ICC'99, June 6-10, 1999, vol. 1, pp. 591-596, Vancouver, British Columbia, ©1999 IEEE

The Value of Optical Bypass in Reducing Router Size in Gigabit Networks

Jane M. Simmons and Adel A. M. Saleh
AT&T Labs - Research

***Abstract* - With the expected explosive growth in data traffic, the required processing demands at network nodes, e.g., in routers, will likely present a significant bottleneck. Optical bypass potentially alleviates this bottleneck by allowing a wavelength to bypass a node if it contains no traffic that needs to be processed there. We explore the value of optical bypass in gigabit networks by measuring the fractional decrease in required router size one can expect. We consider both ring and mesh topologies, and both uniform and distance-dependent traffic. We show that the reduction in router size can be very significant, with the fractional savings increasing as the size of the network increases and the amount of traffic increases, indicating that this technique will scale well in future networks.**

I. Introduction

With the expected explosive growth in data traffic, backbone and access networks will be strained in terms of both transport and processing requirements. Wavelength Division Multiplexing (WDM) has emerged as the favored transport technology for backbone networks, where systems will be deployed with tens of wavelengths, each carrying a 2.5 or 10 Gbit/sec signal. WDM will eventually be extended to the regional access environment, especially to provide access for large business customers. While point-to-point WDM systems will likely meet the enormous transport demands, the processing demands at network nodes, e.g., in routers and switches, are likely to present a significant bottleneck.

There are currently two major approaches for alleviating this nodal bottleneck. One approach is to dramatically increase the throughput of routers and switches, for example, through the development of application specific hardware and the implementation of sophisticated scheduling algorithms [1]. While switches and routers operating at hundreds of Gbit/sec are certainly possible, it is not clear how well this approach will scale. Also, this approach does not reduce the port count, and hence, to a large extent, the physical size, of the switch. Alternatively, one can alleviate the nodal bottleneck to some degree by lessening the processing requirements. For example, in Multiprotocol Label Switching, the processing requirement of IP routers is reduced by assigning labels to traffic streams, and performing more efficient label-based switching as opposed to IP-Header-based routing [2, 3]. Again, there may be scalability issues with such schemes, as the labeling and switching processes themselves may become bottlenecks; also, the overall port count of the "label switching router" is not reduced.

A more scalable solution to the nodal bottleneck is to make use of WDM optical networking (as opposed to point-to-point WDM), and *optically bypass* nodes, as shown in the simple example of Fig. 1. This figure depicts a very small network, with just three wavelengths ($\lambda 1$, $\lambda 2$, and $\lambda 3$) and three routers. Assume that some of the traffic sent by Router 1 on $\lambda 1$ and $\lambda 3$ needs to be processed by Router 2. Further assume that $\lambda 2$ carries no traffic that needs to be delivered to Router 2; all $\lambda 2$ traffic is destined for Router 3. In typical, point-to-point WDM networks, all three wavelengths are delivered to Router 2, as shown in Fig. 1a; thus, $\lambda 2$ is needlessly processed by Router 2. However, by implementing *optical bypass*, $\lambda 2$ completely bypasses Router 2, as shown in Fig. 1b. In this figure, a cross-connect with a full-wavelength granularity directs $\lambda 1$ and $\lambda 3$ to Router 2, but directly delivers $\lambda 2$ to Router 3, thus eliminating any need for Router 2 to process this traffic. (It is possible to eliminate the cross-connect and permanently have $\lambda 2$ optically bypass Router 2, however, the cross-connect allows the network topology to adjust to traffic demands. Also, for a bus or ring network, with only one input and output direction, a Wavelength Add/Drop Multiplexer can replace the cross-connect.)

The overall result of the optical bypass approach is that smaller routers are required, at the expense of adding cross-connects that operate on a much coarser granularity than a router and are thus more efficient at handling large amounts of traffic. In fact, the optical bypass approach performs more efficiently as network traffic increases, since it becomes more likely that full wavelengths worth of traffic will be sent intact over a long distance. Furthermore, as the traffic continues to increase, cross-connects that operate on coarser granularities (e.g., wavelength bands or entire fibers) are feasible, so that the overall port count, of both the router and the cross-connect, remains tractable.

