

MINIMAL COST WDM SONET RINGS THAT GUARANTEE NO BLOCKING

GALEN H. SASAKI
ORNAN GERSTEL

ABSTRACT

Minimal cost networks are presented when SONET UPSR and BLSR networks are deployed over a WDM fiber ring. The primary cost is the number of SONET ADMs, and a secondary cost is the number of wavelengths. The networks have been designed assuming a nonstatistical traffic model, and guarantees no blocking of traffic. Three network architectures are given that have low ADM costs. It is shown that networks that allow traffic to be cross-connected have lower ADM costs than those that do not, and that BLSR can lead to lower costs than UPSR.

1 - INTRODUCTION

Wavelength division multiplexing (WDM) technology exploits the wide communication bandwidth in single mode optical fibers. It enables each fiber to carry multiple optical signals, each at different *wavelengths*. In this way, WDM transforms a fiber into multiple *virtual fibers*. The technology has found commercial acceptance for applications where fiber is precious, such as in long-haul transport or fiber-scarce metropolitan areas.

Recent work [1, 2, 3, 7, 8, 9, 10, 11, 12, 14] has studied how WDM can be applied to optical fiber ring networks to support SONET traffic. Ring topologies are especially relevant because SONET uses it for protection purposes. The standard bodies have defined three types of SONET self-healing rings [6]: *unidirectional line switched rings* (UPSR), and *two- and four-fiber bidirectional line switched rings* (BLSR/2 and BLSR/4), as shown in Figure 1. Each network is composed of counter-rotating fiber rings with a high data rate, e.g., 10 Gbps for OC-192. The networks support full-duplex lower-speed SONET connections, e.g., OC-3 at 155 MBps, which we refer to as the *tributary traffic*. At each node are SONET *add/drop multiplexers* (ADMs), which multiplexes and demultiplexes the tributary traffic onto the optical lines using *time division multiplexing* (TDM). SONET ADMs are a major cost item in the network.

With WDM technology, multiple ring networks may be deployed onto a single set of fibers, with each wavelength corresponding to a "virtual" SONET ring. Figure 2 shows a node in a WDM network of two fiber-rings. There are two SONET ADMs, one per wavelength, each having port-side interfaces for the tributary traffic. The port-side interfaces may be connected to terminating points (sources/sinks) of tributary traffic or cross-connected to each other via connector cables or digital cross-connect systems (DCS).

Though there is a SONET ADM per wavelength, this may be unnecessary if a wavelength carries only *transit* tributary traffic through the node. In that case, the optical signal at the wavelength can be passed-through, avoiding an ADM along with its cost. Note that the cross-connection of ADMs (either by DCSs or cables) facilitates the *grooming* of tributary traffic between wavelengths to avoid ADMs. On the other hand, such cross-connection may mean more port-side interfaces and larger DCSs, which in turn may raise the cost of ADMs and DCSs.

In this paper, we will review some of the WDM ring archi-




FIGURE 1: SONET rings.

architectures that support SONET rings [1, 2, 3, 8, 10]. These architectures minimize cost, of which the primary one is the number of SONET ADMs. Another cost considered is the number of wavelengths. However, costs of port-side interfaces, optical ADMs, and DCSs are ignored.

The architectures are for dynamic, non-statistical models of tributary traffic. In general, traffic may arrive and depart at arbitrary times. Such a traffic model is appropriate when there are no accurate statistical models. The networks are *nonblocking*, i.e., they guarantee no blocking of traffic. Such networks are considered when the tributary traffic streams are high-bandwidth, generating significant revenue. Then provisioning of resources is preferable to blocking.

In Section 1.1, we give a brief description of SONET UPSR and BLSR networks (without WDM), and in Section 1.2, we discuss our network and traffic models. In Sections 2 and 3, we present simple architectures for UPSR and BLSR WDM networks, respectively. They allow traffic to be cross-

connected. It will be shown that this leads to lower ADM costs. In Section 4, we give our final remarks.

For the rest of this part of the introduction, we will discuss related work. Note [7, 9, 11, 12, 14] also present WDM ring architectures with low SONET ADM costs. However, the architectures are for *static* traffic, i.e., unchanging with time. Most deal with specific traffic patterns. The uniform pattern is studied by all. Hot spot traffic, distance dependent traffic, and general traffic patterns are considered in [2, 3], [12, 3], and [14], respectively. In addition, note that though the networks of [1, 8, 10] are for dynamic traffic, they also apply to static traffic with general traffic patterns.

