The Effect of Wavelength Advertisement on the Performance of an Optical Routing Protocol

Bilal Khan* Dardo D. Kleiner† David Talmage* Abdella Battou∘
Center for Computational Science, Naval Research Laboratory, Washington D.C.
http://www.nrl.navy.mil/ccs/project/public/DC/web/

***Abstract*— This paper investigates the efficiency of wavelength selection in an optical network when it is conducted without knowledge of wavelength utilization, and compares it to the case when switches exchange wavelength availability through a routing protocol such as OSPF. We describe a series of experiments to determine the effect of wavelength advertisement on connection blocking probability in heterogeneous networks consisting of both wavelength converter and non-converter switches. Based on these experiments, we describe some consequences of advertising wavelength availability, and quantify when it is advantageous to advertise wavelength availability within the routing protocol.**

I. Introduction

All-optical switches are the future technology for transport backbones in next-generation networks. Scalability issues are very important and intimately tied to the problem of wavelength assignment. The majority of integrated DWDM in switching technologies offer up to 32 wavelengths per fiber. However, it is still not feasible to optically parse packet headers, thus, the control and data planes are decoupled and circuit switching is utilized.

A user requests the establishment of a call, an optical path or trail, via a signaling protocol such as OIF User Network Interface (UNI), RSVP, or CR-LDP. Upon receiving the request, the optical switch selects the outgoing port and a wavelength/interface and forwards the signaling packet to the next adjacent switch towards the specified destination. For networks without wavelength converters, the initial wavelength path is selected by the first switch. If at any point along the route the required wavelength becomes unavailable the call is blocked. Even when wavelength availability is advertised and wavelength conversion possible, blocking can occur because information propagation is not instantaneous.

This paper investigates the efficiency of wavelength selection (in terms of blocking probability) when done without any knowledge of wavelength availability, and compares it to the case where switches exchange this information through a suitably augmented routing protocol (e.g. OSPF-TE using opaque LSA).

Early IETF drafts by Chaudhuri et al. [3] and Basak et al. [1] specified optical network characteristics that should maintained in the switch routing database, classifying it in two categories: (i) *information advertised* using OSPF, e.g. the total number of active channels, preemptable channels, risk groups, etc., and (ii) *information kept locally*, e.g. available link capacity, association between fibers and wavelengths, etc. Later proposals [5], [6] made use of the Opaque OSPF LSAs [4] to implement advertisement of such characteristics.

At present, two routing protocols OSPF-TE and IS-IS-TE have been extended in the IETF to handle optical networks under the Generalized Multiprotocol Label Switching architecture (GMPLS RFCs 3471 and 3473). Since the two routing protocols are both link state protocols, we will only consider OSPF-TE in this paper.

Wang et al. [6] argued that the optical OSPF-TE protocol *should* advertise both the available wavelengths per fiber and the total available bandwidth. The rationale for the approach of Wang et al. is that in a network where a significant fraction of the switches are not wavelength-conversion capable, the probability of selecting a feasible source route decreases dramatically because wavelength continuity constraints at the transit switches render most of these routes infeasible.

On the other hand, RFC 3630 by Katz et al. postulates that an optical adaptation of the OSPF protocol should not advertise wavelength availability in link state advertisements since available wavelengths change frequently, and so presumably any performance increase would not be proportionate to the costs of increased control traffic. In this RFC, the problem of wavelength assignment is postponed to when the lightpaths are being signaled/provisioned.

In this paper, we seek to quantify the tradeoff between wave-

* Advanced Engineering & Sciences, ITT Industries., Alexandria VA
† Computer Integration & Programming Solutions, Corp., MD
∘ Lambda Optical Systems, Reston VA.

length advertisement policies and the performance of the optical network routing protocol.

II. Experiment Design

Our **performance metric** is *blocking probability*, the relative frequency with which a connection could not be established due to wavelength assignment infeasibility. Intuitively, we expect wavelength advertisement to decrease blocking probability. The aim to this paper is to give a quantitative description of the precise extent and circumstances in which this intuition is valid.

We consider the following **experimental parameters**.

1) *λ-Converter Density*. In any network, some percentage μ of the switches are capable of λ-conversion. We expect wavelength advertisement to yield greater benefit in networks where μ is large, since the λ-availability information then has greater potential to be utilized. In our experiments, we consider scenarios when $\mu = 0\%$, 25%, 50%, 75%, and 100%.
2) *Link delay ℓ*, and *mean distance (in hops) between switches d*. The product of these quantities ℓd is proportional to the mean time to propagate changes in routing information. When the product is large, we expect wavelength advertisement to yield less benefit, since λ-availability information will be more likely be stale when it is used in routing.
3) *Mean inter-connection arrival time z*. When the mean inter-connection arrival time is small, we expect wavelength advertisement to yield less benefit, since λ-availability information will be more likely be stale when it is used in routing.

We combine variables of (2) and (3) into a single parameter, the *normalized connection arrival rate δ*, defined as

$$\delta \;=\; \frac{\ell d}{z}.$$

Informally, δ measures the (average) number of connection setups witnessed during the (average) time for routing information to propagate between two switches. In our experiments, we vary δ between 0 and 5 in increments of 0.25.

