Study of GMPLS Lightpath Setup over Lambda-Router Networks

Indra Widjaja and Anwar I. Elwalid
Bell Laboratories, Lucent Technologies
Murray Hill, NJ 07974
iwidjaja@research.bell-labs.com

Abstract— **Generalized Multi-Protocol Label Switching (GMPLS) is currently being specified as the control plane for next-generation optical networks. This paper investigates how GMPLS affects the performance of lambda-router networks. Our performance study reveals the advantage of using a fully tunable laser at each source in the presence of a limited number of wavelength converters. We propose an enhancement to GMPLS which allows a lightpath request to crankback to a lambda router that can use a wavelength converter to complete the lightpath setup successfully. This effectively makes the wavelength converters shared across the network. It is shown that crankback provides a significant benefit when fully tunable lasers at the sources are not universally available.**

Keywords— **GMPLS, optical network, lambda router, crankback and blocking probability.**

I. Introduction

There is presently a great deal of interest in automating lightpath setups and teardowns in an optical transport network. An emerging trend in the industry is to utilize an optical-layer control plane, rather than a management plane as being done traditionally, to provision lightpaths. An intelligent optical-layer control plane is expected to offer several benefits including: rapid circuit provisioning, service flexibility such as bandwidth-on-demand services, enhanced interoperability of network elements from different vendors, and enhanced survivability by providing a dynamic rerouting capability when a failure occurs [1].

It is argued in [2] that the control plane (i.e., routing and signaling) for the optical layer should be based on reusing and leveraging existing control-plane protocols in order to reduce product development cycles and foster rapid deployment of a new class of optical network elements. It has recently become evident that the industry has gravitated toward Generalized MultiProtocol Label Switching (GMPLS) as the control-plane solution for next-generation optical networking. GMPLS is an extension to MPLS which enables Generalized Label Switched Paths (G-LSPs) such as lightpaths to be automatically setup and torn down by means of a signaling protocol [3]. This requires the definition of an MPLS label to be generalized so that a label can also be encoded as a time slot, a wavelength, or a spatial identifier. By taking advantage of the new definition of a generalized label, it becomes apparently clear that MPLS can also be extended to control and configure a TDM crossconnect, a lambda crossconnect, or a fiber crossconnect.

It is of paramount interest to investigate how GMPLS affects the performance of a next-generation network. This paper examines the LSP blocking performance of a lambda-routed optical network under varying scenarios. Insights into how various system parameters, such as the availability of wavelength converters, the tuning capability of a source laser, and the path selection algorithm, affect the blocking performance can be used to optimize the cost of deploying such an optical network. This paper also proposes an enhancement to the GMPLS signaling protocol that can be used to improve blocking performance without additional hardware support.

In Section II, we present a model of a lambda router and describe how GMPLS operates in such a lambda-router network. Section III investigates the performance of a lambda-router network using GMPLS lightpath setup mechanism. We propose an enhancement to GMPLS called crankback in Section IV. Finally, we conclude the paper in Section V.

II. Lightpath Setup in a Lambda-Router Network

A. Lambda Router

We start by presenting the structure of a lambda router, as shown in Fig. 1, that will be used in our study. The lambda router terminates multiple bi-directional fibers consisting of pairs of uni-directional fibers. These fibers are used to inter-connect lambda routers. Each incoming or outgoing fiber typically carries a multitude of wavelengths. Incoming signal from a fiber is first demultiplexed into separate wavelengths before entering an optical switch. The optical switch is capable of routing each wavelength from any input to any output. Each set of out-

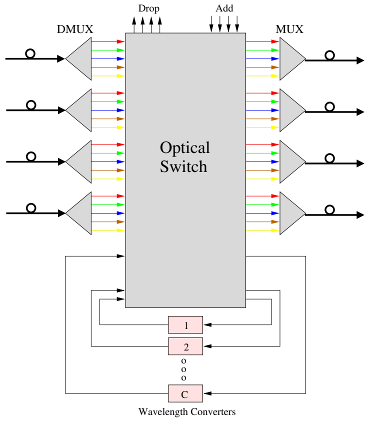

Fig. 1. Structure of a lambda router with a shared pool of wavelength converters.

