OFC'02 Anaheim, CA, March 17-22, 2002, Paper TuG2.

Analysis of Wavelength Conversion in All-Optical Express Backbone Networks

Jane M. Simmons
Corvis Corporation, 7065 Columbia Gateway Drive, Columbia, MD 21046
(301) 310-3410 (T), (301) 310-3017 (F), jsimmons@corvis.com

Abstract: We show that all-optical express networks can provide significant cost savings while achieving network efficiencies similar to those of O-E-O based networks, which provide for full wavelength conversion. We discuss this result from a topological and algorithmic standpoint, as well as demonstrate it quantitatively using the networks of several U.S. and European carriers.

OCIS codes: (060.4510) Optical communications; (060.4250) Networks

1. Introduction

The recent advent of commercially deployed optical networks employing all-optical switching and ultra-long-haul transmission has re-ignited interest in the debate over the role of wavelength conversion [1-3]. Such networks take advantage of optical bypass, where a signal can pass through a network node in the optical domain without being electronically terminated, to eliminate much of the electronics required in the network nodes. This significantly reduces the cost of the overall network, as well as the space and power requirements, and greatly simplifies and accelerates the provisioning process [4,5]. Optical bypass does impose a wavelength continuity constraint: if a signal is transmitted all-optically over a number of hops, it must use the same wavelength on each of the hops (assuming all-optical wavelength conversion is not available). While it is possible to come up with scenarios that show this constraint on wavelength assignment results in a loss of network efficiency, it is important to consider the practical architectures in which ultra-long-haul all-optical systems will be deployed. High capacity, all-optical ultra-long-haul systems are best suited for express backbone networks, which we refer to as all-optical express networks.

In this paper, we explore the physical characteristics of all-optical express networks, as well as their algorithmic component, to show that significant optical bypass can be achieved with almost negligible impact on network efficiency. These results are borne out by simulations using the networks of several U.S. and European carriers, where we compare the network efficiencies of all-optical networks to those of O-E-O based networks. In O-E-O based networks, each wavelength is electronically terminated at every node so that wavelengths can be selected independently on each hop. We conclude that the small amount of wavelength conversion that occurs in all-optical express networks (e.g., due to the small amount of needed regeneration) is sufficient to achieve virtually the same network efficiency as can be achieved with full wavelength conversion at all nodes.

2. Wavelength Conversion in All-Optical Express Backbone Networks

All-optical express networks serve as the lower tier of a two-tier architecture, as illustrated in Figure 1. Two-tier networks have been shown to be a very attractive architecture for networks with large amounts of traffic due to their improved economics and scalability [6-9]. In such architectures, the upper tier is used to collect, multiplex and groom fine granularity traffic, while the lower express tier carries large bundles of traffic over long distances.

The first characteristic to note is that express networks are generally relatively sparse, consisting of 15 to 30 nodes. Thus, connections in this layer do not pass through a large number of hops. Clearly, the fewer the number of hops traversed by a connection in an all-optical network, the easier it is to find a wavelength that is free on all of the hops. It is well documented that the need for wavelength conversion decreases as the number of hops decreases [10].

Consider all-optical express networks that serve as the backbone for networks of large geographic extent (for example, nationwide networks in the U.S., or pan-European networks). Even with ultra-long optical reach, some regeneration is typically required in these networks. For example, with a reach of 3,000 kms (which is significantly greater than the conventional reach of 400-600 kms), a cross-continental connection may require one or two regenerations. At these regeneration points, the signal must be electronically terminated. When the signal is converted back to the optical domain, there is freedom to select any wavelength; i.e., there is no requirement that the 'subconnections' comprising a connection be carried on the same wavelength. Wavelength conversion is achieved essentially 'for free' at regeneration sites.