We should also note that [1, 2, 3, 7] show that cross-connecting traffic may lower ADM costs. The network used to show the cost reduction is the single-hub network, discussed in Section 2.1. In [2, 3, 7], the traffic is static, while in [1], the traffic is dynamic.

1.1 - UPSR AND BLSR NETWORKS

A UPSR network is shown in Figure 1(a). It is composed of two counter-rotating fiber rings; referred to as the *working* and *protection* rings. If nodes *A* and *B* have a tributary traffic between them then *A* and *B* send their data on the working ring, and send copies on the protection ring. Thus, if the traffic is a full-duplex OC-3 connection then it will use 155 Mbps of bandwidth on all fibers in the network. Note that *A* (and *B*) receives two signals, one along each ring. It chooses the one that is better. In this way, when there is a single link failure, it can recover by switching to the available signal.

A BLSR/2 network is shown in Figure 1(b). It also has two counter-rotating fiber rings. Each fiber-pair between two nodes is considered as a full-duplex link. In this link, half the bandwidth carries *working* traffic, and the other half is for protection. If there is a single link failure, the working traffic that was carried on the link is looped-back around the ring using the protection bandwidth of the other links. Figure 1(b) shows a full-duplex tributary traffic stream, and notice that it occupies only one side of the ring. Typically, traffic is routed along shortest paths.

The BLSR/4 network, shown in Figure 1(c), is similar to BLSR/2 except that there are two pairs of counter-rotating fiber rings. One pair is used for working traffic and the other is used for protection.

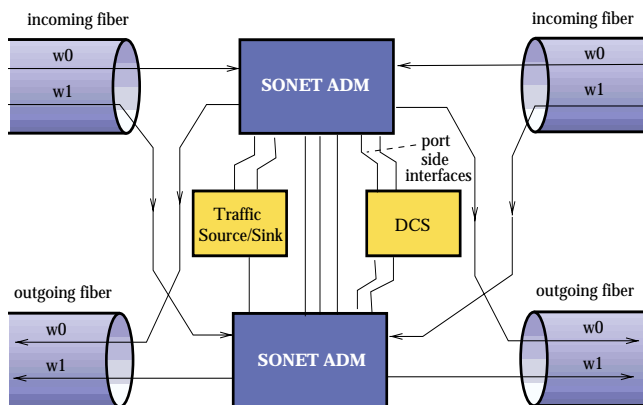


FIGURE 2: A node with a SONET ADM per wavelength.

1.2 - NETWORK AND TRAFFIC MODELS

We will consider two ring network models, one for WDM UPSR and the other for WDM BLSR/2 (we ignore BLSR/4 because it is similar to BLSR/2). The number of nodes in a network is denoted by N , and the nodes are numbered around the ring 0, 1, ..., $N-1$ in the clockwise direction. The number of wavelengths in the network is denoted by W . In a WDM UPSR network, the links are half-duplex and go in the clockwise direction, and in a WDM BLSR/2 network, the links are full-duplex. In a WDM UPSR network, each wavelength corresponds to a UPSR network. In a WDM BLSR/2 network, each wavelength corresponds to a BLSR/2 network. Thus, a WDM channel in a fiber link corresponds to a "virtual link" in a SONET ring network.

Each node has a collection of SONET ADMs corresponding to different wavelengths. There is also an *optical* ADM (OADM) which is used to add/drop wavelengths to the SONET ADMs. The OADM is assumed to be static (i.e., the wavelengths that are added/dropped are fixed). Recall, that port-side interface, OADM, and DCS costs are ignored. This is reasonable if the costs are proportional with the number of ADMs since they can be "absorbed" into the costs of the ADMs. This is also reasonable for DCS costs, if cross-connecting of tributary traffic is done manually with cables. Since OADM costs are ignored, any reference to an ADM will assume a SONET ADM.