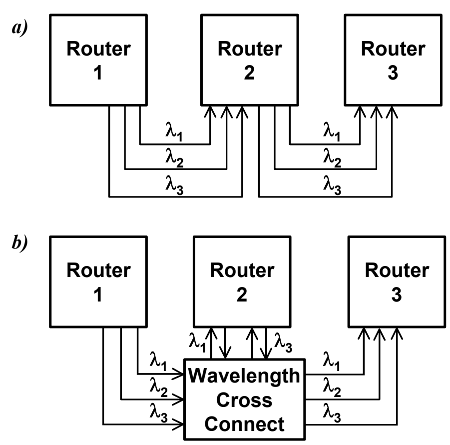

Figure 1. a) In the top figure, all three wavelengths are delivered to Router 2, even though λ_2 carries no traffic that needs to be processed by Router 2. b) A cross-connect (or wavelength add/drop multiplexer) operating at the granularity of a wavelength is used to drop or pass wavelengths at Router 2 as needed. In this figure, λ_2 optically bypasses the router, thus reducing the required capacity of the router.

In this paper, we explore the value that can be provided by optical bypass by measuring the fractional decrease in required router or switch size one can expect. For concreteness, we refer to routers, and consider cross-connects that operate on the granularity of a wavelength; however, the results are very general. We consider two network topologies: ring and mesh; and two traffic patterns: uniform all-to-all traffic and distance-dependent all-to-all traffic. We show that the reduction in router size can be very significant, with the fractional savings increasing as the size of the network increases and the amount of traffic increases. Also, uniform traffic yields greater savings than distance-dependent traffic, for the same amount of traffic.

Related work on Add/Drop Multiplexer savings achievable through optical bypass in rings was presented in [4-7]; these results are used below to analyze router-size savings. Optical bypass in access rings, where the traffic is typically directed to a small number of hubs, was previously analyzed in [5, 8], and is not included here.

II. Router Size

When considering router size in this paper, we in general are referring to the capacity required to handle the sum of the router traffic. To better illustrate this, refer to the router shown in Fig. 2. Here we assume the router receives a total of T_I traffic units from other routers. Of this traffic, T_L is destined for the local subnetworks attached to the router; we refer to this traffic as *Local Traffic*. We assume traffic is symmetric, such that there are T_L local incoming traffic units and T_O traffic units sent out to other routers, where $T_O = T_I$. The required router capacity is then:

$$C = T_I + T_L = T_O + T_L \tag{1}$$

If optical bypass is implemented, there may also be traffic that completely bypasses the router, as represented by T_T in Fig. 2. We refer to this as *Through Traffic*. Without optical bypass, T_T would have to be processed by the router. Thus, the fractional decrease in router-size provided by optical bypass is:

$$\frac{T_T}{T_T + C} \tag{2}$$

This *Through-to-Total Ratio* was similarly defined in [9] as a means of measuring the benefit of optical bypass.

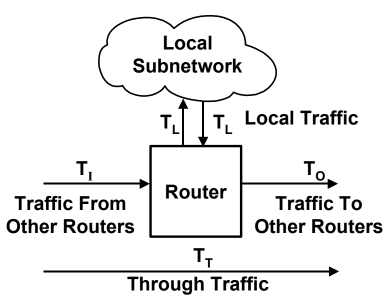

Figure 2. The required capacity of the router is determined by the traffic it receives from other routers and the traffic that is delivered to its local subnetwork. With optical bypass, some of the traffic from other routers may completely bypass the router and thus not contribute to the required capacity.

Our general strategy is to first calculate the through-to-total ratio achievable through optical bypass for various topologies and traffic patterns, assuming the traffic demand between any two nodes is an *integral* number of wavelengths. This ratio is precisely the router-size savings one can achieve. (We assume there is one router per node.) The next step is to consider more moderate traffic, where the internodal demands are less than a full wavelength. In this case, it is necessary to bundle the sub-wavelength internodal demands together in order to form full wavelengths. For example, in Fig. 1, Router 1 may send one half-wavelength worth of traffic to both Routers 2 and 3; both connections can be multiplexed onto λ1. The goal is to bundle the traffic in such a way as to maximize the Through Traffic. By judiciously applying the "super-node approximation" technique of [9], we attain approximate router-size savings for the sub-wavelength internodal problem.

