Deterministic Models and Algorithms for Wavelength Routing Optical Networks

Ori Gerstel
Optical Networking Group, Tellabs Operations,
15 Skyline Drive, Hawthorne, NY 10532,
ori@tellabs.com

Abstract

This paper surveys recent results on the benefits of limited wavelength conversion in WDM ring networks. The maximum traffic load that can be supported is determined as a function of the amount of wavelength conversion provided in the network under a worst-case model where no blocking is allowed. The work addresses three traffic models with increasing level of dynamism: static, incremental and fully dynamic.

Optical wavelength routing networks are based on wavelength division multiplexing (WDM) transmission technology and on optical switching. Simply put, WDM allows for transmission of multiple streams on a single fiber (32 in commercial products) by sending each stream using a laser emitting light at a different frequency; optical switches use moving mirrors and other devices to switch the light to different ports at a node and thus provide optical routing. A combination of these techniques allow for end-to-end optical connections (called lightpaths) of up to 2.5-10 Gbps – orders of magnitude higher than conventional TDM networks (e.g., ATM or SONET). Another important advantage of these networks is that the lightpaths are transparent – insensitive to the protocol and bit rate of the data. See [1, 2] for further description of this technology.

Optical networks have made a huge leap in the last couple of years from a research curiosity to real-world systems that are being deployed by many telecom providers and are offered by many of the leading telecom vendors. This field has also been extensively studied by researchers of various disciplines, from optical device research to network-wide distributed protocols.

The current paper focuses on one of the most fundamental network-level problems: design and allocation of resources, or more specifically, the allocation of wavelengths to lightpaths. Since the wavelength resource is very limited (32 per fiber currently), this has been one of the more popular areas of research. In particular, it is possible to use this resource more efficiently if a lightpath can change the wavelength it uses along its route. A number of recent papers have studied the impact of wavelength conversion in WDM networks. While some of the papers compared full wavelength conversion to having no wavelength conversion at all (e.g., [3]), other papers have studied spare wavelength conversion, in which a small number of nodes has full wavelength conversion capabilities while other nodes have no wavelength conversion capability whatsoever (e.g., [4]).

Relatively few works have studied the effect of having a limited amount of wavelength conversion in all the nodes (see [5] for a probabilistic model). This abstract is based on [6] and surveys our recent results on this topic [9, 8, 12].

For practical reasons we assume bidirectional links and lightpaths which are realized using pairs of counter-propagating connections. For each lightpath we are given a route in the network and must assign a wavelength on each link along its route. Our main focus has been on rings (containing, say N nodes) due to their prevalence in modern telecom networks.

Our results apply to the situation where no blocking whatsoever is allowed; in other words, all lightpath requests must be supported. Moreover we make no assumptions about the traffic behavior, such as Poisson arrivals, exponential holding times etc. Instead, the traffic is characterized by its *maximum load*, L, which is the maximum number of lightpaths that can be present on any link at any time. This parameter represents the maximum utilization that can be achieved on the fiber links. The goal is to determine the maximum load that can be supported as a function of the number of wavelengths available and the degree of wavelength conversion.

We characterize the amount of wavelength conversion by a conversion degree d. Without any conversion $d = 0$. Fixed conversion is denoted by $d = 1$. In this case each node supports conversion of each incoming wavelength to a single predetermined outgoing wavelength. With $d = 2$, each wavelength can be converted to one of two wavelengths at a node. Let W denote the number of wavelengths on each link. The case $d = W$ corresponds to *full* wavelength conversion. Limiting the amount of conversion required can reduce hardware costs whether the nodes are implemented using electronic or optical switching and wavelength converters.

Clearly $L \leq W$ for any wavelength assignment scheme, since all lightpaths that share a link must be assigned different wavelengths. Also clearly $L = W$ if there is full wavelength conversion at each node. For the *static* case, all the lightpath requests are given up front. For the *dynamic* case, requests to add or delete lightpaths arrive in time and must be supported without rerouting already existing lightpaths. Two versions of the dynamic problem are considered: with and without deletions of lightpaths (the latter version is called incremental herein).

Static Case:

Consider first the case with no wavelength conversion. A simple greedy algorithm [7] can support request sets with load $L \leq \lfloor (W+1)/2 \rfloor$. Moreover there are request sets for which no algorithm can support a higher load, even with fixed wavelength conversion [8].

Consider next the case of limited wavelength conversion. [9] shows a network that has fixed conversion at two nodes and no conversion at the others that can support lightpath request sets with load $L \leq W - 1$, which is almost as good as having full wavelength conversion. By providing $d = 2$ limited conversion at two nodes [9], we can have $L = W$, which is as good as having full wavelength conversion.

Incremental:

Next consider the dynamic case where requests arrive but never get deleted. Without wavelength conversion, we can support request sets with load $L \leq \lfloor (W+2)/3 \rfloor$ [10] and there are examples where no algorithm can support higher loads [11]. With d limited conversion, $L \leq \lfloor \frac{W+d}{2} \rfloor$ is possible [12].

Dynamic:

Unfortunately, if deletions are allowed, the load that can be supported becomes much smaller. Without wavelength conversion, there are examples where no algorithm can support a load of more than $\lfloor 2W/\lceil \log_2 N \rceil \rfloor$ – see [8]. The simple first-fit algorithm can support loads L such that $L \leq \lfloor W/(\lceil 2.53 \log_2 N \rceil + 5) \rfloor$ without wavelength conversion [12]. A better algorithm in [12] supports loads of $L \leq \lfloor W/\log_2 N \rfloor$. With $d = 3$ limited conversion, we can support loads such that $L \log_2 L + 5L \leq W$ [12]. Note that a trivial algorithm can support a load of $L \leq d$.