puts of an optical switch are then wavelength-multiplexed before finally sending the signal to an outgoing fiber. Notice that the same wavelength can not be routed to the same outgoing fiber simultaneously. In order to avoid a potential wavelength collision, the lambda router also provides a shared pool of wavelength converters. When an incoming wavelength is to be routed to an outgoing fiber already carrying the same wavelength, the incoming wavelength first needs to be routed to one of the wavelength converters before being re-routed to the outgoing fiber. The wavelength converter translates the incoming wavelength to a new wavelength that is not already used in the outgoing fiber. Note also that the lambda router is capable of dropping and adding certain wavelengths locally to provide connectivities to edge devices such as IP routers. We assume that each add/drop line is not wavelength-multiplexed.

When a lambda router contains no wavelength converters, we call it a transparent lambda router. A lambda router having unlimited number of wavelength converters is called an opaque lambda router. Note that an opaque lambda router is similar to a circuit switch in the sense that an incoming request can always be routed to an outgoing link as long as there is enough capacity available in that outgoing link. The most interesting one is when a lambda router contains a limited number of wavelength converters, which we call a partially transparent lambda router.

B. Lightpath Setup via GMPLS

We now describe how GMPLS signaling establishes a lightpath from a source to a destination, assuming that GMPLS extensions to RSVP-TE is used [4]. The signaling approach follows a "downstream-on-demand" mode whereby an upstream node makes an explicit label request and a downstream node assigns a particular label in response to the request. Moreover, a node can only assign a label to its upstream neighbor if it is the destination node, or if it already receives a label from its next hop. Thus, the signaling phase consists of a label request traversing hop-by-hop from the source to the destination, followed by a label assignment traversing in the opposite direction back to the source. In RSVP-TE, a label request is contained in the Path message, while a label assignment is contained in the Resv message. In the context of this paper, a particular label is synonymous to a particular lambda.

When a source needs to establish a lightpath to a particular destination, it first determines the route/path toward the destination. Once the route has been determined, the source initiates an RSVP Path message containing a Label Request Object, an Explicit Route Object (ERO) which records the route that has been determined, and other relevant objects. GMPLS also allows an upstream node to restrict the set of labels that a downstream node can choose. This is done by including a Label Set Object in the Path message. Label set is very useful for a node that cannot perform wavelength conversion, and ensures that a downstream node will assign a label that is acceptable to an upstream node.

Fig. 2 illustrates the sequence of steps leading to a lightpath establishment. In step (a), the source includes a label set $\{\lambda_1, \lambda_2, \lambda_3, \lambda_4, \lambda_5, \lambda_6\}$ in its Path message indicating that its source laser is capable of tuning its wavelength to any of the six possible lambdas. Upon receipt of a Path message, the first lambda router takes the intersection of the incoming label set and the set of available labels on its outgoing link, yielding an outgoing label set $\{\lambda_1, \lambda_4, \lambda_5, \lambda_6\}$, which will be forwarded in a Path message to the next hop. Each lambda router performs the same operation until the label set is propagated to the destination, which then selects one of the acceptable labels. After the destination has selected a label as indicated in step (b), label λ_5 is assigned at each hop toward the source. When all labels along the path have been assigned, a lightpath with λ_5 between the source and the destination is established (step (c)).