To simplify the discussion, we assume the tributary traffic streams all have the same rate (e.g., all at OC-3). The number of traffic streams that may be multiplexed into a wavelength is indicated by a parameter g (for *granularity*). For example, if a wavelength carries data at the OC-48 rate (2.5 Gbps) and the tributary traffic streams are OC-3 (155 Mbps) then $g = 16$ because OC-48 rate = 16 x OC-3 rate. In the case of UPSR, the tributary traffic and WDM channels are considered unidirectional. However, the traffic is assumed to be symmetric so that for any pair of nodes A and B , if there are c unidirectional streams from A to B then there are c unidirectional streams from B to A . These unidirectional streams are paired to form c full duplex streams between the nodes. The number of unidirectional tributary traffic streams that can be multiplexed into a WDM channel is g . In the case of BLSR/2, tributary traffic and WDM channels are considered bidirectional, i.e., full-duplex. Actually, these bidirectional traffic streams and WDM channels are composed of two unidirectional ones in opposite directions. The number of full-duplex tributary traffic streams that can be multiplexed into a full-duplex WDM channel is $g/2$ because 50% of the bandwidth is reserved for protection. Notice that the tributary traffic provides full-duplex communication between end nodes.

The traffic is *dynamic*, i.e., traffic may arrive and terminate at arbitrary times. A special case of dynamic traffic is *incremental* traffic, where traffic streams never terminate. This models high-bandwidth connections that are unlikely to terminate in the near future.

The traffic is parameterized by integer values $\{t(i)\}$ for the nodes i . The value $t(i)$ is the maximum number of full-duplex tributary streams that may terminate at node i . We refer to this as the *node traffic constraint*. The $\{t(i)\}$ parameters indicate the number of ADMs required at a node, where larger $t(i)$ means less possible transit traffic and more required ADMs.

The network architectures for dynamic traffic described in the following sections will be *nonblocking*, i.e., an arriving traffic is never blocked. The traffic is always assumed to satisfy the node traffic constraint. A network will be referred to as *rearrangeably nonblocking* if it can always accommodate new tributary traffic and it is allowed to rearrange existing tributary traffic. It will be referred to as *wide sense nonblocking* if it can always accommodate new tributary traffic without disturbing existing tributary traffic, provided that the traffic set up according to some algorithm. It is assumed that any mechanism that cross-connects traffic at a node (whether manually or DCS) is wide-sense nonblocking. Notice that the $\{t(i)\}$ parameters provide lower bounds on the number of ADMs to ensure no blocking. In particular, node i requires at least $\lceil t(i)/g \rceil$ ADMs.

Wide-sense nonblocking is the more practical scenario today because many real systems will not disturb an existing communication connection and risk disruption of service. However, rearrangeably nonblocking leads to lower cost and resources. In addition, rearrangeable systems may be acceptable for Internet traffic, which has less stringent requirements for quality of service.

2 - WDM UPSR NETWORKS

We will consider two kinds of WDM UPSR networks. First, we consider those that do not have any constraints on W . A simple architecture will be given called the *single-hub* network. It has a small number of ADMs. Then we will consider networks that have $W = W_{\min}$, where W_{\min} is the minimum number of required wavelengths.

2.1 - UNLIMITED W

The single-hub network is shown in Figure 3. This is a wide-sense nonblocking network that has been well studied in [1, 2, 3, 7, 8, 9]. It has one *hub* node, which is node 0 in the figure. The hub has an ADM on every wavelength, and cross-connects tributary traffic between nodes. The other nodes, i.e., the *non-hubs*, send their tributary traffic to the hub, where the traffic is possibly cross-connected, and finally forwarded to its destination. Each node i reserves bandwidth to accommodate its possible $t(i)$ traffic. A simplifying special case is for each node i to have $t(i)$ be a multiple of g . Then each node i has $t(i)/g$ wavelengths dedicated to it, and has an ADM at each

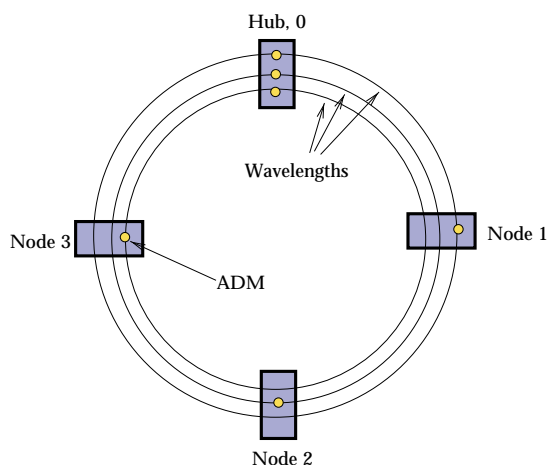


FIGURE 3: A single-hub WDM UPSR network with three wavelengths.




