An implementation of an OSPF-TE to support GMPLS-controlled All-Optical WDM Networks

Luiz Gustavo Zuliani, *Student, IEEE,* Marcio Savasini, *Student, IEEE,* Gustavo Sousa Pavani, *Student, IEEE,*
Rafael Pasquini, *Student, IEEE,* Fábio L. Verdi, *Student, IEEE,* Maurício Magalhães, *Member, IEEE*

Abstract— **GMPLS-controlled all-optical networks are the promise to handle the increasing volume of IP traffic. The GMPLS routing and signaling protocols, mainly OSPF and RSVP, work in such a way that the route calculation for optical circuits does not take into account the label (lambda) availability. This is not optimal in terms of network usage and blocking probability of new circuits. To deal with this scenario, more efficient RWA algorithms could be used to calculate the route and the wavelength assignment at one time. These RWA engines need to know the optical topology in a way that is not described by the current OSPF standards. This paper proposes Traffic Engineering extensions to the OSPF protocol to enable the GMPLS control plane to take advantage of the most effective RWA classes. A prototype was developed and deployed in an optical-simulated copper-based network to verify its feasibility based on the bandwidth overhead generated in the control plane.**

Index Terms— **Constraint-based routing, GMPLS, optical networks, OSPF-TE, routing protocols, wavelength assignment (RWA).**

I. INTRODUCTION

The explosion of the Internet traffic in the recent past years has increased the need for high speed IP networks. Telecom companies and Internet Service Providers (ISP) have invested to deploy WDM networks to satisfy this need. Currently, there is a number of protocols between the IP and WDM layers, such as ATM and SONET/SDH, used primarily to support Quality of Service (QoS) and rapid fault restoration, respectively. These networks have a complex management, because each protocol layer must have its own configuration and resource provisioning, which can take hours or even weeks [1]. Also, these protocols add an extra traffic overhead in the Optical Line System (OLS). The next generation optical networks will carry the IP protocol directly over a WDM plane (IP over WDM). They are basically composed by IP routers, optical transponders, all-optical (or Photonic) Cross-Connects (OXC) and the OLS, which in turn is composed by fibers and amplifiers or regenerators. In these networks, an optical circuit called *lightpath* is created by assigning a wavelength (or *lambda*) to a set of optical links. A lightpath is then used as a point-to-point link between a pair of IP routers.

To automate the management, configuration and resource provisioning of the next generation optical networks, Generalized MultiProtocol Label Switching (GMPLS) has been defined [2]. GMPLS is an Internet Engineering Task Force (IETF) proposal standard, and it extends MPLS to optical networks in several ways, such as: the distinction between data plane and control plane, the possibility of link identification without IP addresses (*unnumbered links*), link agregation (*link bundling*), and dissemination of a LSP as a Traffic Engineering (TE) link (*Forward Adjacencies*). Depending on the carrier service model, the IP and Optical layers could have different levels of integration. In the *Overlay Model*, the layers are separated and interact in a Client-Server fashion. The *Peer Model* runs a single control plane over both layers, and the *Augmented Model* is located somewhere between the previous models, where each layer has its own routing instances, but routing information is exchanged between them.

The Link Management Protocol (LMP) [3] was designed to address the issues related to optical link management. Also, new extensions were made in the routing and signaling protocols to attend optical domains. The *de facto* GMPLS protocols adopted are the Open Shortest Path First (OSPF) [4] for routing and the Resource reSerVation Protocol (RSVP) [5] for signaling. When a new lightpath is needed, an online Constraint-Based Routing (CBR) algorithm is used to calculate a route, based on the TE link properties flooded by the OSPF protocol. Then, is up to the RSVP protocol to ask to one of the border routers a label (or a lambda), and then reserve it along the nodes of the route. The border router can suggest one or more lambdas to RSVP, but even in this case its decision is taken based on local lambda availability. In fact, currently there are no GMPLS nor OSPF documents that propose the flooding of the state of every single lambda. Therefore, the path calculation process is made completely independent of the wavelength assignment, which is not optimal in terms of blocking probability and network utilization. It is possible to improve these metrics using more sophisticated online Routing and Wavelength Assignment (RWA) algorithms to calculate the path and the wavelength to be assigned to it simultaneously. The most effective RWA engines require a *global knowledge* of the network consisting of the state of the available optical channels, the physical layer constraints [6], and the WDM plane logical restrictions (such as the number of available transponders, or the OXC switching matrix capa-