Fig. 3 illustrates a different scenario where wavelength conversion is needed. As before, the source includes a label set $\{\lambda_1, \lambda_2, \lambda_3, \lambda_4, \lambda_5, \lambda_6\}$ in its Path message. The first lambda router performs the intersection operation, and generates an outgoing label set $\{\lambda_1, \lambda_5, \lambda_6\}$. The second lambda router generates an outgoing label set $\{\lambda_1\}$. When the third lambda router performs the intersection operation between its incoming label set and its available labels on its outgoing link, the resulting set is empty. If the lambda router is not capable of performing wavelength conversion, it must abort the setup process by sending a PathErr message back to the source. If the lambda router has a wave-

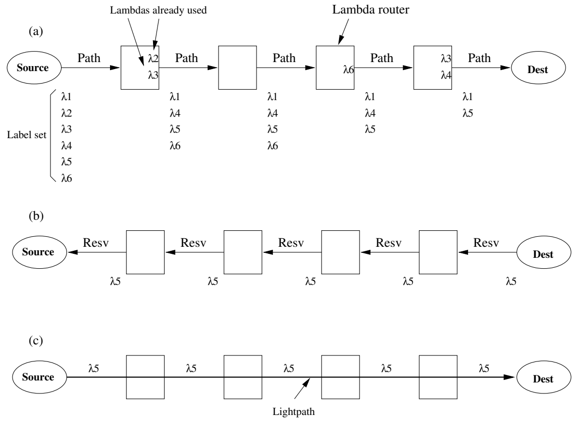

Fig. 2. How a lightpath is established with RSVP-TE signaling.

length converter, as shown in step (a), then it can instead perform the intersection operation between the set of labels that its wavelength converter can tune to and its available labels on its outgoing link. Thus, wavelength conversion has the advantage of completing the lightpath request that may otherwise be blocked. Note that in this scenario, the lightpath consists of two segments with different lambdas (i.e., λ_1 and λ_2).

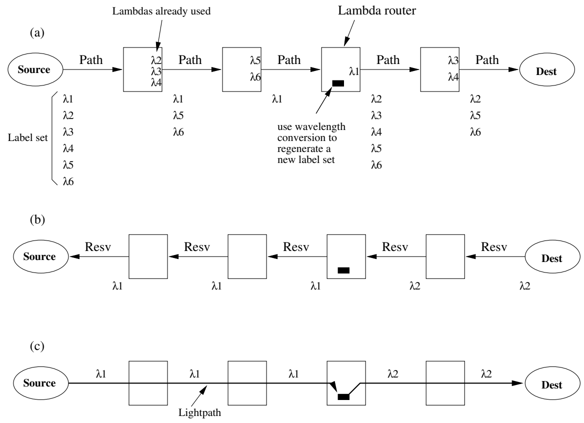

Fig. 3. The need of wavelength conversion.

III. GMPLS PERFORMANCE

In this section, we investigate the performance of GMPLS lightpath setup over a lambda-router network by means of simulation. We first describe the assumptions and the network topology used in the study. We then present the results and discussions.

A. Assumptions

We summarize the following assumptions in our simulation model:

- Dynamic environment: lightpaths arrive and depart at random.

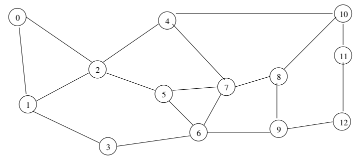

Fig. 4. Example of a meshed network.

- Lightpath requests arrive according to a Poisson process, and the lightpath holding time is exponentially distributed.
- Blocked lightpath request is cleared.
- The route of a lightpath follows either shortest path first (SPF) or constrained shortest path first (CSPF).
- Destination chooses a wavelength randomly.

In SPF, a source computes the shortest path to a destination based on static topology. In CSPF, a source first relies on dynamic link metrics to prune links that do not have available capacity or meet certain constraints. The source then computes the shortest path to a given destination based on the pruned topology. Note that SPF is implemented in OSPF, while CSPF may use an extension specified in OSPF-TE [5].

B. Network topology

We consider a typical meshed topology depicting a U.S. backbone network connecting major cities, as shown in Fig. 4. We assume that each bi-directional link consists of M fibers in each direction. Each fiber in turn is capable of handling up to W wavelengths. We also assume that each lambda router has C wavelength converters.