III. Ring Topology

First, we consider implementing optical bypass in a ring topology with bi-directional routing; we assume traffic is always routed over the shortest distance. We assume the number of nodes in the ring, represented by N, is odd, so that shortest distance is unambiguous; the results are similar for N even.

A. Uniform Traffic

We first consider uniform all-to-all traffic, where every node sends precisely 1/G of a wavelength to every other node, for some $G \geq 1$. For full-wavelength internodal traffic, i.e., where G equals 1, T_I is N-1 wavelengths, all of which is delivered locally to the router (i.e., T_I equals T_L), and $T_I + T_T$ is $\frac{N^2-1}{4}$ wavelengths [10]. Thus, T_T is $\frac{(N-1)(N-3)}{4}$. Using (2), the through-to-total ratio of each router is:

$$\frac{N-3}{N+5} \tag{3}$$

This ratio is precisely the fractional savings in router size that can be attained through the use of optical bypass with full-wavelength internodal traffic. The savings increases with N since the relative amount of through traffic expands as the ring is enlarged; for N equal to 10, the fractional savings is close to 50%.

Next, we consider more moderate traffic, were $G > 1$. Using the technique of [9], we partition the N nodes into groups of $\sqrt{G}$ nodes and form 'super-nodes' from each of these groupings, such that we have $N/\sqrt{G}$ super-nodes, as illustrated in Fig. 3 for N=20 and G=16. Consider sending 1/G of a wavelength worth of traffic from each of the $\sqrt{G}$ nodes comprising one super-node to each of the $\sqrt{G}$ nodes comprising any other super-node. In total, then, there are G connections, with the aggregate traffic sent from one super-node to the other being exactly a full wavelength.

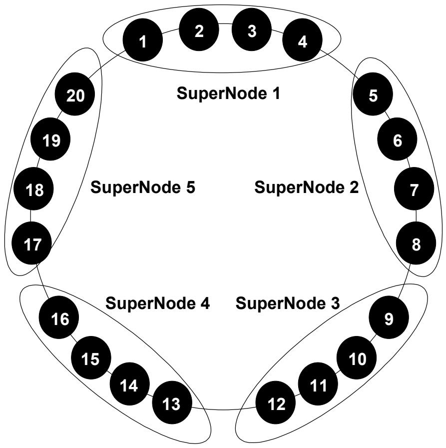

Figure 3. The super-node approximation, for 20 nodes and granularity 1/16 of a wavelength. The 20 nodes are partitioned into 5 super-nodes, each with 4 nodes. Each node within a super-node sends 1/16 of a wavelength worth of traffic to each node in every other super-node; thus, a total of one full wavelength is sent from each super-node to every other super-node.

Replacing N by $N/\sqrt{G}$ in the formulas for full-wavelength traffic yields: $T_I = \frac{N}{\sqrt{G}} - 1$ and $T_T = \frac{1}{4}\left(\frac{N}{\sqrt{G}} - 1\right)\left(\frac{N}{\sqrt{G}} - 3\right)$. Only $1/\sqrt{G}$ of the traffic delivered to each super-node is delivered locally to any single node; thus, $T_L = T_I/\sqrt{G}$. Using (2), the through-to-total ratio of each router is:

$$\frac{N - 3\sqrt{G}}{N + \sqrt{G} + 4} \quad (4)$$

Note that the super-node construction does not account for traffic between the nodes comprising a super-node, thus we expect this approximation to slightly overestimate the possible savings. Also, the quantities involved (e.g., $\sqrt{G}$) do not always represent integers. Nevertheless, it can be shown, using the enumerative methodology of [6], that the optimal fractional savings are extremely close to that of (4).

It is also possible to use the nodes at the extreme edge of a super-node (e.g., Nodes 1 and 4 in Fig. 3) to groom traffic such that more traffic bypasses the interior nodes of the super-node (e.g., Nodes 2 and 3). For example, the extreme nodes could use individual wavelengths to send to the interior nodes only the traffic that is local to them, possibly at the expense of using more wavelengths. The fractional router-size savings at the interior nodes would then be larger.