FIGURE 4: Average number of ADMs per node for WDM UPSR networks.

of its wavelengths. The general case when $t(i)$ is not necessarily a multiple of g is considered in [3, 9]. Then the problem of minimizing the number of ADMs is NP-Complete [3, 9].

Figure 4 shows the number of ADMs per node for the single-hub network when $g = 16$ and $t(i) = N-1$ for all nodes i . The figure also has a couple of lower bounds, labeled LB_{no-cc} and LB_{cc} . Both of these bounds were designed for static uniform traffic, where each pair of nodes has the same number of tributary traffic streams between them. They also apply to nonblocking networks as long as the $\{t(i)\}$ are at least the number of traffic streams terminating at the nodes.

LB_{cc} is valid if traffic is allowed to be cross-connected at nodes, and LB_{no-cc} is valid if traffic is disallowed from cross-connecting. As expected LB_{cc} is smaller than LB_{no-cc} . Single-hub is below LB_{no-cc} for most instances of N . This means that it has a smaller number of ADMs than any WDM UPSR network that disallows traffic from being cross-connected. Therefore, cross-connecting traffic can lower ADM costs, and in some cases by a significant amount. For example, when $N = 16$, the single-hub network has 63% of the ADMs than the lower bound LB_{no-cc} . Also note that the single-hub can be very efficient. For example, when $N = 16$, it has nearly the same number of ADMs as the lower bound LB_{cc} .

It is unclear how tight the lower bound LB_{no-cc} is. Thus, in the figure we provide the number of ADMs of a network that efficiently uses ADMs but disallows traffic from being cross-connected [2, 3]. This network is designed for static uniform traffic. It efficiently uses ADMs by partitioning the nodes into groups, and have pairs of groups share wavelengths to communicate. We refer to this network as the *Grouped* architecture. Notice that the single-hub ADM costs can be significantly less than the Grouped architecture. For $N = 16$, it can have less than half the cost. In addition, the single-hub network is nonblocking for dynamic traffic, while the

Grouped architecture is designed for static traffic.

Finally, we should mention that the single-hub network has also been proposed for another kind of traffic where all tributary traffic is to be connected to the hub node. This can model a local access network, where the hub is the access point to a backbone network [2, 3]. It can also model the case where one node is a *hot spot* for traffic [1].

2.2 - LIMITED W

The single-hub network has a small number of ADMs, but can have a large number of wavelengths. To illustrate this, consider the following simple lower bound on the number of required wavelengths in a WDM UPSR network:

$$W_{min} = \left\lceil g^{-1} \left\lfloor \sum_i t(i)/2 \right\rfloor \right\rceil \quad (1)$$

This bound follows from the observation that each full-duplex tributary traffic stream between two nodes takes up bandwidth around the entire ring. Since the total number of tributary traffic streams is $\lfloor \sum_i t(i)/2 \rfloor$, and g of them can fit in each wavelength, W_{min} is a lower bound. Getting rid of the “floors” and “ceilings”, W_{min} is approximately $\sum_i t(i)/2g$. This is about half of $\sum_{i \neq 0} t(i)/g$, which is the number of wavelengths in a single-hub network if $\{t(i)\}$ are multiples of g .

A network that has W_{min} wavelengths and efficiently uses ADMs is the *incremental network* described in [8, 10]. It is rearrangeably nonblocking. (It is also wide-sense nonblocking for incremental traffic.) A simple version of this network is when even nodes have ADMs on all W_{min} wavelengths, while odd nodes i have ADMs on $\lceil t(i)/g \rceil$ wavelengths. Therefore, around the ring, nodes alternate between having the maximum and minimum numbers of ADMs. Figure 5 shows an example when $\frac{t(i)}{g} = 1$ for all nodes i .

