

Quantifying the Value of Wavelength-Add/Drop in WDM Rings

Jane M. Simmons and Adel A. M. Saleh
 AT&T Labs - Research, 100 Schulz Drive, Red Bank, NJ 07701

Abstract

In WDM rings, the value of wavelength-add/drop rises steeply with ring size and/or internodal demand. We present a bound on attainable equipment savings, and present a methodology for performing the bundling for rings carrying uniform all-to-all traffic, and rings carrying uniform hubbed traffic.

Key Words

(060.4250) Networks; (060.4510) Optical communications; (060.4230) Multiplexing

1. Introduction

There is a great deal of interest in designing networks comprised of Wavelength Division Multiplexed (WDM) Synchronous Optical Network (SONET) rings, due to their large capacity and inherent reliability. One drawback to such rings is the potentially large amount of equipment necessary for their deployment. In a conventional WDM SONET ring, each wavelength is dropped at every node in the ring, and is terminated in a SONET Add/Drop Multiplexer (ADM). However, deploying a Wavelength Add/Drop Multiplexer (WADM) in a node allows wavelengths to be selectively dropped. If a wavelength carries no traffic originated at, or destined to, a particular node, then that wavelength can optically bypass the node, thereby eliminating a costly SONET ADM. This paper quantifies the benefit of optical bypass for two different traffic scenarios. First, we consider uniform all-to-all ring traffic, as initially presented in [1]. Second, we examine rings with uniform hubbed traffic. In both scenarios, optical bypass can significantly reduce the number of required ADMs; the savings increases rapidly with the number of nodes in the ring and the amount of traffic.

2. Uniform All-to-All Traffic

Consider N nodes on a bidirectional ring, where each node sends one λ/T worth of traffic to each other node, for some

integer T (e.g., T equal to 1 means that a full wavelength of traffic is sent from each node to every other node). For simplicity, we assume that N is odd, and that traffic is always routed over the shortest path. We assume the SONET ADM operates on a full wavelength, where each wavelength is partitioned into T time slots of granularity λ/T . In general, for uniform all-to-all traffic and N odd, the minimum number of time slots required to carry all of the traffic is [2]:

$$\frac{N^2 - 1}{8} \quad (1)$$

In Figure 1, one particular time slot assignment is shown for a ring with N equal to 7. Each dashed circle represents a time slot, with the squares indicating the add/drop locations on that time slot. For example, a time slot with squares at nodes 1, 4, and 5 indicates that on this time slot, Node 1 sends λ/T worth of traffic to Node 4, Node 4 sends λ/T to Node 5, and Node 5 sends λ/T to Node 1. Note that the add/drop locations in the six time slots in the figure allow each node to transmit one λ/T to every other node.

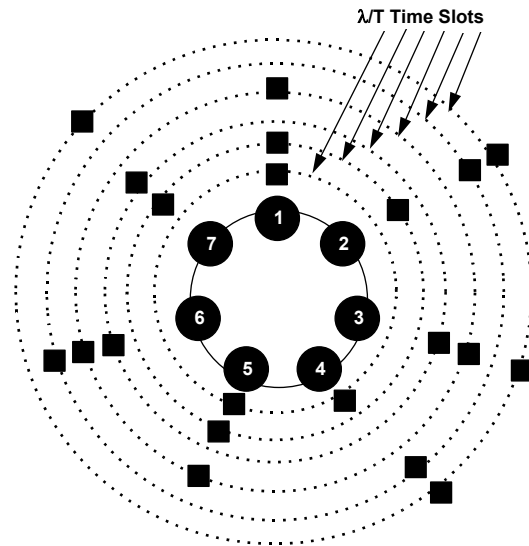


Figure 1. Schematic illustration of a time-slot assignment for a 7-node SONET ring. Time slots are added and dropped at locations marked by a square.