Network topologies. We generated 250 random networks of 25, 100, and 200 switches respectively. Each of these 750 networks is constructed by the following process: (a) Place all of the switches at random positions inside the unit square. (b) Add links via the following random process: (b_1) Choose a random pair of switches that both have fewer than 8 links, (b_2) With probability that decays exponentially with the euclidean distance between u and v add a link (u, v) with 32 wavelengths. Repeat (b_1) and (b_2) until all switches have a degree of at least two and the network is connected. This process of sampling random networks is based on the techniques of Waxman [7].

Network traffic. For each experiment, we generate 25000 connection requests between random pairs of switches using a Poisson process. The connection holding time is exponentially distributed, with mean 1. All times (including mean inter-connection arrival time) are interpreted as values normalized to this fixed unity.

Our experiments are conducted using the Toolkit for Routing in Optical Networks (TRON), a simulation library for experimenting on routing protocols for optical networks [2]. TRON combines both peer and overlay models for routing information exchange, by making the router aware of the optical network details and collecting information to compute source routes. We extend TRON to explore the impact of advertising wavelength availability on the performance of optical link state routing protocols.

III. Results

The chart in Figure 1 describes how blocking probability changes as δ is increased from 0 to 5 in homogeneous networks (consisting entirely of converters or non-converters) under varying λ-advertisement policies.

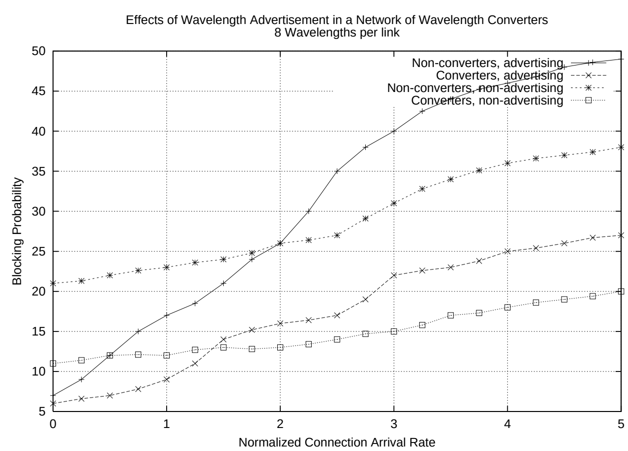

Fig. 1. Blocking probabilities in homogeneous λ-advertising and non-λ-advertising networks.

Comparing the *heights* of the curves in Figure 1, we conclude that (for either advertisement policy) λ-converter networks enjoy a lower blocking probability that is between 40-55% of what is experienced in non-λ-converter networks.

Comparing the *slopes* of the curves in Figure 1, we conclude that networks which do not advertise λ-availability are significantly less sensitive to δ. Specifically, networks of non-converters (resp. converters) which advertise λ-availability see blocking probability increase at roughly 8.4% (resp. 4.2%) for each unit increase in δ. In contrast, similar networks of that do not advertise λ-availability witness only an increase of approximately 3.4% (resp. 1.8%).

Considering the *intersections* of the curves in Figure 1, we conclude that the λ-advertising network outperforms the non-λ-advertising network whenever δ is smaller than a certain threshold T. The graphs show that $T_0 = 1.95$ for networks where 0% of the switches were λ-converters, and $T_{100} = 1.3$ for networks where 100% of the switches were λ-converters. Informally, the crossover point of the advertising/non-advertising curves exists because when $\delta \gg T$, wavelength availability information is stale by the time it is used, causing selection of λ-specific source routes which cannot be fulfilled.

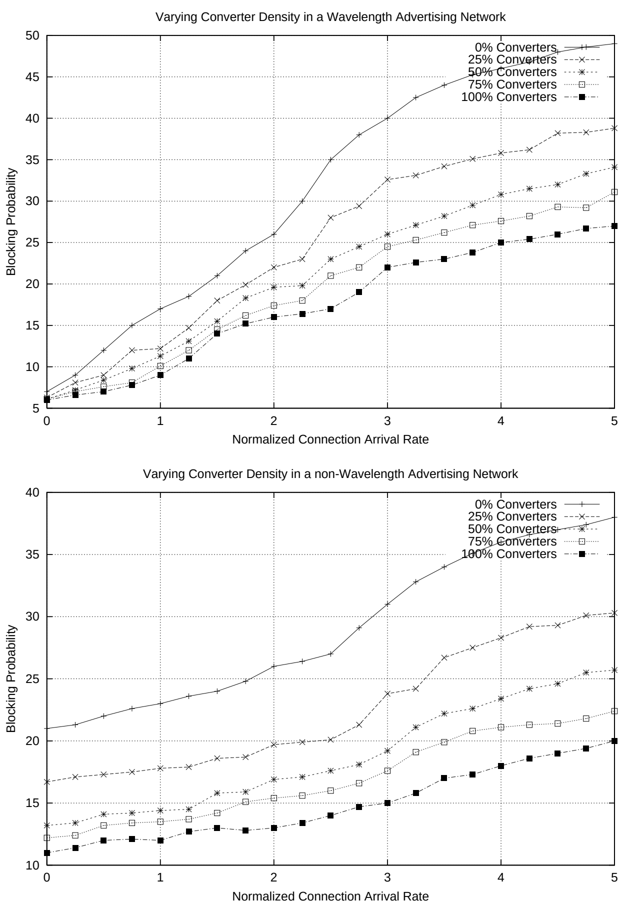

Fig. 2. Varying converter density in heterogeneous λ-advertising and non-λ-advertising networks.