L. G. Zuliani (e-mail: zuliani@dca.fee.unicamp.br), M. F. Magalhães (e-mail: mauricio@dca.fee.unicamp.br), Rafael Pasquini (e-mail:pasquini@dca.fee.unicamp.br) and Fábio L. Verdi (e-mail:verdi@dca.fee.unicamp.br) are with the Department of Computer Engineering and Industrial Automation (DCA). M. Savasini (e-mail: savasini@decom.fee.unicamp.br) and G. S. Pavani (e-mail: pavani@decom.fee.unicamp.br) are with the Communications Department (DECOM). Both departments are with the School of Electrical and Computer Engineering (FEEC), State University of Campinas (UNICAMP), Campinas-SP, Brazil (DCA-FEEC-UNICAMP. C. P. 6101. Campinas-SP. Brazil. CEP 13081-970).

bility). Some proposals [7], [8] were made to flood via OSPF the total number of wavelengths and the number of available wavelengths, but this information is not sufficient to the class of RWA algorithms that require a global knowledge of the network, precisely the most effective class. In this paper, we propose a new OSPF-GMPLS extension, which floods not only the state of every single available lambda, but also a complete set of information related to the WDM plane. This information is used by complex RWA engines to calculate the path and the wavelength assignment simultaneously. In this way, only one label is suggested to the RSVP protocol (the best one), that is already known to be available during the path calculation. This could result in even lower blocking probability and better network resources utilization. We implemented the proposed extensions to the OSPF-TE protocol, which was deployed in a copper-based testbed to verify its feasibility based on the bandwidth overhead generated in the GMPLS control plane.

The remainder of this paper is organized as follows: in Section II we present an overview of the OSPF protocol, from its origins to its use by the GMPLS architecture. In Section III we describe the RWA routine. Later, in Section IV a detailed description of the OSPF-GMPLS implementation is presented. Also, an evaluation is made based on the traffic load generated by the implementation. Finally, we conclude the paper in Section V.

II. OSPF OVERVIEW

The OSPF development began in 1987, and in 1991 the first specification was published. The motivations for a brand new protocol was the flaws of OSPF's predecessor, the Routing Information Protocol (RIP). By that time, when the size of the Autonomous Systems (AS) began to increase in a fast pace, RIP's convergence times and the band consumed in the process had began to be unacceptable. RIP is a *distance vector* routing protocol, where the *metric* used to calculate the routes is the *distance* from the other subnets. On the other hand, OSPF is a link-state protocol, which uses a more flexible metric. A *cost* is attributed to each link, usually related to the link bandwidth capacity. Further than deal with RIP's problems, OSPF introduced new functionalities, such as equal-cost multipath, routing hierarchy, internal and external routes separation and improved security.

Each OSPF router advertises the state of its links in the form of Link State Advertisements (LSA). LSAs are distributed across the network via a complex mechanism called *reliable flooding*. This mechanism assures that all the routers in a network (or OSPF area) will have the same set of LSAs called Link State Database (LSDB). There are five types of LSAs, but only the first two are used in single OSPF area (the area 0.0.0.0 is called *backbone*, and is always present): the type 1 is called router-LSA. Each router originates a single instance of it describing all active interfaces and neighbors. A network composed solely of point-to-point links has its LSDB formed by LSAs of type 1 only. LSAs of type 2 are called network-LSAs. They are used in broadcast and NonBroadcast MultiAccess (NBMA) networks, where the flooding process is slightly different, and more efficient in terms of bandwidth overhead. Taking into account that in these networks each router can communicate directly with any other router attached in the access medium, the number of adjacencies can be reduced from $n * (n-1) / 2$ to n, where n is the number of routers. All the routers, instead of sending LSAs to and receiving them from all other routers, they send to and receive from only two specific routers, called Designated Router and Backup Router (the backup is present only for robustness purposes).