B. *Distance-Dependent Traffic*

In real networks, traffic demand has historically been correlated with distance, with nodes that are closer in distance exchanging more traffic. In this section, we incorporate this feature using the distance-dependent traffic demand curve represented by the solid line in Fig. 4, which is chosen because of its analytic simplicity and reasonable depiction of reality.

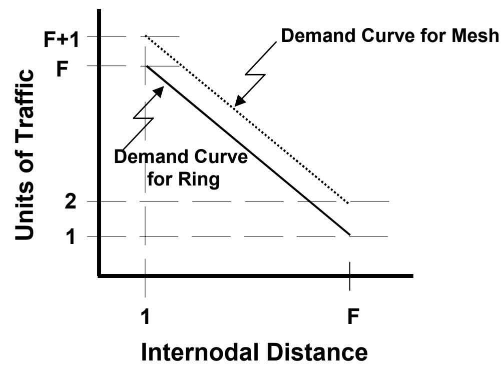

Figure 4. Distance-dependent traffic demand patterns used in the ring and mesh analyses, where F represents the distance of the nodes that are furthest apart. In the ring topology demand curve, the nodes furthest apart exchange one unit of traffic, and the internodal traffic demand increases by one unit as the internodal distance decreases by one link. The mesh demand curve is shifted up by one unit to simplify the analysis, as described in the text.

The amount of traffic between the most distant nodes is one unit, and the traffic demand increases by one unit as the distances decrease by one link. We consider bi-directional rings, so that the shorter path between two nodes is used as the internodal distance.

For the case of full-wavelength internodal traffic (G = 1), it can be shown that T_I is $\frac{N^2-1}{4}$ wavelengths, all of which is delivered locally to the router (i.e., T_I equals T_L), and T_T is $\frac{(N^2-1)(N-3)}{24}$. Using (2), the through-to-total ratio, and hence the fractional router-size savings, is:

$$\frac{N-3}{N+9} \quad (5)$$

For the case of G > 1, we again use a super-node approximation. If a super-node is comprised of X nodes, then using the ring internodal demand pattern of Fig. 4 yields a total of X^3 traffic units sent from any super-node to the super-node most distant from it (for any N). As the super-nodes get one unit closer, there is an additional X^3 traffic units sent. Thus, in Fig. 3, where X equals 4, the total amount of traffic sent from super-node 1 to 3 is 64 units (i.e., 4^3) and from super-node 1 to 2 is 128 units, where each unit is 1/G of a wavelength.

The super-node size should be chosen such that each super-node sends an integral number of wavelengths to every other super-node. Thus, for 1/G of a wavelength worth of internodal traffic, the super-node size, X, is chosen to be $G^{1/3}$, yielding an inter-super-node traffic pattern that precisely mimics the solid line in Fig. 4 (with the x-axis representing super-node distance). Using the full-wavelength results yields:

$T_I = \frac{1}{4}\left[\left(\frac{N}{G^{1/3}}\right)^2 - 1\right]$ and $T_T = \frac{1}{24}\left(\left(\frac{N}{G^{1/3}}\right)^2 - 1\right)\left(\frac{N}{G^{1/3}} - 3\right)$.

Only $1/G^{1/3}$ of the traffic delivered to each super-node is delivered locally to any single node; thus, $T_L = T_I/G^{1/3}$. Using (2), the through-to-total ratio, or the fractional size saving, of each router is:

$$\frac{N - 3G^{1/3}}{N + 3G^{1/3} + 6} \tag{6}$$

The router-size savings predicted by (6) somewhat underestimates the savings that were obtained in the best manual constructions. The difference is larger when N is large and G is small (excepting G =1, which is exact). The savings underestimate occurs because fewer than all $G^{1/3}$ nodes in each super-node may need to process a particular inter-super-node wavelength. For example, assume G equals 8, such that the unit of traffic is 1/8 of a wavelength and the super-node size is two. As the inter-super-node distance grows, a point is reached (assuming N is large enough) where one node sends 8 units of traffic (i.e., *a full wavelength*) to a node in another super-node. This particular wavelength only needs to be processed by one node, not both nodes, of each super-node. As N increases and G decreases, there is a greater occurrence of inter-super-node wavelengths that need to be processed by only a fraction of the nodes comprising the super-nodes. Thus, (6) underestimates the attainable router-size savings. Nevertheless, for simplicity, we use (6) as a reasonable approximation to the savings.

