2). Second, it selects an edge node among them randomly, e.g., E4. In the third step, E2 removes the first two sub-objects (itself and v-ABR #1) from the ERO of the original received Path message, and inserts v-ABR #2 at the top, then forwards the modified Path message to E4.

When E4 receives the Path message and finds the 1st sub-object in the received ERO is v-ABR #2, together with a loose second sub-object, R8, it knows that it should find an explicit path between these two sub-objects. As shown in Fig. 4b, E4 is capable to do the resolving work because E4 and R8 reside in the same expanded area, Area #2. E4 finds the optimized explicit path v-ABR #2→E5→R8. E4 then replaces the ERO object in the received Path message with a new ERO object that stores the resolved explicit route (E5 R8). Finally, E4 forwards the new modified Path message to E5 as if it were forwarded from E2 by using E2's data (IP address, etc.). We call this process a Path message handoff. At the same time, E5 also sends an acknowledge message (containing the resolved path) to E2 (Fig. 5). From the above handoff process, we can see that only the area-specific reachability (not TE) information needs to be exchanged among area. In our proposal, TE information is organized within each area.

Edge node E5 receives the Path message and believes it is from E2. Since all the sub-objects in the received ERO are strict, E5 processes this Path message in a standard way, just as R3 did in Area #1, and then forwards the processed Path message to R8.

When the destination, R8, gets the Path message, it responds to this establishment request by sending a RSVP Resv message. The purpose of this response is to have all routers along the path perform the Call Admission Control (CAC), make the necessary bandwidth reservations and distribute the label binding to the upstream router.

The label is distributed using the Label Object in the Resv message. The labels sent upstream become the output labels for the routers receiving the label object. The label that a router issues upstream become the inbound label used as the lookup into the hardware output tag table. The reservation-specific information is stored in the local reservation state block (RSB) of each router.

When the AAPN edge node E5 receives the Resv message from downstream (R8), it starts internal AAPN signalling to ask the core to set-up a connection from E2 to E5 (omit v-ABR #1&2). If bandwidth is available for this connection, the core informs both E2 and E5. E5 then sends a Resv message to E1. The label could be, for instance, a timeslot number or a wavelength, on the links between E2, the core, and E5. Note that there is no timeslot exchange in the core.

The Resv message makes its way upstream (see Fig. 5), hop by hop, and when it reaches the source LSR, R1, the inter-area path is set-up: R1→R3→E2→v-ABR#1→v-ABR#2→E5→R8. Now, a globally-optimized inter-area LSP is set-up. It can be maintained and torn-down just as any normal intra-area LSP tunnel.

3.4 Further discussions

Compared with other inter-area proposals, our proposal can provide globally-optimized inter-area routing and does not require any changes on existing traditional IP/MPLS routers, hardware or software, to implement (good backward compatibility). Furthermore, there is no node having global TE information. Instead, the TE information is distributed on per-area basis and only area-specific reachability (not TE) information is exchanged among areas. Global optimization is achieved through cooperation and interaction between AAPN edge nodes in different areas (Path message handoff). In addition, for the 2nd half of an inter-area LSP (in the tail-end area), the optimized routing computation is done randomly by an AAPN edge node in the tail-end area. Hence, load-sharing among these edge nodes is achieved.

Under our proposed framework, inter-area routing can be dynamic. In addition, re-optimization of an inter-area TE LSP can also be implemented, either locally within an area (by the head-end LSR for the 1st half or by an edge node for the 2nd half of LSP) or globally by the head-end LSR (end-to-end re-optimization).

Regarding inter-area QoS, there is not much work left. Current single area QoS mechanisms can be expanded directly to multiple areas and to AAPN [7].

As seen in Fig. 4b, our proposal keeps OSPF's hierarchical structure and just expands non-backbone areas a little. Hence the scalability of our proposal is as good as OSPF/OSPF-TE.

The routing concepts discussed in this paper are based on the assumption that there are a number of edge nodes (that are connected with other routers - through traditional Internet technology - and belonging to the same OSPF area) and these edge nodes can establish optical connections between one another in an agile manner and can adjust the bandwidth of each connection in an agile manner according to the varying bandwidth that is required by the IP traffic. We think any agile optical switching technology (burst switching, TDM (Time-Division-Multiplexing), routed wavelength (with less bandwidth flexibility)) may be used.

Furthermore, our proposal is not limited to AAPNs, it is actually applicable in a much larger context. The fundamental ideas abstracted from our proposal are: (1) a "load-symmetrical" network (optical mostly) as backbone, (2) overlap between backbone and none-backbone areas, and (3) virtual-ABR. A load-symmetrical optical network is a network that can provide one or several optical connections for each edge node pair (source-destination pair) and the load among the several optical connections of each edge node pair is balanced. Hence the "bundle" concept [5] can be used and a single core can represent arbitrary optical network topology. Note that "load-symmetrical" does not mean the loads among any distinct edge node pairs are balanced; it only refers to the load-balancing within the connections of one edge node pair. Load symmetrical networks do not need to be symmetrical physical network topology, although a symmetrical physical network topology (such as AAPN, Petaweb [10] and PON (Passive Optical Networks)) is easier to be load-symmetrical. Similar ideas can also be applied to IP over large clouds (e.g., IP over ATM).