LSAs are carried between routers by OSPF packets. They are encapsuleted into IP packets with type 89 in the type field of the IP header. The destination IP address is always set to the neighbor's IP address or to one of the OSPF multicast addresses, in case of a broadcast network. There are five types of OSPF packets. The first one (Type 1) is the Hello packet, used to discover and maintain neighbor relationship. All the other four are used in the LSDB synchronization.

The routing table of each router is calculated by applying a Shortest Path First (SPF) algorithm in a reachability tree constructed from its LSDB. This results in paths to each known subnet, and the next hop to each destination is inserted in the routing table. Once the route to a certain subnet is calculated, all related traffic to that destination will use that path, no matter if the links that compose the route have or not sufficient bandwidth to handle that traffic (even if exist other links with available capacity in the network that could reach the same destination). This occurs because it is not possible to perform Traffic Engineering in a network using only the link cost parameter. Other link properties must be taken into account. In the late nineties the IETF MPLS working group proposed enhancements to the OSPF protocol to allow it to carry not only the cost metric, but also other TE link properties, as bandwidth parameters, local and remote IP addresses, and administrative group/resource class/color. The OSPF protocol with these TE extensions is called OSPF-TE. The TE-link metrics, used by CBR engines to calculate pseudo-circuits (or tunnels) routes across a MPLS cloud, are carried by OSPF-TE using a new type of LSA, called Opaque LSA. They are slightly different from standard LSAs (just a couple of fields were added to the LSA header), and were conceived to allow applications to use the OSPF flooding mechanism to spread its data throughout the OSPF topology, or some part of it. Three types of Opaque LSAs were defined, and OSPF-TE uses only the type 10, that floods the information through all OSPF area (there is no standard defined to allow TE in multiple OSPF hierarchies). The TE information transported by an opaque LSA - a TE LSA in this case - is organized in structures called Type-Length-Value (TLV) triplets. It is an extensible method to carry protocol data. A TLV is composed of three fields: The first field is the *type* field. It describes the type of information being carried, and its value is an application specific code. The second one is the *length* field, that informs the numbers of octets (or bytes) of the last field, the *value* field. It is this last field that really transports the data described by the first two fields. The type and length fields always have a size of 16 bits each. TLVs could be nested, i.e. a TLV could carry another TLV. The OSPF-TE RFC specifies two top TLVs, and they cannot be carried by the same TE-LSA. They are the

Router Address TLV and the Link TLV. The former carries an IP address that is always reachable as long as there is any connectivity to the router, and the later is formed by 9 sub TLVs that describe the TE link properties.

With the advent of GMPLS, new extensions were proposed to the OSPF protocol [4], this time to support Time Division Multiplexing (TDM) and optical links. These extensions are in the form of four new sub TLVs of the Link TLV. They are Link Local/Remote Identifiers, Link Protection Type, Interface Switching Capability Descriptor and Shared Risk Link Group sub TLVs. The first one deals with the fact that a GMPLS link could be an unnumbered link, i.e., could be addressed with a non-IP identification. The link Protection Type sub TLV is related to the level of resilience of a TE-link. The possible choices are unprotected, shared, dedicated 1:1 and dedicated 1+1. The Interface Switching Capability Descriptor sub TLV contains one of the following values: Packet-Switch Capable 1 through 4, Layer-2 Switch Capable, Time-Division-Multiplex Capable, Lambda-Switch Capable and Fiber-Switch Capable. Finally, the Shared Risk Link Group sub TLV carries the Shared Risk Link Group information, discussed in [4]. Although well documented and precise, these extensions are not enough to describe all the needs of the most efficient RWA algorithms.

III. RWA

The optical networks huge capacity will serve different applications from a variety of client networks. An optical network connection consists of the communication between a source and a destination node, traversing network links, forming a lightpath. Most optical nodes are not capable of converting wavelengths. Due to this restriction, a lightpath must have the same wavelength on every link throughout its route.