In addition, as discussed for the uniform traffic scenario, grooming could be performed by the extreme nodes of the super-node so that the interior nodes realize a greater savings.

C. *Uniform vs. Distance-Dependent Ring Traffic*

In order to compare the savings that can be obtained using optical bypass with uniform traffic as opposed to distance-dependent traffic, we normalize the traffic such that there is an equal amount of traffic sourced at each node. For 1/G uniform internodal traffic, the total amount of traffic sourced at each node is $\frac{N-1}{G}$; for 1/G distance-dependent traffic, the total is $\frac{N-1}{2}\frac{N+1}{2}\frac{1}{G}$. Thus, for comparison purposes, we use 1/G uniform traffic and 1/G′ distance-dependent traffic, where G′ equals $\frac{N+1}{4}G$.

The comparison between the two traffic demand patterns is shown in Fig. 5 for a range of ring sizes and for G equal to 1, 4, and 16. The savings with uniform traffic is greater than or equal to that of distance-dependent traffic in all cases. This is expected because distance-dependent traffic favors connections that are close, resulting in less through traffic. The differences are smaller for finer granularity traffic (i.e., larger G); more connections need to be bundled together to form a full wavelength so that it is less likely a wavelength can bypass a node in either traffic demand scenario.

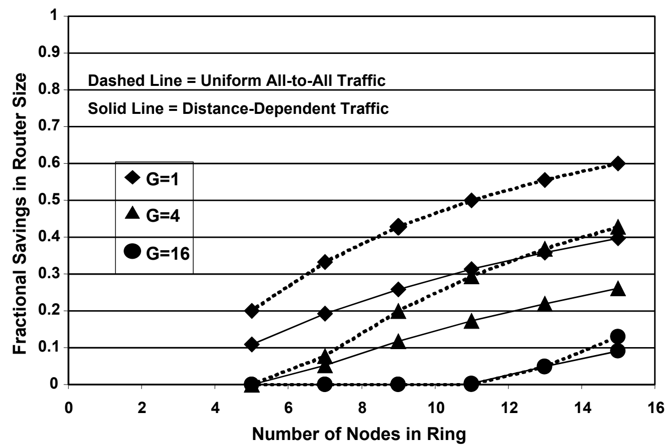

Figure 5. Fractional router-size savings with uniform traffic vs. distance-dependent traffic on a ring. The internodal demand for the uniform traffic scenario is 1/G of a wavelength, for G = 1, 4, 16. The internodal demand for the distance-dependent traffic is adjusted accordingly so that the amount of traffic sourced/sunk at each node equals that of the uniform traffic scenario.

IV. MESH TOPOLOGY

We next consider optical bypass in a mesh topology, where N nodes are arranged in a $\sqrt{N} \times \sqrt{N}$ grid, and each node is directly connected to four neighboring nodes (except for the edge nodes), as shown in the 6x6 example of Fig. 6. The source-to-sink routing strategy that we follow is to route by row and then by column. As with rings, we consider both uniform and distance-dependent traffic. The fractional router-size savings depends on the position of the node in the mesh – the savings are greater for the nodes nearer the center of the mesh since they are passed by more through traffic than the nodes near the edge. The calculations below are for the node at the very center of the mesh, for N odd (the results are similar for N even). However, the fractional router savings drop off very slowly as one moves away from the center so that the formulas approximately hold for most of the nodes in the mesh.

A. *Uniform Traffic*

We first consider uniform full-wavelength traffic, where every node in the mesh sends precisely one wavelength worth of traffic to every other node. It can be shown that, for the center node of the mesh, T_I is N-1 wavelengths, all of which is delivered locally to the router (i.e., T_I equals T_L), and T_T is $(N-1)(\sqrt{N}-1)$. Using (2), the through-to-total ratio is:

$$\frac{\sqrt{N}-1}{\sqrt{N}+1} \tag{7}$$

Equation 7 represents the fractional router-size savings that can be achieved, for full wavelength traffic, at the center node by taking advantage of optical bypass. The savings increase with the size of the mesh.