The traffic demand matrix $T = [T_{ij}]$ is an $N \times N$ matrix, where T_{ij} is the offered load from node i to node j and N is the number of nodes in the network. We use use a flexible representation of the traffic demand matrix given by $T_{ij} = m\,\bar{T}_{ij}$, for all i, j, where $\bar{T} = [\bar{T}_{ij}]$ is the base traffic demand matrix, and m is the multiplier. We employ the following base traffic demand matrix in our study:

$$\begin{pmatrix}
0 & 2 & 0 & 4 & 5 & 0 & 8 & 0 & 0 & 0 & 10 & 5 & 7 \\
2 & 0 & 2 & 5 & 3 & 3 & 8 & 4 & 1 & 1 & 10 & 7 & 10 \\
0 & 2 & 0 & 5 & 5 & 5 & 10 & 3 & 0 & 0 & 5 & 11 & 10 \\
4 & 5 & 5 & 0 & 5 & 5 & 10 & 5 & 5 & 5 & 10 & 5 & 8 \\
5 & 5 & 5 & 5 & 0 & 0 & 10 & 0 & 10 & 0 & 5 & 10 & 5 \\
0 & 3 & 5 & 5 & 10 & 0 & 6 & 5 & 10 & 6 & 5 & 6 & 10 \\
8 & 8 & 10 & 10 & 10 & 0 & 0 & 0 & 0 & 0 & 10 & 6 & 10 \\
0 & 4 & 3 & 5 & 0 & 5 & 0 & 0 & 5 & 10 & 8 & 8 & 10 \\
0 & 1 & 0 & 5 & 10 & 10 & 0 & 5 & 0 & 0 & 5 & 7 & 10 \\
0 & 1 & 0 & 5 & 0 & 6 & 0 & 10 & 0 & 0 & 7 & 5 & 10 \\
10 & 10 & 5 & 10 & 5 & 5 & 10 & 8 & 5 & 7 & 0 & 10 & 10 \\
5 & 7 & 11 & 5 & 10 & 6 & 6 & 8 & 7 & 5 & 10 & 0 & 5 \\
7 & 10 & 10 & 8 & 5 & 10 & 10 & 10 & 10 & 10 & 10 & 5 & 0
\end{pmatrix}$$

C. Simulation Results

We first consider the effect of path selection on the path blocking performance. Specifically, we compare SPF with CSPF. Here, we assume that each link has the same weight so that shortest path is equivalent to minimum hop.

In Fig.5, we plot the path blocking probability versus the total offered load for the SPF and CSPF cases, assum-

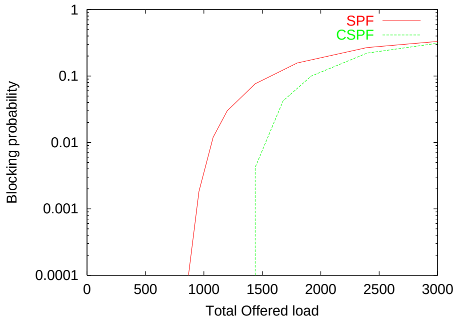

Fig. 5. Effect of path selection.

ing that wavelength converters are always available. As expected, path selection using CSPF outperforms that using SPF since CSPF takes advantage of the dynamic nature of the network state. In the following simulation study, we assume that the path selection algorithm is based on CSPF.

Fig.6 shows the effect of tunability of the source laser on the blocking probability, assuming that each lambda router does not have a wavelength converter. “No tuning” implies that the source laser is not tunable. In the “limited tuning” case, the source laser is assumed to be capable of selecting one of eight possible wavelengths. In the “full tuning” case, the source laser is capable of selecting any possible wavelength.

When the source laser is capable of tuning to more wavelengths, the source node can offer a larger label set to a lambda router, thereby reducing the chance of being blocked. It is interesting to see from the figure that even limited tuning offers a significant gain compared to no tuning. Since a full-tunable laser is much more costly than a limited-tunable one, the system designer may want to trade off the cost versus performance. Finally, note that if unlimited number of wavelength converters are available, then all three cases give identical performance.