It is hard to predict how the traffic statistical properties will behave in an optical network. Lightpath configuration will be decided by a routing and wavelength assignment algorithm (*RWA*). Clients will request connections between two network nodes and a route and wavelength will be assigned to the lightpath, regarding imposed restrictions, such as, wavelength continuity, transmission power limitation or any other physical impairment.

Network traffic can be defined as static or dynamic. For the first case, all the connection requests are known *a-priori*. The other input to the RWA algorithm is the physical topology. The objective is to accommodate all the requests minimizing the use of wavelengths or scarce resources. The results of the algorithm may vary for different network device disposals.

The static case RWA may be formulated as a integer linear program (*ILP*) [9], being an example of a NP-hard problem [10]. The optimal solution is calculated offline, before any lightpath is established. However, a practical optimal solution is only feasible for limited size networks.

In the dynamic case, requests must be attended without any change in the configuration of the already established connections. The route and wavelength are chosen taking into consideration the momentary network state. The objectives are similar to those from the static case: accept the most number of future requests and minimize the use of scarce resources. But now the information about every other ongoing connection must be known. An easy and practical solution for the dissemination of the network state information is use the OSPF protocol.

A complete strategy to find a route and a wavelength for a connection request is the Shortest Path on the Available Wavelength Graph (*SPAWG*) [11]. It consists in creating a Wavelength Graph (*WG*), formed by different “wavelength planes”, where vertices and edges correspond to the network nodes and links on each plane. If wavelength conversion is possible on a node, the two corresponding vertices on the wavelength planes get connected. To find the route and wavelength for a request, one may use the Dijkstra algorithm.

The idea of link and conversion costs are introduced. Each edge in the WG receives an utilization cost. Building costs or physical impairments (e.g. attenuation, ASE, PMD, FWM) in the fiber can be modeled as link costs. The use of a converter can be understood to be the cost of changing wavelength planes. The Dijkstra algorithm analyzes which is the lowest cost route in the WG, from source to destination.

Depending on the size of the network, the dynamic RWA can last longer than the time interval between connection requests. To reduce its computational complexity, the RWA problem can be divided into two subproblems: the routing and the wavelength assignment. The three most common routing strategies are:

1) Fixed Routing - only one route, usually the shortest, can be used to connect the source and destination nodes. If there is no wavelength available on that route, the request is blocked. The routes for all source-destination nodes are pre-calculated.
2) Fixed-Alternate Routing - a list of paths is designated to all source-destination pairs. If the first route on the list cannot accommodate the request, the next route must be tested, until the end of the list. If all routes are not able to accommodate the request, it will be then blocked.
3) Adaptive Routing - the route will be selected taking into account the network configuration. Normally, the instant shortest path will be chosen.

The most simple wavelength assignment technic consist of picking up randomly a wavelength among all the available ones (to ballance the wavelength usage). Another simple method is to use a priority wavelength list. This way the last wavelengths listed are left idle for a long time, mitigating the blocking probability of future connections. The list could be static or dynamic. In the later case, the wavelengths are sorted by its usage on the network, where the most used wavelength is on top of the list. More complex strategies, such as MaxSum [12], allocate the wavelength that will leave the highest number of possible future connections available.

For this work the RWA algorithm presented in [13] was adapted to cope with the OSPF dissemination of information. All the information about the physical and the virtual topology is gathered by each node to start the RWA. The algorithm takes into account some power constraints to find the solution. The objectives are to minimize the blocking probability of the connections by routing, assigning wavelengths and maintaining an acceptable level of optical power and adequate Signal-Noise Ratio (SNR) all over the network. The minimum power constraint (sensitivity level) assures that the optical signal can be detected by each optical device on the selected route. The maximum power constraint guarantees the mitigation of non-linear physical impairments limiting the aggregate power on every link.

To reduce the complexity, the problem was separated into three subproblems. For the routing subproblem, the fixed-alternate approach was used, forming the list with routes according to Yen's algorithm [14]. For the wavelength assignment subproblem a priority list was chosen by numbering every wavelength randomly. The first wavelength available on the list is selected. This approach is used not only due to its simplicity and low computation cost, but also due to its good performance in terms of blocking probability and fairness.