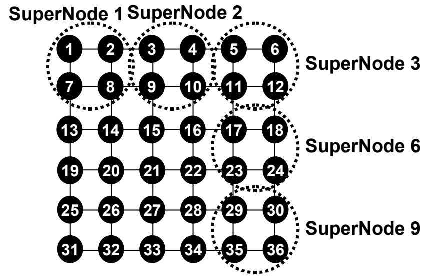

Figure 6. A 6x6 square mesh topology. For 1/16 of a wavelength uniform inter-nodal traffic, the 36 nodes are partitioned into 9 super-nodes, each with 4 nodes. One full wavelength is sent from each super-node to every other super-node.

The super-node approximation can be applied to the mesh for sub-wavelength traffic. Nodes are grouped into square-shaped super-nodes as shown in Fig. 6. For 1/G of a wavelength worth of internodal traffic, the super-nodes are comprised of $\sqrt{G}$ nodes, which yields one full wavelength worth of traffic between each super-node pair (same bundling size as the ring topology). Applying the full-wavelength results to the center super-node yields:

$$T_I = \frac{N}{\sqrt{G}} - 1 \text{ and } T_T = \left(\frac{N}{\sqrt{G}} - 1\right)\left(\sqrt{\frac{N}{\sqrt{G}}} - 1\right).$$

If we assume there is no grooming within the super-node, then each node in the center super-node receives T_I amount of traffic, of which a fraction $1/\sqrt{G}$ is local. The through traffic for each node in the center super-node is approximately $T_T/G^{1/4}$ (the super-node through traffic does not pass by all nodes in the super-node; on average, a through wavelength passes through $G^{1/4}$ nodes out of the $\sqrt{G}$ nodes in the super-node). Using (2), the through-to-total ratio of the nodes in the center super-node is approximately:

$$\frac{\sqrt{N} - G^{1/4}}{\sqrt{N} - G^{1/4} + \sqrt{G} + 1} \tag{8}$$

If grooming is performed, which is very likely to be the case in a real network, the router size of all the nodes in the super-node is decreased, and greater savings is achieved.

B. *Distance-Dependent Traffic*

Next, we consider distance-dependent traffic on a mesh. The distance between two mesh nodes is defined as the sum of the horizontal distance and the vertical distance (i.e., the number of links that must be crossed). The traffic demand pattern that we consider is represented by the dashed line in Fig. 4; compared to the demand pattern used for the ring topology, the amount of traffic between the most distant nodes is 2 units as opposed to 1. This adjustment allows for the application of the super-node approximation as will be discussed below.

With this traffic demand pattern, and full-wavelength internodal traffic, it can be shown that for the center node of the mesh, T_I is $\frac{3}{2}\sqrt{N}(N-1)$ (T_I equals T_L), and T_T is $\frac{1}{4}(N-1)(\sqrt{N}-1)(5\sqrt{N}-1)$. Using (2), the through-to-total ratio at the center node is:

$$\frac{(\sqrt{N}-1)(5\sqrt{N}-1)}{(\sqrt{N}+1)(5\sqrt{N}+1)} \approx 1 - \frac{2}{\sqrt{N}} \tag{9}$$

Next, consider applying the super-node approximation for sub-wavelength internodal demand ($G > 1$). Assume that the super-node bundle size is X, i.e., an $X^{1/2} \times X^{1/2}$ square. Using the mesh internodal demand curve of Fig. 4, the number of traffic units between the most distant super-nodes is $2X^{5/2}$. The traffic increases by $X^{5/2}$ as the super-nodes get one unit closer; the increase is slightly less when moving to a super-node that lies in the same row or column. To illustrate this, we consider Fig. 6, where X=4.