Figure 6 shows the average number of ADMs per node for the simplified incremental network for different N , when $g = 16$, and for all nodes i , $t(i) = 2(N-1)$. Also in the figure are two lower bounds, $LB_{re,no-cc}$ and $LB_{ws,no-cc}$ from [1]. The bounds are for nonblocking networks that disallow cross-connected traffic and use W_{min} wavelengths. $LB_{re,no-cc}$ is for rearrangeably nonblocking networks, and $LB_{ws,no-cc}$ is for wide-sense nonblocking networks. Note that the number of ADMs for the incremental network is smaller than both bounds which means that cross-connecting traffic can lower ADM costs. For example, when $N = 16$, the reduction is by 17%




FIGURE 5: An incremental network with $W_{min} = 3$.

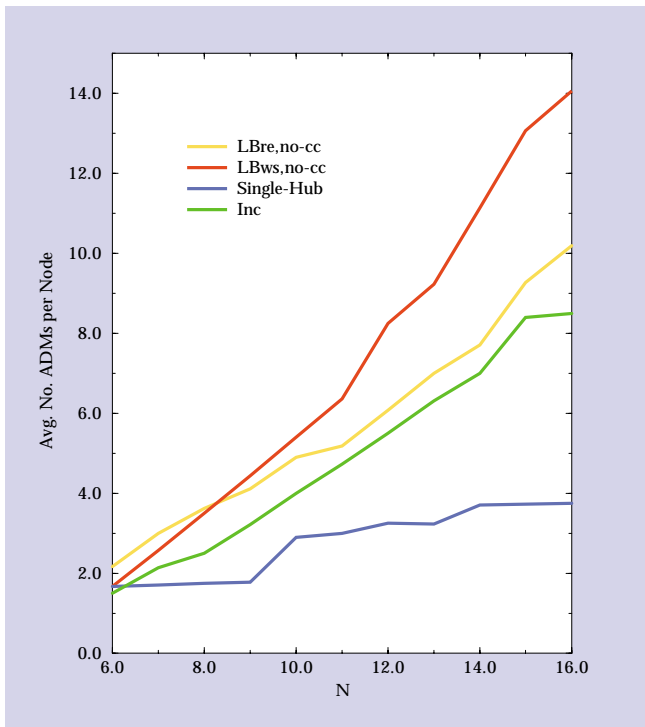


FIGURE 6: Average number of ADMs per node for WDM UPSR networks with a minimum number of wavelengths.

from $LB_{re,no-cc}$. Also note that $LB_{re,no-cc}$ can be significantly lower than $LB_{ws,no-cc}$ indicating that rearrangeably nonblocking networks can have much lower costs than those that are wide sense nonblocking.

The figure also has the number of ADMs for the single-hub network. The single-hub network has a considerably lower number of ADMs compared to the lower bounds or the incremental network. However, recall that the single-hub network uses approximately twice as many wavelengths.

Finally, we should note that [1] describes nonblocking WDM UPSR networks with W_{min} wavelengths when traffic is not cross-connected. They are for the case when $t(i)$ is constant over all nodes i . Design algorithms were provided that determine where to remove ADMs and still keep the network rearrangeably nonblocking.

3 - WDM BLSR/2 NETWORKS

In this section, we consider WDM BLSR/2 networks. As in the previous section, we will consider those without constraints on wavelengths and those that use minimal numbers of wavelengths. We will compare the ADM requirements with those for UPSR.

3.1 - UNLIMITED W

The single-hub network architecture can be employed in WDM BLSR/2 networks. It is wide-sense nonblocking. It has a hub node, say node 0, and all other nodes connect their traffic streams through it. Note that each wavelength can support g full-duplex tributary traffic streams from a node, just as in WDM UPSR networks. Thus, it requires about the same number of wavelengths and ADMs as a single-hub WDM UPSR network. At present, we do not know of any other architecture that is nonblocking and has a smaller number of ADMs.

The advantage BLSR/2 has over UPSR is that tributary traffic may be routed on either side of the ring. Efficient routing typically leads to a lowering of the bandwidth requirements. This is true, as we shall see in the next section.

3.2 - LIMITED W

In this section, we will consider a network architecture that has approximately the same number of ADMs as a single-hub network. However, it requires much less numbers of wavelengths. The network is the *double-hub* network [8]. It has two *hub* nodes and, without loss of generality, let them be nodes 0 and h , for some h . Each hub has ADMs on every wavelength. Each node has its tributary traffic go directly to the hubs. An example of a double-hub network is shown in Figure 7.