Even if a route and a wavelength are assigned to a request, the power constraints must be verified. An iterative method for finding the transmitting power for each lightpath is used. The algorithm starts with -30dBm of transmitting power and at each iteration the power is increased by 1dBm. The iteration process ends when the power in every component are above the sensitivity level or the transmitting laser has reached its maximum power.

IV. OSPF-GMPLS IMPLEMENTATION

The goal of this work is to develop an OSPF system capable of efficiently disseminating wide and accurate information regarding the optical plane. Such information is then used as the input for RWA engines. The set of information flooded by this OSPF implementation is a new extension of the OSPF-TE for GMPLS optical networks. It covers the physical impairments and logical restrictions of the DWM plane, and aims at satisfying the needs of the most complex and effective classes of RWA algorithms. In the follow subsections an analysis of the implementation is made.

A. Implementation Details

The OSPF-GMPLS implementation was coded in C, using linux with the 2.6 kernel series. Each linux box running the application acts as an Optical Network Element (ONE) controller, i.e., the controller of a node in an optical network. Each ONE is composed by an OXC and an OLS, that can have multiple fibers. The system is composed by two programs, the OSPF-GMPLS daemon itself and a Lightpath Manager. The OSPF GMPLS daemon is the system core, a fully decentralized multithreaded piece of software. It is responsible for:

- Parsing the Initial Setup File that describes in what manner the daemon will perform its operations. The file supplies the refresh and timeout times for TE-LSAs, and the Lightpath Manager Address used for its socket connection (the Lightpath Manager is described later in this section). It also draws the way the TE-LSAs will be reflooded: in a periodical fashion or in a hybrid fashion, where the TE-LSAs are reflooded from time to time but also immediately after changes in its contents;
- Parsing the ONE Startup Configuration File that describes the local optical equipments asset, and translating it to the TLVs form, to allow it then be transported in TE-LSAs;
- Controlling the TE-LSAs flooding process;
- Keeping the LSDB updated, by controlling the remote TE-LSAs timeouts;
- Listening to the Lightpath Manager connection, to receive local TE-properties modifications due to setup or teardown of lightpaths, in which the current node is part of the lightpath route.

Both configuration files are written in the eXtensible Marketup Language (XML). In a real scenario, the data in these files are supplied by the network administrator and/or by the LMP protocol. The LSDB is also an in-memory XML tree, that in fact is a collection of the Startup Configuration File of all ONE's in the network. To accomplish this, all foreign TE-LSAs are converted to XML form after being received. Also, after every change in the LSDB, it is dumped in a XML file to be used by the RWA engine when needed. All the infrastructure necessary to originate and capture TE-LSAs was created using the OPSF-API that comes with the Quagga Routing Suite [15].

The Lightpath Manager is responsible for setting up and tearing down lightpath circuits. It supports two manners to create a lightpath: 1) via a RWA engine explained in Section III, or 2) by explicit route. In the former case, it is only necessary to supply the source-destination pair for the new lightpath, the bandwidth and the Q-factor [16], an optical quality factor. To calculate the route, the RWA engine uses the latest LSDB snapshot. In the second case, it is necessary to inform, in addition to the bandwidth and the Q-factor, all the nodes and links that will compose the route; the transponders with their respective operation powers and the lambda that will be used in the border nodes. Once the lightpath information is obtained, no matter if it comes from the RWA or from a manual process, the Lightpath Manager informs each node that is part of the route about the TE-properties that were changed. The communication is made using TCP socket connections established with all OSPF-GMPLS daemons running in the network. This is a centralized signaling method, since we are not using a distributed signaling such as RSVP. Each daemon that receives a notification from the Manager updates its LSDB local parts and prepares updated TE-LSAs to be flooded. The process ends with the updating of the LightPath DataBase (LPDB), a XML file that describes the virtual topology. Such database could be used to add grooming and Forward Adjacencies functionalities in the IP plane in a real scenario, as part of an Integrated Routing approach. Fig. 1 describes the process of creating a lightpath and how the modules interact to each other. First of all, the new lightpath data is gathered, using the build-in RWA engine or by informing manually the explicit route (step 1). Then, each OSPF-GMPLS daemon that is part of the route receives a message from the Lightpath Manager with the updated TE parameters - such as "lambda 3 in fiber 4 is now in use" (step 2a). At the same time, The Lightpath Manager also updates the LSDB (step 2b). The LSDB is updated, and new TE-LSAs with the new modifications are prepared to be flooded. As these new TE-LSAs are locally injected (step 3) and received (step 4) by all OSPF-GMPLS daemons that compose the optical core, their LSDB are updated after each reception. This way, all routers will have the same LSDB, that will be used by the RWA engine for the next lightpath calculation.