2.1 Full Wavelength Traffic

If each node transmits a full wavelength to each other node ($T=1$), then a time slot is equivalent to a full wavelength, i.e., it is not necessary to bundle together multiple time slots to form a wavelength's worth of traffic. In this scenario, the minimum number of required ADMs is $\frac{N(N-1)}{2}$, one for each source/destination pair. By contrast, in the absence of wavelength-add/drop (i.e., an ADM in every node for every wavelength), one requires a total of $\frac{N(N^2-1)}{8}$ ADMs. Thus, for full wavelength internodal demand, wavelength-add/drop allows one to save a fraction $\frac{N-3}{N+1}$ of the ADMs.

2.2 Sub-Wavelength Traffic

For more moderate demand ($T > 1$), time slots need to be bundled together to form wavelengths. We use the notation $[i, j, \dots k]$ to indicate that a given time slot add/drops at nodes $i, j, \dots k$. Suppose that each node has a half-wavelength for each other node, so that two time slots must be combined to form a wavelength ($T = 2$). For N odd and uniform all-to-all traffic, it can be shown that each wavelength requires a minimum of 5 add/drops. Furthermore, it can be shown that the most 'efficient' union that yields 5 add/drops is obtained by combining a 3-add/drop slot with a 4-add/drop slot (not all such combinations yield 5 add/drops). Thus, the optimal ring bundling arises by forming all wavelengths from such combinations, if possible.

As an example, for $N = 7$, one optimal bundling is:

Wavelength #1 =

Union of $[1, 4, 5]$ & $[1, 2, 5, 7] = [1, 2, 4, 5, 7]$

Wavelength #2 =

Union of $[3, 6, 7]$ & $[1, 3, 5, 6] = [1, 3, 5, 6, 7]$

Wavelength #3 =

Union of $[2, 4, 6]$ & $[2, 3, 4, 7] = [2, 3, 4, 6, 7]$

This bundling requires a total of 15 ADMs. Without optical bypass, an ADM is required at each of the 7 nodes, on each of the 3 wavelengths. Wavelength-add/drop thus permits a terminal-equipment savings of 6/21. This is optimal, but it

is less than the $\frac{N-3}{N+1} = \frac{1}{2}$ attainable with $T=1$, illustrating

that finer routing granularity undermines the value of wavelength-add/drop.

This combinatoric approach can be extended to traffic granularities finer than $\lambda/2$. For example:

Granularity $\lambda/4$: any combination of four time slots yields at least 7 add/drops, with the most efficient grouping being a union of three 4-add/drop time slots and one 3-add/drop time slot.

Granularity $\lambda/8$: any combination of eight time slots yields at least 9 add/drops, with the most efficient grouping

being a union of seven 4-add/drop time slots and one 3-add/drop time slot.

Using the above approach, one can obtain bounds on the fractional terminal-equipment savings obtainable from wavelength-add/drop. These are plotted in Figure 2 for rings with between 3 and 15 nodes, roughly the range of practical interest. The savings are seen in general to rise steadily as both the size of the network and the inter-node demand increases. Wavelength-add/drop is seen to be particularly valuable in rings whose inter-office traffic exceeds about one-quarter of a wavelength. Indeed, the saturation of these curves suggests that at such high traffic levels, there is little incentive to extend the SONET ring node limit beyond 16. However, for inter-node demand below one-quarter wavelength, only modest savings are seen. Indeed, for inter-node demand of $1/16$ wavelength, no savings occur in rings with fewer than 11 nodes. To extract significant value from wavelength-add/drop in this low-traffic-volume regime would require an extension of the SONET ring standard beyond its current limit of 16.

Figure 2 in general suggests that as inter-node traffic volume increases, as is expected to occur swiftly in the public network, the value of wavelength-add/drop will increase accordingly. We should note, finally, that for all points in Figure 2 excepting the three encircled ones, we were able to explicitly construct traffic bundlings exemplifying the computed optima.