Traffic Units Between Super-nodes 1 and 9 = $2X^{5/2}$ = 64
Traffic Units Between Super-nodes 1 and 6 = $3X^{5/2}$ = 96
Traffic Units Between Super-nodes 1 and 3 = $3.75X^{5/2}$
= 120 (1 and 3 are in same row)

We round up the traffic between super-nodes in the same row and column to the nearest multiple of $X^{5/2}$ (in essence, this assumes there is *slightly* more traffic between some super-nodes than is actually present). (If the solid line in Fig. 4 had been used as the internodal demand curve, the amount of traffic between the most distant nodes would be $2X^{5/2} - X^2$; the traffic would still increase by $X^{5/2}$ as the super-nodes get closer.)

For internodal traffic granularity of 1/G, we choose X to be $G^{2/5}$, to yield an inter-super-node demand pattern that exactly mimics the dashed line in Fig. 4. Applying the full-wavelength results to the center super-node yields:

$$T_I = \frac{3}{2}\sqrt{\frac{N}{G^{2/5}}}\left(\frac{N}{G^{2/5}} - 1\right) \text{ and}$$

$$T_T = \frac{1}{4}\left(\frac{N}{G^{2/5}} - 1\right)\left(\sqrt{\frac{N}{G^{2/5}}} - 1\right)\left(5\sqrt{\frac{N}{G^{2/5}}} - 1\right).$$

If we assume there is no grooming within the super-node, then each node in the center super-node receives T_I amount of traffic, of which a fraction $1/G^{2/5}$ is local. The through traffic at any node in the super-node is approximately $T_T/\sqrt{G^{2/5}}$ (same reasoning as for uniform traffic). The fractional router-size savings for 1/G internodal distance-dependent traffic is then approximately:

$$\frac{(\sqrt{N} - \sqrt{G^{2/5}})(5\sqrt{N} - \sqrt{G^{2/5}})}{(\sqrt{N} - \sqrt{G^{2/5}})(5\sqrt{N} - \sqrt{G^{2/5}}) + 6\sqrt{N}(G^{2/5} + 1)} \tag{10}$$

As discussed for the distance-dependent ring scenario, when N is large and G is small (excepting G=1), (10) somewhat underestimates the router-size savings due to some inter-super-node wavelengths not needing to be processed by all nodes comprising the super-node.

Additionally, if grooming were performed at the nodes, the savings would be even greater than indicated by (10).

TABLE 1. SUMMARY OF RESULTS

	N-Node Ring Topology		N-Node Mesh Topology	
	Uniform	**Distance-Dependent**	**Uniform**	**Distance-Dependent**
Fractional Router-Size Savings for Integer Number of Wavelengths Traffic	$\frac{N-3}{N+5}$	$\frac{N-3}{N+9}$	$\frac{\sqrt{N}-1}{\sqrt{N}+1}$	$1-\frac{2}{\sqrt{N}}$
Fractional Router-Size Savings for 1/G Granularity Traffic	$\frac{N-3\sqrt{G}}{N+\sqrt{G}+4}$	$\frac{N-3G^{1/3}}{N+3G^{1/3}+6}$	$\frac{\sqrt{N}-G^{1/4}}{\sqrt{N}-G^{1/4}+\sqrt{G}+1}$	Equation (10)
Total Amount of Traffic Sourced for 1/G Granularity Traffic	$\frac{N(N-1)}{G}$	$\frac{N(N-1)}{2}\frac{(N+1)}{2}\frac{1}{G}$	$\frac{N(N-1)}{G}$	$\frac{4N(N-1)\sqrt{N}}{3}\frac{1}{G}$

C. *Uniform vs. Distance-Dependent Mesh Traffic*

We compare the savings that can be obtained with uniform traffic as opposed to distance-dependent traffic on the mesh topology. For 1/G uniform internodal traffic, the total amount of traffic sourced at *all* nodes is $\frac{N(N-1)}{G}$; for 1/G distance-dependent traffic, the total amount of traffic sourced at *all* nodes is $\frac{4N(N-1)\sqrt{N}}{3}\frac{1}{G}$. Thus, in order to equate the total sourced traffic, we use 1/G uniform traffic and $1/G'$ distance-dependent traffic, where G' equals $\frac{4\sqrt{N}}{3}G$.