Figure 1 plots the maximum supported load under different scenarios for $W = 32, N = 16$. Loads achieved by our algorithms, as well as proven upper bounds on the loads are plotted. The following are our main conclusions: (a) If worst-case guarantees are required, wavelength conversion helps significantly in rings, in contrast to results for probabilistic models [3]. (b) Very limited wavelength conversion helps a lot for the static case. In other words, rearranging existing lightpaths to accommodate new ones will enable us to support much higher loads, with limited conversion. (c) The guaranteed loads that can be supported is smaller for the dynamic assignment case without lightpath deletions and much smaller if deletions are allowed.

For dynamic traffic our results use a low conversion degree ($d < 4$) The main open issue is to design algorithms that efficiently make use of a larger conversion degree.

References

[1] R. L. Cruz, G. R. Hill, A. L. Kellner, R. Ramaswami, and G. H. Sasaki, eds., *IEEE JSAC / IEEE/OSA JLT: Special Issue on Optical Networks*, June 1996.

[2] R. Barry, ed., *Network Magazine, Special issue on optical networks*, Nov. 1996.

[3] R. Barry and P. Humblet, "Models of blocking probability in all-optical networks with and without wavelength changers," *IEEE JSAC / IEEE/OSA JLT: Special Issue on Optical Networks*, vol. 14, no. 5, pp. 858–867, 1996.

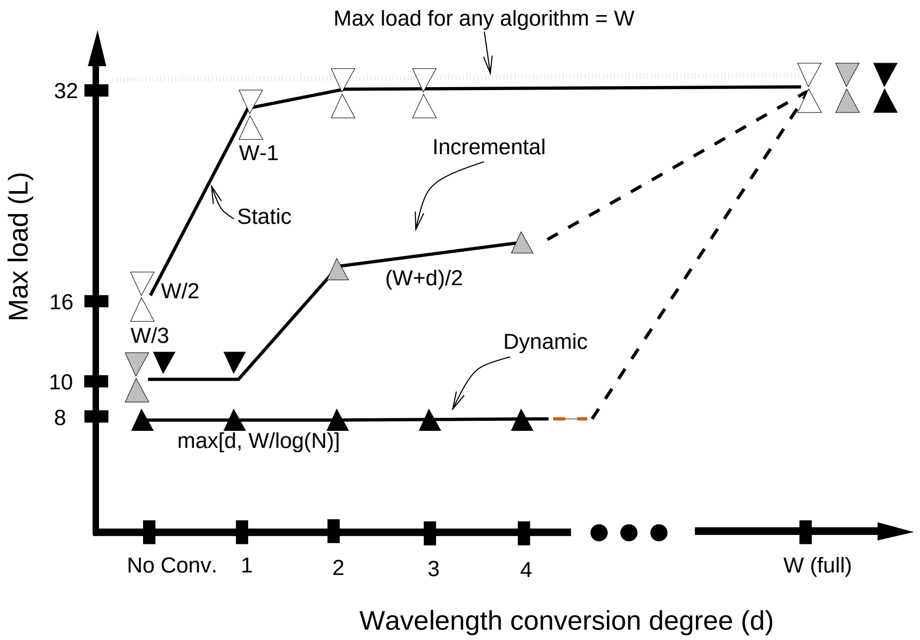

Figure 1: Upper and lower bounds on the load in a ring network with $W = 32$ wavelengths and $N = 16$ nodes. Upward pointing triangles correspond to achievable loads while downward pointing triangles correspond to upper bounds on the achievable load. White, gray, and black colors correspond to static, incremental, and dynamic models respectively.

[4] S. Subramaniam, M. Azizoglu, and A. K. Somani, "Connectivity and sparse wavelength conversion in wavelength-routing networks," in *IEEE Infocom'96*, pp. 148–155, 1996.

[5] J. Yates, J. Lacey, D. Everitt, and M. Summerfield, "Limited-range wavelength translation in all-optical networks," in *IEEE Infocom'96*, pp. 954–961, 1996.

[6] O. Gerstel, R. Ramaswami, and G. Sasaki, "Benefits of limited wavelength conversion in wdm ring networks," in *OFC'97 Tech. Digest*, 1997.

[7] A. Tucker, "Coloring a family of circular arcs," *SIAM Journal on Applied Math*, vol. 29, no. 3, pp. 493–502, 1975.

[8] O. Gerstel and S. Kutten, "Dynamic wavelength allocation in WDM ring networks," in *ICC'97*, June 1997. Also IBM research report No. RC 20462, May, 1996.

[9] R. Ramaswami and G. Sasaki, "Multiwavelength optical networks with limited wavelength conversion," in *IEEE Infocom'97*, 1997.

[10] M. Ślusarek, "Optimal online coloring of circular arc graphs," *Informatique Theoretique et Applications*, vol. 29, no. 5, pp. 423–429, 1995.

[11] H. Kierstead and W. Trotter, "An extremal problem in recursive combinatorics," *Congressus Numerantium*, vol. 33, pp. 143–153, 1981.

[12] O. Gerstel, G. Sasaki, and R. Ramaswami, "Dynamic channel assignment for wdm optical networks with little or no wavelength conversion," in *The 34th Allerton conf. on Communication, Control, and Computing*, 1996.