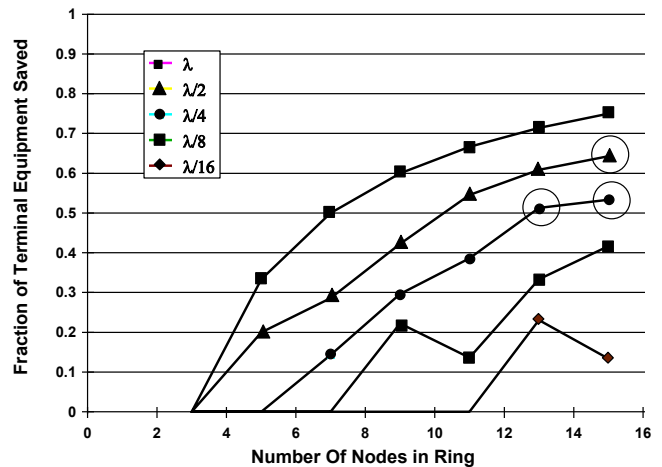


Figure 2. Maximum fractional terminal-equipment savings achievable via wavelength-add/drop in a WDM SONET ring. The parameter represents the fraction of a wavelength from each node to each other node. Terminal equipment savings rise with node-to-node demand and ring size.

2.3 Super-Node Approximation

In Section 2.2, we described a combinatoric approach for bounding the number of ADMs required for the case of uniform all-to-all λ/T traffic, where $T > 1$. Should rings with yet larger number of nodes become important in the

future, this exact combinatoric approach would likely prove unwieldy. In this section, we present an alternative technique that performs extremely well in approximating the number of required ADMs.

For the case of full-wavelength uniform traffic, the minimum number of required ADMs was easily seen to be precisely $\frac{N(N-1)}{2}$, one for each source/destination pair.

This holds for the case of N even as well as N odd. Using the $\frac{N(N-1)}{2}$ bound, we can generate an approximation to the number of required ADMs for the scenario where the granularity is finer than a full wavelength.

Assume the granularity is λ/T , for some $T > 1$. Using the technique of [3], we arbitrarily partition the N nodes into groups of \sqrt{T} nodes and form ‘super-nodes’ from each of these groupings, such that we have N/\sqrt{T} super-nodes, as illustrated in Figure 3 for N=10 and T=4. Consider sending λ/T worth of traffic between each of the \sqrt{T} nodes comprising one super-node and each of the \sqrt{T} nodes comprising any other super-node. In total, then, there are T connections between each super-node pair, with the aggregate traffic between each pair being exactly a full λ . Treating each super-node as if it were a single node, we have N/\sqrt{T} nodes with a full wavelength of traffic between each pair; we know that the minimum number of ADMs required for such an arrangement is:

$$\frac{N}{\sqrt{T}} \left(\frac{N}{\sqrt{T}} - 1 \right) \quad (2)$$

Each super-node is comprised of \sqrt{T} real nodes, thus, each ADM required at a super-node translates into \sqrt{T} required ADMs at real nodes. The total number of required ADMs is then:

$$\frac{N \left(\frac{N}{\sqrt{T}} - 1 \right)}{2} \quad (3)$$

Although the quantities involved (e.g., \sqrt{T}) do not always represent integers, we use (3) to approximate the minimum number of ADMs required for arbitrary N and T. Note that the super-node construction does not account for traffic between the nodes comprising a super-node, thus we expect this approximation to underestimate the number of required ADMs.

Figure 4 plots the number of required ADMs using both the exact enumeration technique of the previous section and the approximating super-node method. There is relatively good agreement between the curves, although, as expected, the super-node approximation underestimates the minimal required number of ADMs. What is remarkable is the accuracy one nevertheless obtains from the super-node approximation over the plotted range.




Figure 3. The super-node approximation, for 10 nodes and granularity $\lambda/4$. The 10 nodes are partitioned into 5 super-nodes, each with 2 nodes. Each node within a super-node sends one $\lambda/4$ to each node in every other super-node, for a total of one full λ between each super-node pair. Only the traffic between super-nodes 1 and 2 is shown in the figure.




Figure 4. Number of required ADMs in WDM SONET rings. The solid lines indicate a lower bound on the number of required ADMs using the exact enumeration technique, while the dotted lines represent an approximation to the minimal required number, using the super-node approximation.