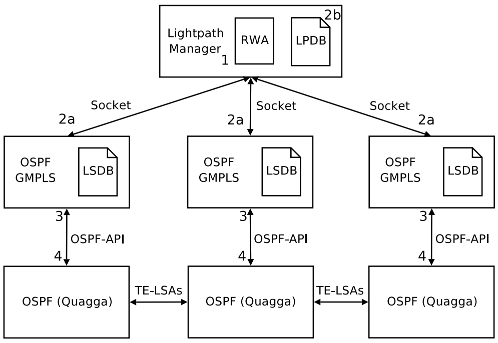

Fig. 1. OSPF-GMPLS modules interaction

The two parts that form the system use local and remote syslog facilities. To keep the linux routers synchronized, all of them use the Network Time Protocol (NTP).

B. Optical Plane equipments and features mapping to TLVs

An ONE can be divided into an OXC and an OLS. Therefore, we proposed two new top TLVs, the Optical Connection Controller (OCC) TLV and the Fiber Link TLV. Both have a set o sub TLVs that are described bellow. The details related to the optical features carried by each sub TLVs can be found in [13].

The OCC sub TLVs are:

- *OCC Router Address*: Used for identification purposes, and its value must be unique in the network. Value length is 4 octets.
- *OXC Switching Matrix Capability*: Denote the OXC switching matrix capability, for example 40x40 channels. Value length is 4 octets.
- *Input and Output Amplifiers Small Signal Gain*. Value length is 4 octets each.
- *Input and Output Amplifiers Spontaneous Emission Factor (NSP)*. Value length is 4 octets each.
- *Input and Output Amplifiers Saturation Power*. Value length is 4 octets each.
- *Transponder ID*: Each transponder must have a unique identification number in a node. Value length is 4 octets.
- *Transponder Maximum Transfer Rate*: The maximum transfer rate of a transponder, for example 10 Gb/s. Value length is 4 octets.
- *Transponder Tunable Lambdas List*: A list of channels that a transponder is capable of tuning. A transponder with a non-tunable laser will have only one lambda listed here. Value length is 2 octets per channel listed.
- *Transponder Minimum Operational powers*. Value length is 4 octets.
- *Transponder Maximum Operational powers*. Value length is 4 octets.
- *Transponder Current Operating Lambda*: If the transponder is being used, this sub TLV indicates in which channel its tuned. Value length is 4 octets.
- *Transponder Current Operating Power*: If the transponder is being used, this sub TLV indicates the laser power. Value length is 4 octets.

The Fiber Link sub TLVs are:

- *Fiber Link ID*: Each fiber must have a unique identification number in a node. Value length is 4 octets.
- *Local and Remote OCCs*: Describes the OCCs IDs at each end of the fiber. Value length is 4 octets each.
- *Length*: Denotes the length of a fiber, in Km. Value length is 4 octets.
- *Attenuation*: Denotes the attenuation of a fiber, in dB/Km. Value length is 4 octets.
- *Inline Amplifier ID*: Each inline amplifier must have a unique identification number in a node. Value length is 4 octets.
- *Inline Amplifier Small Signal Gain*. Value length is 4 octets.
- *Inline Amplifier Spontaneous Emission Factor (NSP)*. Value length is 4 octets.
- *Inline Amplifier Saturation Power*. Value length is 4 octets.
- *Inline Amplifier Position*: The distance between the amplifier and the node, in Km. Value length is 4 octets.
- *Optical Bandwidth*. Value length is 4 octets.
- *Channel and Total Maximum Power*. Value length is 4 octets each.
- *Lambdas In-use List*: The list of channels being used in a fiber. Value length is 2 octets per channel listed.