Figure 2 summarizes the results of similar experiments for heterogeneous networks, where λ-converter density was varied between the two homogeneous extremes (i.e. from $\mu = 0\%$ to $\mu = 100\%$) with policies of λ-advertisement (top graph) and non-advertisement (bottom graph). These two charts demonstrate that networks consisting of only 25% λ-converters already enjoy an improvement of blocking probabilities that is nearly 50% of the optimal achieved in networks where all switches are λ-converters. Networks where 50% of the switches are λ-converters enjoy blocking probabilities that are nearly 80% of this optimal.

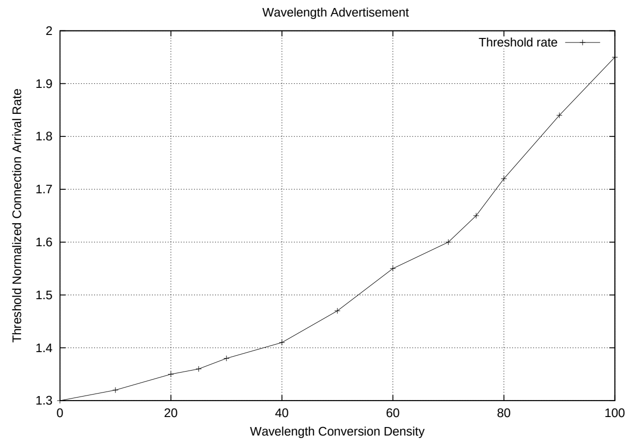

Fig. 3. the Threshold Normalized Connection Arrival Rate versus λ-Converter Density.

Finally, the curve in Figure 3 shows how T_μ (the Threshold Normalized Connection Arrival Rate) varies as a function of μ (the λ-Converter Density). Points on the curve are derived based on intersection points of corresponding curves from the two graphs in Figure 2. Suppose we are given a network which has λ-converter density μ, and for which we have estimated δ (based on statistical analysis of traffic or modeling, for example). Then Figure 2 provides a decision curve: If $\delta < T_\mu$, wavelength advertisement yields lower blocking probabilities; but if $\delta > T_\mu$ then it is more advantageous to not advertise wavelength availability.

IV. Conclusion

In conclusion, advertising wavelength is good strategy for networks where $\delta < T_\mu$, while networks where $\delta > T_\mu$ will benefit from not advertising wavelengths. Unfortunately it is not clear how δ will evolve as the technologies continue develop in the long term. For a given core network topology, increases in network link speeds drive the value of δ downward. On the other hand, connection arrival rates will increase over time, providing upward pressure on δ. The balance between the growth rates of these two quantities determines how δ will evolve over time, and as a consequence, whether advertising wavelength availability will be advantageous or not, in the long term.

References

[1] D. Basak, D. Awduche, J. Drake, and Y. Rekhter. Multi-Protocol Lambda Switching: Issues in Combining MPLS Traffic Engineering Control with Optical Corssconnects. Internet Draft, January, 2000.

[2] G. B. Brahim, B. Khan, A. Battou, M. Guizani, and G. Chaudhry. TRON: The Toolkit for Routing in Optical Networks. Proceedings of IEEE Globecom 2001, San Antonio, Texas USA; November, 2001.

[3] S. Chaudhuri, G. Hjalmtysson, and J. Yates. Control of Lightpaths in an Optical Network. Internet Draft, February, 2000.

[4] R. Coltun. The OSPF Opaque LSA Option. Request for Comments No. 2370, July, 1998.

[5] K. Kompella, Y. Rekhter, D. Awduche, J. L. G. Hjalmtysson, S. Okamoto, D. Basak, G. Bernstein, J. Drake, N. Margalit, and E. Stern. Extensions to IS-IS/OSPF and RSVP in support of MPL(ambda)S. Internet Draft, October, 1999.

[6] G. Wang, D. Fdyk, V. Sharma, K. Owens, G. Ash, M. Krishnaswamy, Y. Cao, M. Girish, H. Ruck, S. Bernstein, P. Nquyen, S. Ahluwalia, L. Wang, A. Doria, and H. Hummel. Extensions to OSPF/IS-IS for Optical Routing. Internet Draft, March, 2000.

[7] B. M. Waxman. Routing of multipoint connections. *IEEE Journal on Selected Areas in Communications*, 6:1617 – 1622, 1988.