It is important to note that it is not necessary to designate certain nodes as regeneration sites as is often done in studies on wavelength conversion. It is more cost-effective to regenerate on an as-needed basis rather than pre-determining that all traffic that passes through a particular node be regenerated. Moreover, intelligent selection of a regeneration site for a given connection can decrease conflicts arising from wavelength continuity requirements. For example, assume a connection traverses four links, each of distance 1,000 kms, and assume the optical reach is 3,000 kms. One regeneration is required for this connection; the regeneration can occur at any of the three intermediate sites of the connection. The regeneration point should be selected such that the resulting subconnections best align with the subconnections already in existence in the network. This minimizes the wavelength 'mixing' effect, which should result in fewer wavelength assignment conflicts.

Furthermore, it is not mandatory that the full optical reach of the system always be utilized. A small number of regenerations can be added solely for the purpose of improving network utilization. Obviously, if too much regeneration is added, the benefits of ultra-long reach are reduced. However, as is shown in the next section, a very small number of extra regenerations can appreciably increase network utilization.

Another important factor in mitigating the effects of the wavelength continuity requirement is assigning wavelengths in a judicious manner. Wavelength assignment strategies based on combinations of the well-known 'first-fit' and 'most-used' algorithms perform well. These strategies can be readily modified for use with ultra-long reach systems to account for the property that on some fiber types not every wavelength can attain the full optical reach.

3. Simulation Results

In order to quantify the role of wavelength conversion, simulations were performed using the topology of four carrier networks (three national U.S. networks and one pan-European network), as summarized in Table 1. Four scenarios were modeled for each of the four networks. First, we considered O-E-O point-to-point systems with 600-km reach and full wavelength conversion at every node. Second, we considered O-E-O point-to-point long-reach systems. In these systems, it is assumed the reach is long enough to span the inter-nodal distances but that full O-E-O wavelength conversion still occurs at each node. We also modeled two all-optical networking scenarios with 3,000 km optical reach: in one scenario, the optical reach is always utilized as fully as possible; in the other scenario, extra regenerations are permitted to be added in order to increase the network utilization (up to one extra regeneration per 'subconnection' is permitted). In all four scenarios, alternative path routing is permitted, subject to the restriction that alternate routes be no more than 3,000 kms longer than the shortest path (for the optical networking scenarios, this essentially means a circuitous path is allowed if it incurs no more than one extra regeneration). In each of the scenarios, the maximum capacity of the system is 40 wavelengths. (Note that for systems with more wavelengths, the relative performance of all-optical networks improves.)

In the first study, one thousand sets of demands were randomly generated for each network, where the selection of the demand endpoints was weighted based on the amount of add/drop the carrier forecast for each node. (The per-node traffic was considered as opposed to the actual demand node-pairs forecast by the carriers to better capture the uncertainty of traffic demand prediction). The demands were added to the network one by one; i.e., the routing and wavelength assignment occurred for each demand individually.

For each scenario, the network utilization at the 1% blocking level was determined; the results are shown in Table 2. The network utilization was calculated as the total bandwidth-distance product of the successfully routed demands (the distance was based on the shortest possible path for the demand, not necessarily the route taken). The results in Table 2 are all normalized to 1.0 for the full-wavelength conversion scenarios. (Both scenarios with full-wavelength conversion yield the same utilization.) As shown in the table, taking advantage of the full 3,000-km reach results in a loss of approximately 10% of the capacity. However, if a small number of additional regenerations are added, the loss of efficiency is on the order of about 2%. The results are consistent across the various networks.

The results in Table 3 demonstrate the benefit of optical bypass. The average number of transmitter/receiver (Tx/Rx) modules needed per connection in each scenario is shown in the table (one demand carried per connection). Compared to a conventional 600-km reach system, the optical networking scenarios eliminate approximately 75% of the Tx/Rx cards. Compared to an O-E-O point-to-point long-reach system, optical bypass eliminates approximately 60% of the Tx/Rx cards. Also, there is a very small difference in the number of Tx/Rx cards per connection in the two optical networking scenarios, indicating that very little extra regeneration is needed in order to achieve the higher network efficiency.

To further validate these results, a second study was performed where demands were continually added to and deleted from each of the four networks. This models the scenario of constant network churn. The demands were

Poisson distributed with exponential holding times. Enough processes were added to each system to achieve 1% blocking in steady state. Again, we compared network utilization and the average number of Tx/Rx modules per connection. The results are very similar to those shown in Tables 2 and 3.