It is assumed that for each node i , $t(i)$ is a multiple of 4. Nodes 0, 1, ..., $h-1$ will be referred to as *side 1* of the ring, and the rest of the nodes will be referred to as *side 2*. Hub 0 is part of side 1, and hub h is part of side 2. Each non-hub node i reserves bandwidth on its side of the ring to accommodate a possible $t(i)/2$ traffic to each hub. Each hub node i reserves bandwidth on its side of the ring to accommodate a possible $t(i)/2$ traffic to the other hub. A simplifying special case is when, for each node i , $t(i)$ is a multiple of g . (The general case when $t(i)$ is not necessarily a multiple of g is discussed in [8].) Then node i has $t(i)/g$ wavelengths dedicated to it on its side of the ring, and it has an ADM on each of its wavelengths. The number of wavelengths in the network is

$W_{dbl} = \max \{ \sum_{i=0}^{h-1} t(i)/g, \sum_{i=h}^{N-1} t(i)/g \}$, and the number of ADMs is $2W_{dbl} + \sum_{i \neq 0,h} t(i)/g$. Notice that the hub nodes can be selected to optimize the number of wavelengths and ADMs. If the $t(i)$ values are the same, the hubs should be chosen to be at opposite ends of the ring.

The network has been shown to be rearrangeably nonblocking [8]. This follows from the observation that it can operate as a three-stage Clos network, which is rearrangeably nonblocking [4, 13].

Figures 8 and 9 show the numbers of wavelengths and ADMs per node, respectively, required by the double-hub network when $g = 16$ and for each node i , $t(i) = 2(N-1)$. They are labeled *Double-bslr* in the figures. Notice that whenever N is even, $t(i)$ is not a multiple of 4, which is required by the

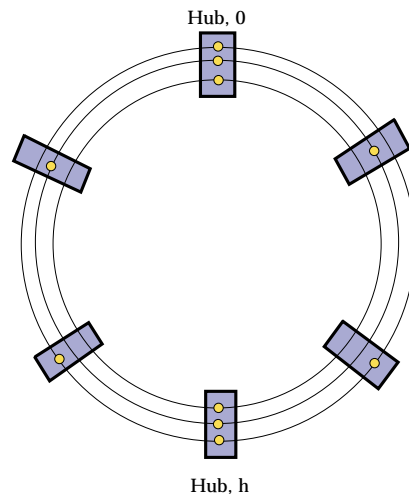


FIGURE 7: A double-hub WDM BLSR/2 network with three wavelengths.




FIGURE 8: Wavelength requirements of the double-hub BLSR/2 network compared with lower bounds.

network. Thus, when N is even, the actual value of $t(i)$ used to compute the wavelength and ADM requirements for the network is a slightly larger value of $2N$.

Figure 8 compares the number of wavelengths in the double-hub network with two lower bounds, LB_{blsr} and LB_{upsr} . LB_{blsr} is for BLSR/2 networks that allow cross-connected traffic. It follows from a worst case traffic where nodes communicate only with nodes that are directly opposite on the ring, making each stream traverse $\lfloor N/2 \rfloor$ links. Since there are $N(N-1)$ traffic streams, the average number of traffic streams over a link is $\lfloor N/2 \rfloor (N-1)$. Since there are at most $g/2$ streams in a wavelength, the lower bound is $\lceil \lfloor N/2 \rfloor \frac{2(N-1)}{g} \rceil$.

LB_{upsr} is for WDM UPSR networks that allow cross-connected traffic. It is the same as W_{min} in Subsection 2.2. This bound does not apply to the double-hub network, which is a BLSR/2 network. However, it is in the figure to compare the double-hub with WDM UPSR networks.

Both lower bounds have about the same values, though LB_{blsr} can sometimes be a little lower. The double-hub network requires a few more wavelengths than the lower bounds, and sometimes has the minimum number of wavelengths, such as when N is 8 and 12.