In summary, the basic idea of our proposal can be considered an interesting framework for implementing inter-area MPLS Traffic Engineering.

4. Conclusions

We have studied how MPLS flows are routed over a centrally-controlled agile all-photonic star WDM network (AAPN), and focus on OSPF/OSPF-TE routing information exchange. We present several schemes to improve the scalability when deploying OSPF over an AAPN. We show that small-sized/medium-sized AAPN can use full-mesh or star-type exported topology, while large AAPN should adopt a hierarchical structure to organize routing information exchange, e.g., introducing the idea of virtual routers.

AAPN is more suitable to be used in multi-area network environment due to its agility at the core and large capacity. Based on deploying AAPN in multi-area networks, we proposal a novel framework that aims to implement inter-area MPLS traffic engineering. Compared with other inter-area routing proposals, our proposal has two distinguishing characteristics:

1. Our proposal can provide globally-optimized inter-area routing;

2. There will be no change, hardware or software, on existing traditional IP/MPLS routers to implement our proposal.

Furthermore, our proposal is not limited to AAPNs; it is actually applicable to any load-symmetrical (optical) network with arbitrary physical network topology. Indeed, our proposal can be considered as a highly competitive candidate solution to the recently published RFC by IETF network working group, RFC 4105 (Requirements for Inter-Area MPLS Traffic Engineering). One of our future research topics is to investigate the impact of our proposal on the selection of MPLS protection paths in the context of an underlying AAPN.

Acknowledgements

This work was supported by the Natural Sciences and Engineering Research Council (NSERC) of Canada and industrial and government partners, through the Agile All-Photonic Networks (AAPN) Research Network. This work is part of an AAPN Partnered Research Project sponsored by TELUS. The authors thank Rainer Iraschko for useful discussions.

References

[1] G. v. Bochmann, M.J. Coates, T. Hall, L. Mason, R. Vickers and O. Yang, “The Agile All-Photonic Network: An architectural outline”, Proc. Queen's University Biennial Sym-posium on Communications, 2004, pp.217-218

[2] L.G. Mason, A. Vinokurov, N. Zhao and D. Plant, "Topological Design and Dimensioning of Agile All Photonic Networks", accepted to "Computer Networks", special issue on Optical Networking,

[3] J. Moy, “OSPF Version 2”, IETF RFC 2328, April, 1998.

[4] D. Katz, K. Kompella, and D. Yeung, “Traffic Engineering (TE) Extensions to OSPF Version 2,’’ IETF RFC 2328 April, 2003

[5] K. Kompella and Y. Rekhter, “OSPF Extensions in Support og Generalized MPLS”, IETF Internet Draft, Oct., 2003

[6] Awduche, D., et al. "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, December 2001.

[7] A. Vinokurov, X. Liu, and L. Mason, "Resource sharing for QoS in Agile All Photonic Networks", OPNETWORK 2005, Washington D.C., August 2005.

[8] Trevor J. Hall, Sofia A. Paredes, Gregor v. Bochmann, "An Agile All-Photonic Network", invited paper, The International Conference on Optical Communications and Networks, ICOCN 2005; Bangkok, Thailand, 14-16 December 2005.

[9] K. Kompella et al. “Link Bundling in MPLS Traffic Engineering,” IETF Internet draft, work in progress.

[10] R. Vickers and M. Beshai, “PetaWeb Architecture,” Networks 2000, Toronto, Canada, 2000.

[11] Dongmei Wang, John Strand, Jennifer Yates, “OSPF for Routing Information Exchange Across Metro/Core Optical Networks” Optical Networks Magazine, Vol. 3, Issue 5 , September/October 2002.

[12] T. Miyamura, T. Kurimoto, M. Aoki and A. Misawa,“An Inter-area SRLG-disjoint Routing Algorithm for Multi-segment Protection in GMPLS Networks,” Proc. ICBN 2004, 2004.

[13] Takashi Miyamura, “A Multi-layer Disjoint Path Selection Algorithm for Highly Reliable Carrier Services”, GLOBECOM '04.

[14] D. L. Truong and B. Thiongane, "Dynamic routing for shared path protection in multidomain optical mesh networks," J. Opt. Netw. 5, 58-74 (2006)

[15] Tarek Saad and Hussein T, Mouftah, “Inter-Domain Wavelength Routing in Optical WDM Networks”, 11th International Telecommunications Network Strategy and Planning Symposium, NETWORKS 2004, 2004

[16] Mohamed Boucadair, “A PCE-based Approach for providing inter-AS MPLS-based QoS tunnels”, draft-boucadair-pce-interas-00.txt, 2004

[17] M. Choe, J. Wybenga, and B. C. Kang, “A Routing Coordination Protocol in a Loosely-Coupled Massively Parallel Router”, In Proc. of HPSR02, May 2002.

[18] Le Roux, et al., “Requirements for Inter-Area MPLS Traffic Engineering” , IETF RFC 4105, June 2005