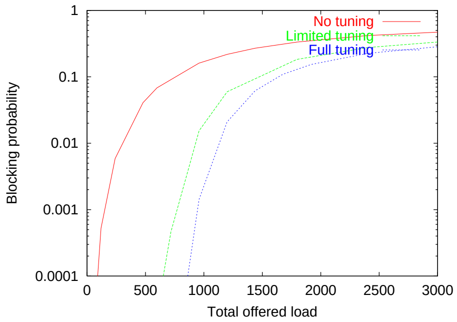

Fig. 6. Effect of source laser’s tunability.

We now study how the availability of wavelength converters affects the performance of a lambda-router network. In order to this, we fixed the offered load and vary the number of wavelength converters, C, at each lambda router. Fig. 7 shows the effect of wavelength converters on the blocking performance. Observe that the improvement offered by limited tuning over no tuning is roughly constant as C is varied. However, the improvement offered by full tuning over no tuning increases as C increases.

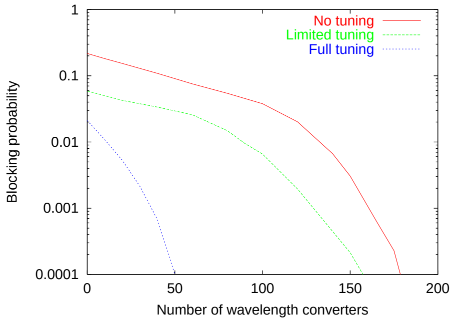

Fig. 7. Effect of wavelength converters.

It is interesting to see if adding wavelength converters beyond a certain number yields a diminishing return. Fig. 8 indicates that after a certain point, blocking due to lambda exhaustion dominates blocking due to insufficient number of wavelength converters.

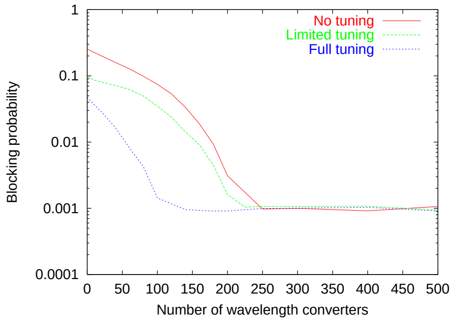

Fig. 8. Effective number of wavelength converters.

IV. Crankback

In this section, we describe an extension to the current GMPLS specification that is aimed to reducing the lightpath blocking probability. Recall that in GMPLS, a lightpath is blocked at a lambda router if the outgoing label set is empty and there is no more wavelength converter available at that lambda router. However, observe that if some upstream lambda router applies wavelength conversion, it is possible that the lightpath can be established.

Fig. 9 illustrates how crankback can successfully establish a lightpath that otherwise would fail. Observe that in step (a), the outgoing label set would be empty at the third lambda router if no wavelength converter is available. Thus, GMPLS will block such a request. To this end, we propose adding a new codepoint in the PathErr message of RSVP-TE to indicate to upstream lambda routers that

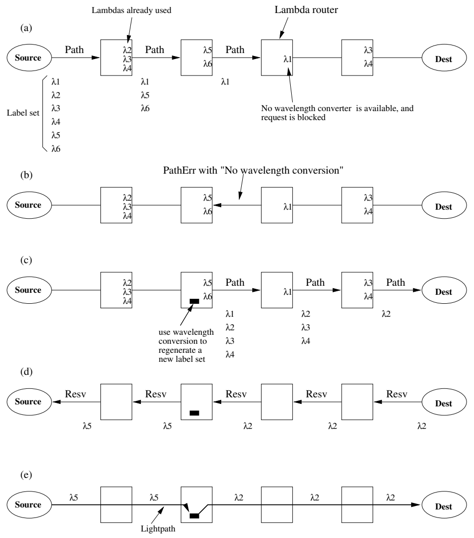

Fig. 9. Benefit of crankback.