4. Conclusion

Sparse wavelength conversion is essentially achieved for free in all-optical express networks by taking advantage of wavelength conversion at regeneration sites. As has been shown in numerous studies, sparse wavelength conversion is sufficient for achieving utilization levels close to those of networks with full wavelength conversion [10]. We discussed qualitatively why this should be expected for all-optical express networks, and further demonstrated this using four carrier networks. Thus, the significant savings in capital and operating expenses and the fast-provisioning benefits of all-optical networks can be attained with very little loss of network efficiency.

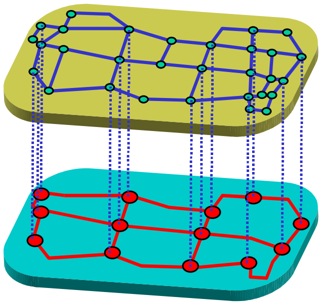

Fig. 1. A two-tier architecture, with a dense fine-granularity collection/grooming tier and a sparse, high-capacity express tier.

Table 1. Statistics of the Four Networks used in the Study

	Network 1	Network 2	Network 3	Network 4
Number of Nodes	13	18	25	19
Number of Links	15	20	29	28
Avg Link Distance (kms)	889	813	821	691
Max Link Distance (kms)	1750	1241	2054	1463

Table 2. Normalized Network Utilization

Wavelength Conversion Policy	Optical Reach	Network 1	Network 2	Network 3	Network 4
Full	600 km	1.00	1.00	1.00	1.00
Full	Max Internodal Distance	1.00	1.00	1.00	1.00
Convert at Regens	3,000 kms	0.92	0.90	0.89	0.96
Convert at Regens w/ extra Regens	3,000 kms	0.96	0.97	0.98	0.99

Table 3. Average Number of Transmitter/Receiver Modules Needed per Connection

Wavelength Conversion Policy	Optical Reach	Network 1	Network 2	Network 3	Network 4
Full	600 km	12.9	13.3	15.8	9.0
Full	Max Internodal Distance	6.9	8.0	8.7	5.7
Convert at Regens	3,000 kms	3.0	3.1	3.5	2.2
Convert at Regens w/ extra Regens	3,000 kms	3.1	3.2	3.6	2.3

[1] W. Van Parys, et al., "Quantifying the benefits of Selective Wavelength Regeneration in ultra long-haul WDM networks," *OFC'01*, March 19-22, 2001, TuT4.

[2] S. Chaudhuri, E. Bouillet, G. Ellinas**,** "Addressing Transparency in DWDM mesh survivable networks," *OFC'01*, March 19-22, 2001, TuO5.

[3] J. Strand, R. Doverspike, and G. Li, "Importance of Wavelength Conversion in an Optical Network," *Optical Networks*, May/June 2001.

[4] A. A. M. Saleh, "Transparent Optical Networking in Backbone Networks," *OFC '00*, March 5-10, 2000, ThD7, pp. 62-64.

[5] R. E. Wagner, et al., "The potential of optical layer networks," *OFC'01*, March 17-22, 2001, TuT3.

[6] A. Lardies and G. Ester, "Optimal network design of ultra-long-haul transmission networks," *Alcatel Telecommunications Review*, 2nd Quarter, 2001, pp. 142-147.

[7] E. Dotaro and A. Jourdan, "Optical Backbone Topology Design under Traffic Aggregation and Flexibility Constraints," *OFC'01*, March 19-22, 2001, TuG6.

[8] S. Singh, S. Chinburg, and M. Wendel, "Efficient network design using mesh and ultra long haul configuration," *NFOEC'01*, pp. 1006-1011.

[9] J.M. Simmons, "Economic and Architectural Benefits of Hierarchical Backbone Networks," *LEOS Summer Topicals*, July 24-28, 2000.

[10] R. Ramaswami and K. Sivarajan, *Optical Networks: A Practical Perspective,* Morgan Kaufmann Publishers, San Francisco, CA, 1998, Ch 8.