Figure 9 shows the average number of ADMs per node for the double-hub network. Also shown in the figure are the ADM costs for the single-hub WDM UPSR network, labeled *Single-upsr*. Note that the double-hub network has about the same ADM cost as the single-hub WDM UPSR, which is fairly minimal. The figure also has a lower bound for WDM UPSR and it is labeled $LB_{re,no-cc,upsr}$. This bound is from [1] and applies to rearrangeably nonblocking WDM UPSR networks that disallow cross-connected traffic and use minimum wavelengths. In other words, it is the bound $LB_{re,no-cc}$ in Section 2.2. Notice that the ADM costs for $LB_{re,no-cc,upsr}$ is much higher than double-hub. This implies that BLSR/2 tech-

nology and cross-connecting traffic can decrease ADM costs considerably while keeping wavelength requirements minimal. For instance, when $N = 16$, the double-hub network has 37% of the ADMs of $LB_{re,no-cc,upsr}$, while requiring only one additional wavelength over LB_{upsr} (a less than 7% increase). In the case of $N = 12$, the double-hub has 52% of the ADMs of $LB_{re,no-cc,upsr}$ and the same number of wavelengths as LB_{upsr} .

4 - FINAL REMARKS

We considered nonblocking WDM UPSR and BLSR networks that have a minimal number of ADMs and allow traffic to be cross-connected. These are the single-hub, double-hub, and incremental networks. The cross-connection allows the single-hub and incremental networks to have ADM costs that are lower than any WDM UPSR network that disallows cross-connection.

If wavelengths are plentiful then the single-hub network leads to a low number of ADMs for both UPSR and BLSR. It is also wide-sense nonblocking, which makes it more practical. However, it uses a lot of wavelengths, about twice the minimum.

For UPSR, the incremental network uses the minimum number of wavelengths. It is rearrangeably nonblocking. It is also wide-sense nonblocking for incremental traffic. Thus, it is closer to being practical than if it were only rearrangeably nonblocking. Unfortunately, it has many more ADMs than single-hub, though it has less ADMs than networks that disallow traffic cross-connection. For BLSR, the double-hub network demonstrates that BLSR technology can be better than UPSR. The number of its wavelengths is at or near the minimum number of wavelengths for any WDM UPSR network, and its ADM cost is close to the single-hub network.

An interesting open problem is to find wide-sense nonblocking networks that have a minimal number of wave-

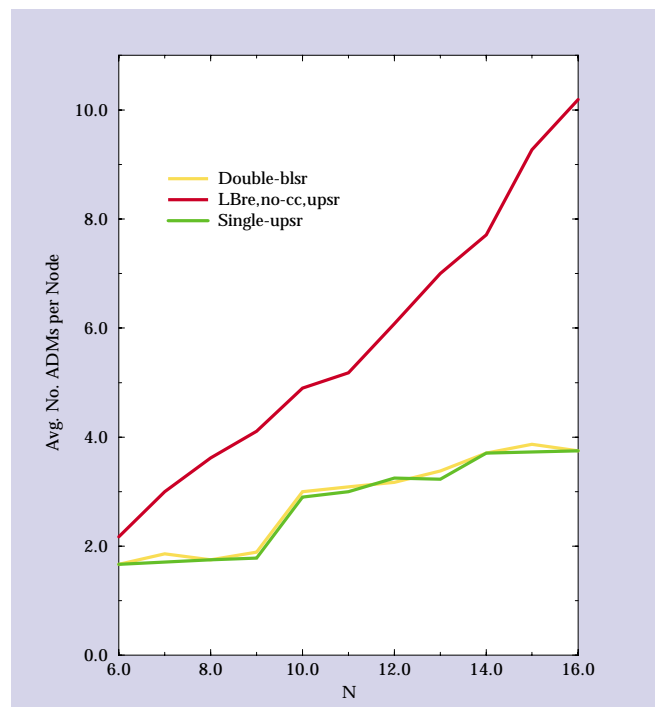


FIGURE 9: The number of ADMs per node for the double-hub WDM BLSR/2 network compared with a lower bound and the single-hub WDM UPSR network.

lengths and small ADM costs. Lower bounds in [1] indicate that WDM UPSR, wide-sense nonblocking networks have inherent high ADM costs. However, the bounds assume no traffic cross-connection. Hence, it would be interesting if efficient networks result from traffic being cross-connected.

ACKNOWLEDGEMENTS

Effort sponsored in part by the Defense Advanced Research Projects (DARPA) and Rome Laboratory, Air Force Materiel Command, USAF, under agreement number F30602-97-0342. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright annotation thereon. The views and conclusions contained therein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Defense Advanced Research Projects Agency (DARPA), Rome Laboratory, or the U.S. Government.