C. Evaluation

To evaluate the prototype, we deployed it in a network consisting of six linux routers and nine copper Fast and Giga Ethernet links, as shown in Fig 2. The objective is to simulate an optical network with an out-of-band in-fiber control plane. This way, it is possible to analyze the OSPF-GMPLS traffic overhead in the control plane. Each linux router runs the OSPF protocol implementation that is part of the Quagga Suite with the Opaque-LSA and OSPF-API options enabled. To simulate fiber links, each Ethernet link is part of a single VLAN (as they were crossovered point-to-point cables), and were set as point-to-point links in the Quagga OSPF configuration, overriding the autodetection option (ethernet links are autodetected as broadcast links). Therefore, there are no Designated and Backup Routers present in the network (consequently, there are no network LSAs being flooded). Thus, the flooding process is exactly the same as if the copper links were fiber links. This assures that the traffic measured in the simulated network is the same as if it were measured in an optical network, no matter if there is no real fiber nor OXC present.

Each Linux router was configured to emulate ONEs with four transponders, each one capable of tuning in 8 channels. Some fibers (varying from 2 to 4 depending on the router) are long enough to require inline amplifiers. To stress the control

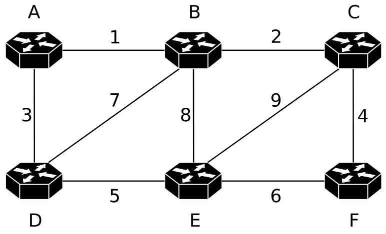

Fig. 2. Network topology used to evaluate the prototype.

plane with OSPF packets, each OSPF-GMPLS daemon was configured to reflood its local TE-LSAs at the highest rate allowed by the prototype - which is at every second -, no matter if the TE-LSAs contents were changed or not (due to lightpaths setups or teardowns). Indeed, reflooding LSAs at every second is a very high rate. Doing it at every 10 seconds should be enough to keep the LSDB synchronized in most networks. The traffic was measured with *Ethereal* [17], a network protocol analyzer, for more than four hours in a set of links. All of them had the same behavior. The results are shown if Fig. 3. The traffic measured is essentially due to TE-LSAs, because the traffic related to OSPF Hello packets is negligible. The peak rate detected was 50 Kbps at rare instants, with 15 Kbps average rate. The LSDB convergence time after a lightpath setup or teardown usually was less than 2 seconds, and the worst case detected was 3 seconds. These are plausible numbers, considering the average capacity of the control plane in GMPLS networks (usually a supervision channel or a legacy network with a minimum bandwidth of 155 Mb/s).

The traffic generated by the prototype is close to the one produced by a standard OSPF-TE implementation to MPLS in [18]. A more precise comparison is difficult to state, considering the different topologies, number of nodes and amount of data flooded by both implementations.

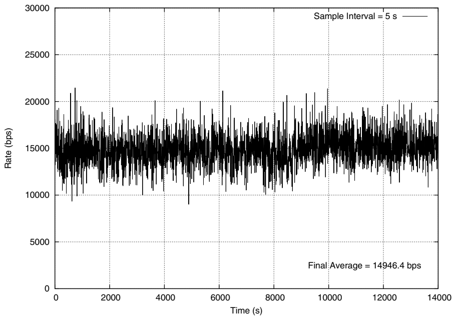

Fig. 3. Control Plane overhead related to OSPF-GMPLS traffic.

V. CONCLUSION

The TE extensions proposed in this paper in the form of new TLVs to the OSPF-TE protocol allow more efficient RWA engines to be used in GMPLS-controlled optical networks. A prototype was build and deployed in a testbed network to evaluate its practicability based on the overhead traffic measured in the control plane. While the prototype was operating in conditions to provide the maximum routing accuracy allowed by it (stressing the links with TE-LSAs at every second) the traffic measured was considered reasonable, considering the capacity of the control plane in GMPLS networks. Further studies are necessary to evaluate the implications of these TE extensions when deployed in networks with more nodes and different topologies. Also, other environments will be addressed considering Dense Wavelength Division Multiplexing (DWDM) with dozens of channels per fiber and different resilience capacities.