3. Uniform Hubbed Traffic

Above we considered rings with uniform all-to-all traffic, where each node sends λ/T to each other node in the ring. In this section, we consider a ring with hubbed traffic, which is typical of an access ring, and which approximates the traffic in many backbone rings. We assume the ring has N nodes, of which one is designated as the hub node, or Central Office (CO), and the remaining $N-1$ nodes are referred to as Access Nodes. Figure 5 shows a ring with a total of 5 nodes, including one CO.

It is assumed that Access Nodes can communicate directly only with the CO. The CO is likely to be an egress point to a backbone network; thus, any traffic destined for the backbone must first be sent to the CO. If traffic needs to be exchanged between two Access Nodes, the CO serves as an intermediary node. For example, in Figure 5, if Node 1 needs to communicate with Node 4, the traffic is sent from Node 1 to the CO (where it may be groomed with other traffic destined for Node 4), and then from the CO to Node 4. We assume the ADMs are bidirectional, 2-fiber (i.e., they can source/sink one wavelength in the clockwise direction and one wavelength in the counter-clockwise direction).




Figure 5. A hubbed ring with a total of 5 nodes, including one Central Office (CO). Only the traffic between Node 1 and the CO is shown.

3.1 Bidirectional Hubbed Rings

We consider the case where a total of λ/T worth of traffic is exchanged between each Access Node and the CO, for some integer T (e.g., T equal to 1 means that a full wavelength of traffic is sent from each Access Node to the CO, and a full wavelength of traffic is sent from the CO to each Access Node). In this section, we consider bidirectionally routed rings, where working traffic is routed over the shortest path. Protection capacity is also provisioned; protection traffic is routed in the opposite direction as the working traffic, as shown in Figure 5.

The minimum number of bidirectional wavelengths needed to carry all of the traffic is:

$$\left\lceil \frac{N-1}{T} \right\rceil \quad (4)$$

With optical bypass, the minimum number of required ADMs is 1 ADM at each of the Access Nodes and

$\left\lceil \frac{N-1}{T} \right\rceil$ ADMs at the CO. Compared to requiring an ADM at every node on every wavelength, optical bypass yields a fractional savings of:

$$1 - \frac{N-1 + \left\lceil \frac{N-1}{T} \right\rceil}{N \left\lceil \frac{N-1}{T} \right\rceil} \geq 1 - \frac{T+1}{N} \quad (5)$$

This lower bound is plotted in Figure 6 for various values of N and T . The savings provided by optical bypass can be very significant, and increases with larger rings (i.e., larger N) and coarser traffic granularity (i.e., smaller T).

We next consider the more general scenario where M of the N total nodes are designated as Central Offices, with N being an integral multiple of M , and the COs being evenly spaced around the ring. We assume the M COs are redundant, such that Access Nodes direct their traffic to the closest CO and protection traffic is routed to the nearest CO in the reverse direction (it is likely M is small). In this arrangement, the maximum fraction of ADMs that can be saved using optical bypass, compared to an ADM at every node on every wavelength, is approximately $1 - \frac{M(T+1)}{N}$.

Again, the savings increases with larger ring size and coarser traffic granularity; it decreases with increasing M .




Figure 6. Fraction of ADMs saved in a hubbed ring with one CO. The percentage savings increases rapidly with increasing ring size and with coarser traffic granularity.

3.2 Unidirectional Hubbed Rings

Access rings often are designed to support only unidirectional routing, where all working traffic is routed in one direction, and all protection traffic is routed in the opposite direction. For a ring with one central office, optical bypass provides the same benefit as it does with bidirectional routing; unidirectional and bidirectional routing are equally efficient in this scenario. With M COs, and unidirectional routing, the fractional ADM savings provided by optical bypass is approximately $1 - \frac{M+T}{N}$.

Note that the savings decreases more slowly with increasing M as compared to the bidirectional case, due to unidirectional routing remaining inefficient as M increases (i.e., it does not take advantage of shortest path routing).