REFERENCES

- [1] R. Berry and E. Modiano, "Minimizing electronic multiplexing costs for dynamic traffic in unidirectional SONET rings," *Proc. ICC*, June 1999.
- [2] A. Chiu and E. Modiano, "Reducing electronic multiplexing costs in unidirectional SONET/WDM ring networks via efficient traffic grooming," *Proc. Globecom 98*, Sydney Australia, Nov. 1998.
- [3] A. Chiu and E. Modiano, "Traffic grooming algorithms for reducing electronic multiplexing costs in WDM ring networks," *J. Lightwave Technology*, vol. 18, no. 1, Jan. 2000, pp. 2-12.
- [4] A. Deguid, "Structural properties of switching networks," Brown University, Progress Report BTL-7, 1959.
- [5] M. Garey and D. Johnson, *Computers and Intractability: a guide to the theory of NP-Completeness*, W.H. Freeman and Co., New York, 1979.
- [6] I. Haque, W. Kremer, and K. Raychauduri, "Self-healing rings in a synchronous environment," *SONET/SDH: a sourcebook of synchronous networking*, Eds. C.A. Siller and M. Shafi, IEEE Pres, New York, pp. 131-139, 1996.
- [7] O. Gerstel, P. Lin, and G. Sasaki, "Combined WDM and SONET network design," *Proc. Infocom 99*, New York, NY, Apr. 1999.
- [8] O. Gerstel, R. Ramaswami, and G. Sasaki, "Cost effective traffic grooming in WDM rings," *Proc. Infocom 98*, San Francisco, CA, Apr. 1998.
- [9] E. Modiano and A. Chiu, "Traffic grooming algorithms for minimizing electronic multiplexing costs in unidirectional SONET/WDM rings networks," *Proc. CISS '98*, Princeton, NJ, Feb. 1998.
- [10] G. Sasaki, O. Gerstel, and R. Ramaswami, "A WDM ring network for incremental traffic," *Proc. Thirty-Sixth Annual Allerton Conference on Communication, Control, and Computing*, Monticello, IL, Sept. 1998.
- [11] J. Simmons, E. Goldstein, and A. Saleh, "On the value of wavelength add/drop in WDM rings with uniform traffic," *Proc. OFC '98*, San Jose, CA, Feb. 1998.
- [12] J. Simmons, E. Goldstein, and A. Saleh, "Quantifying the benefits of wavelength add-drop in WDM rings with distance independent and dependent traffic," *J. Lightwave Technology*, vol. 17, no. 1, Jan. 1999, pp. 48-57.
- [13] D. Slepian, "Two theorems on a particular crossbar switching," unpublished manuscript, 1952.
- [14] X. Zhang and C. Qiao, "An effective and comprehensive solution to traffic grooming and wavelength assignment in WDM rings," *SPIE Proc. of Conference on All-Optical Networking*, vol. 3531, Nov. 1998, pp. 221-232.

Galen H. Sasaki

sasaki@spectra.eng.hawaii.edu

Galen Sasaki received the B.S. degree in electrical engineering from the University of Hawaii in 1981, and the M.S. and Ph.D. degrees in electrical and computer engineering from the University of Illinois at Urbana-Champaign in 1984 and 1987, respectively. From 1987 to 1992, he was an Assistant Professor with the Department of Electrical and Computer Engineering, University of Texas at Austin. Since 1992, he has been an Associate Professor in Electrical Engineering at the University of Hawaii. His research interests are in communication networks (with an emphasis on optical and high speed networks), performance evaluation, and optimization algorithms. Dr. Sasaki has served as an editor for IEEE journals, and has been on the program committee for Infocom.

Ornan Gerstel

Ornan (Ori) Gerstel received the B.A., M. Sc., and D.Sc. degrees from the Technion, Israel. After finishing his D.Sc. he joined the Optical Network Systems Group at IBM T.J. Watson Research Center, then moved with the group to develop optical networking products with Tellabs Operations. There he served as the system and software architect for the Tellabs Optical Networking Group, building the TITAN 6100 metro DWDM product line. Recently he has left Tellabs to join Xros, a startup building all-optical cross-connects. Ori has served on the program committees of Infocom and OFC, and has published more than a dozen journal papers, and a similar number of patents. He also served as a guest editor for a IEEE JSAC issue on optical networks, and as an editor for the IEEE Communication Surveys journal. His research interests include network architecture, fault-tolerance and protection, and network design problems in optical networks.