ACKNOWLEDGMENTS

The authors would like to thank CAPES, FAPESP and the GIGA Project and its partners for their support.

REFERENCES

[1] M. Kontak and D. Sobar, “IP based optical control plane,” *Proceedings of the 7th International Telecommunications*, vol. 2, pp. 427–433, June 2003.

[2] E. Mannie, “Generalized Multi-Protocol Label Switching (GMPLS) Architecture,” RFC 3945 (Proposed Standard), Oct. 2004. [Online]. Available: http://www.ietf.org/rfc/rfc3945.txt

[3] J. Lang, “Link Management Protocol (LMP),” RFC 4204 (Proposed Standard), Oct. 2005. [Online]. Available: http://www.ietf.org/rfc/rfc4204.txt

[4] K. Kompella and Y. Rekhter, “OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS),” RFC 4203 (Proposed Standard), Oct. 2005. [Online]. Available: http://www.ietf.org/rfc/rfc4203.txt

[5] L. Berger, “Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions,” RFC 3473 (Proposed Standard), Jan. 2003, atualizada pelas RFCs 4003, 4201, 4420. [Online]. Available: http://www.ietf.org/rfc/rfc3473.txt

[6] J. Strand, A. L. Chiu, and R. Tkach, “Issues for routing in the optical layer,” *IEEE Commun. Mag.*, vol. 39, no. 2, pp. 81–87, February 2001.

[7] N. Yamanaka, M. Katayama, K. Shiomoto, E. Oki, and N. Matsuura, “Multi-layer traffic engineering in photonic-GMPLS-router networks,” *Global Telecommunications Conference*, vol. 3, pp. 2731–2735, November 2002.

[8] Y. Yin and G. S. Kuo, “An improved OSPF-TE in GMPLS-based optical networks,” *Workshop on High Performance Switching and Routing*, pp. 241–245, May 2005.

[9] R. Ramaswami and K. Sivarajan, “Routing and wavelength assignment in all-optical networks,” *IEEE/ACM Trans. Networking*, vol. 3, no. 5, pp. 489–500, October 1995.

[10] I. Chlamtac, A. Ganz, and G. Karmi, “Lightpath communications: an approach to high-bandwidth optical WAN’s,” *IEEE Trans. Commun.*, vol. 40, no. 7, pp. 1171–1182, July 1992.

[11] I. Chlamtac, A. Farago, and T. Zhang, “Lightpath (wavelength) routing in large WDM networks,” *IEEE Trans. Commun.*, vol. 14, no. 5, pp. 909–913, June 1996.

[12] S. Subramaniam and R. Barry, “Wavelength assignment in fixed routing WDM networks,” *IEEE Int’l Conf. Commun.*, pp. 411–415, June 1997.

[13] G. S. Pavani and H. Waldman, “Dynamic routing and wavelength assignment with power constraints,” in *Proceedings of the XXI Simpósio Brasileiro de Telecomunicações*, September 2004.

[14] J. Yen, “Finding the k shortest loopless paths in a network,” *Management Science*, vol. 17, no. 11, pp. 712–716, July 1971.

[15] “Quagga Routing Suite,” June 2006. [Online]. Available: http://www.quagga.net/

[16] M. Yoon, J. Shin, C. Jeong, J. Jo, O. Kang, and S. Kim, “Optical-LSP Establishment and a QoS Maintenance Scheme Based on Differentiated Optical QoS Classes in OVPNs,” *Photonic Network Communications*, vol. 7, no. 2, pp. 161–178, March 2004.

[17] “Ethereal: A Network Protocol Analyzer,” June 2006. [Online]. Available: http://www.ethereal.com/

[18] H. M. Alnuweiri, L. Y. K. Wong, and T. Al-Khasib, “Performance of new link state advertisement mechanisms in routing protocols with traffic engineering extensions,” *IEEE Commun. Mag.*, vol. 42, no. 5, pp. 151–162, May 2004.