4. Comparison Between Uniform All-to-All Traffic and Uniform Hubbed Traffic

We can compare the fraction of ADMs that can be saved in single-CO and double-CO hubbed rings to the fraction that can be saved in the uniform all-to-all traffic scenario. In all cases, we assume that the total number of nodes, N , is the same. We also assume that the total traffic is the same. The uniform all-to-all scenario includes traffic between all possible source/destination pairs, while the hubbed traffic scenario is limited to communication pairings only with the COs. Thus, we would expect that the fraction of ADMs that can be saved in the all-to-all traffic scenario is less than that in the hubbed scenario.

To illustrate this, we consider the all-to-all traffic case where each of the N nodes in the ring sends $\lambda/8$ to every other node in the ring. The total amount of traffic, counting both directions, is: $N(N-1)(\lambda/8)$. In order for the total traffic to be the same in the single-CO hubbed case, each of the $N-1$ access nodes must send $N\lambda/16$ to the CO (and the CO sends an equal amount back). In the double-CO hubbed case, each of the $N-2$ access nodes must send/receive $[(N-1)/(N-2)] N\lambda/16$ to/from the CO.

We focus on the regime where N falls between 9 and 15, inclusive; i.e., where the fractional ADM savings in the uniform all-to-all case is non-zero (refer back to Figure 2). For the corresponding hubbed rings, each of the access nodes sends and receives greater than $\lambda/2$ worth of traffic. In general, we showed above that the fractional savings decreases as the traffic granularity becomes finer; thus, we can lower bound the savings in the hubbed cases by considering a traffic granularity of $\lambda/2$.

Using the results from Sections 2.2 and 3.1, we compare the fractional ADM savings in Figure 7. As expected, the fractional savings for the hubbed ring scenarios are greater, especially for a single CO ring.




Figure 7. Fractional equipment savings compared for uniform all-to-all $\lambda/8$ traffic, and single- and double-CO hubbed traffic. The total traffic sent in all scenarios is the same.

5. Tradeoff Between Minimizing ADMs and Minimizing Wavelengths

In the uniform all-to-all and uniform hubbed scenarios enumerated above, where the inter-nodal traffic demand is λ/T , for some integer T , the minimum number of ADMs can be attained using the minimum number of wavelengths. However, for more general traffic demand, where T is not an integer, it may be necessary to deploy more wavelengths in order to achieve the minimum number of ADMs.

Consider, for example, a hubbed-ring with 4 Access Nodes, 1 Central Office, and nodal traffic demand of $3\lambda/4$. The minimum number of required wavelengths is 3. However, in order to achieve this minimum, at least one Access Node must split its traffic onto multiple wavelengths, thereby requiring more than one ADM. For example, three Access Nodes could each send $3\lambda/4$ on a single wavelength and one Access Node could send $\lambda/4$ on each of three wavelengths; a total of 9 ADMs (including the 3 at the CO) is required in this scenario. If, instead, each Access Node sends $3\lambda/4$ on a separate wavelength, the number of required wavelengths is 4, but the number of required ADMs is only 8.

Thus, as this example illustrates, if T is not an integer, there is potentially an architectural tradeoff between minimizing ADMs and minimizing wavelengths.

6. Conclusions

Optical bypass potentially provides significant benefit in decreasing the amount of terminal equipment required in WDM SONET rings. For the uniform all-to-all scenario and the uniform hubbed scenario, it was shown that the benefit increases rapidly with ring size and with the amount of traffic. For the hubbed scenario, the benefit decreases as the number of hubs in the ring increases. In addition to quantifying the maximum terminal-equipment savings attainable using wavelength-add/drop, we devised a traffic-bundling methodology that enables one to produce constructions that achieve the optimal savings.

7. Acknowledgments

We would like to thank Evan Goldstein for several useful discussions related to this work.

8. References

1. J.M. Simmons, E.L. Goldstein, A.A.M Saleh, "On The Value of Wavelength-Add/Drop in WDM Rings With Uniform Traffic," *OFC'98*, San Jose, CA, February, 1998.
2. A.F. Elrefaie, "Multiwavelength Survivable Ring Network Architectures," *ICC'93*, Geneva, Switzerland, Paper 48.7, 1993.
3. A.A.M. Saleh, "Transparent Optical Networks for the Next Generation Information Infrastructure," *OFC'95*, San Diego, CA, February 26-March 3, 1995.