the request cannot be completed because the sender has "no wavelength conversion". Upon receipt of this message, a lambda router checks if it has an available wavelength converter. If it cannot do wavelength conversion, it propagates the message upstream. If none of the upstream lambda routers can do wavelength conversion, then the request will have to be blocked. If an upstream lambda router can perform wavelength conversion, it will use its wavelength converter and issues a new outgoing label set taken from the intersection operation between the set of labels that its wavelength converter can tune to and its available labels on its outgoing link. Step (b) illustrates this process where the outgoing label set at the second lambda router is enlarged from $\{\lambda_1\}$ to $\{\lambda_1, \lambda_2, \lambda_3, \lambda_4\}$. Note that with this larger label set, the request can be successfully completed.

To illuminate the benefit of crankback, we consider a ring network consisting of seven lambda routers having 50 wavelength converters each interleaved with seven lambda routers with no wavelength converters. We also assume that the traffic is uniformly distributed; that is a lambda router has an equal probability of making a request to setup a lightpath to any other lambda routers. We further assume that $M = 4$ and $W = 40$.

Fig. 10 shows that the advantage of crankback is relatively small when fully-tunable source lasers are employed. The reason is that much of the blocking is due lambda exhaustion at an output link rather than due to insufficient number of wavelength converters.

When source lasers have no tuning capability, Fig. 11 indicates that crankback provides a significant improvement. This is due to the fact the crankback provides a large of degree of sharing of wavelength converters across the network.

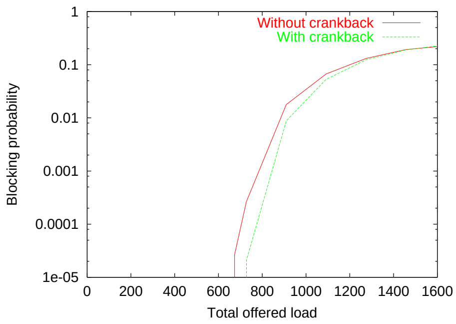

Fig. 10. Effect of crankback: full tuning.

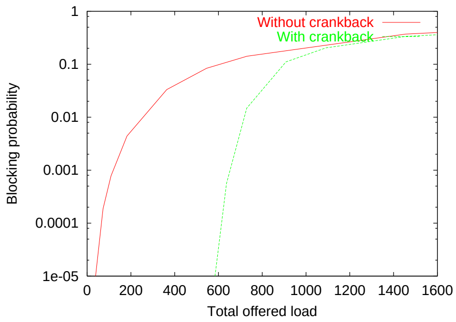

Fig. 11. Effect of crankback: no tuning.

V. Conclusion

The performance of GMPLS and its crankback enhancement has been investigated in the lambda-router network environment. We conclude that: 1) in a fully transparent network where no wavelength conversion is available, tunable lasers at the source provide significant performance benefit, 2) fully tunable lasers at each source coupled with a limited number of wavelength converters at each lambda router provides performance comparable to that of a fully opaque network, and 3) when tunable lasers are not universally available at the sources, crankback can provide a significant performance gain.

References

[1] Y. Xue et al., "Carrier Optical Services Requirements," IETF Draft, Work in Progress, Jul. 2001.

[2] D. Awduche et al, "Multi-Protocol Lambda Switching: Combining MPLS Traffic Engineering Control with Optical Crossconnects," IETF Draft, Work in Progress, Apr. 2001.

[3] E. Mannie et al., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture," IETF Draft, Work in Progress, Jun. 2001.

[4] P. Ashwood-Smith et al., "Generalized MPLS Signaling - RSVP-TE Extensions," IETF Draft, Work in Progress, Jul. 2001.

[5] K. Kompella et al., "OSPF Extensions in Support of Generalized MPLS," IETF Draft, Work in Progress, Jul. 2001.