The comparison of savings between the two traffic demand patterns is shown in Fig. 7 for a range of mesh sizes and for G equal to 1, 4, and 16. As with rings, the router-size savings with uniform traffic is greater than or equal to that of distance-dependent traffic in all cases. However, for either traffic demand scenario, the savings can be quite significant, particularly for a large mesh.

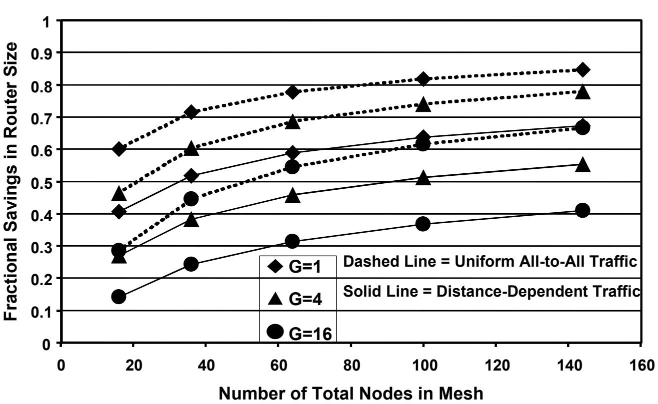

Figure 7. Fractional router-size savings with uniform traffic vs. distance-dependent traffic on a mesh. The internodal demand for the uniform traffic scenario is 1/G of a wavelength, for G = 1, 4, 16. The internodal demand for the distance-dependent traffic is adjusted so that the total amount of traffic sourced at all the nodes is the same in the two scenarios. The savings is greater with uniform traffic, however, the savings can be significant in either scenario.

V. SUMMARY

We have analyzed the benefit that can be provided by employing optical bypass in gigabit networks. The formulas corresponding to the various topologies and traffic demand patterns are summarized in Table 1. The fractional router-size savings for integer number of wavelengths traffic (i.e., $G = 1$) are shown in the first row. The approximate fractional router-size savings for $G > 1$, obtained using the super-node approximation technique, are shown in the second row. (Note that setting G equal to 1 in the formulas of the second row yields the exact results of the first row.)

We have shown that optical bypass can result in significantly smaller routers. Furthermore, the relative savings increase as the size of the network and the internodal traffic increases, indicating that this technique will scale well in future networks.

REFERENCES

[1] McKeown, N., Izzard, M., Mekkittikul, A., Ellersick, W., and Horowitz, M., “The tiny Tera: a packet switch core,” *IEEE Micro*, Jan/Feb 1997, pp. 26-33.

[2] Rosen, E.C., Viswanathan, A., and Callon, R., “Multiprotocol Label Switching architecture,” IETF, Internet Draft, July, 1998.

[3] Viswanathan, A., Feldman, N., Wang, Z., and Callon, R., “Evolution of Multiprotocol Label Switching,” *IEEE Communications Magazine*, May, 1998, pp. 165-173.

[4] Simmons, J. M., Goldstein, E.L., Saleh, A.A.M., “On the value of wavelength-add/drop in WDM rings with uniform traffic,” *OFC’98*, San Jose, CA, February 22-February 27, 1998, pp. 361-362.

[5] Ramaswami, R. And Sivarajan, K., *Optical Networks: A Practical Perspective,* Morgan Kaufmann Publishers, San Francisco, CA, 1998, pp. 359-389.

[6] Simmons, J. M., Goldstein, E.L., Saleh, A.A.M., “Quantifying the benefit of wavelength-add/drop in WDM rings with distance-independent and dependent traffic”, *IEEE/OSA Journal of Lightwave Technology,* Vol. 17, No. 1, January, 1999, pp. 48-57.

[7] Modiano, E. and Chiu, A., “Reducing electronic multiplexing costs in unidirectional SONET/WDM ring networks via efficient traffic grooming,” *Globecom*, Sydney, Australia, Nov. 8-12, 1998.

[8] Simmons, J.M. and Saleh, A.A.M., “On the value of wavelength-add/drop in WDM rings,” *Trends in Optics and Photonics: Optical Networks and Their Applications*, Vol. 20, June, 1998, pp. 179-183.

[9] Saleh, A.A.M., “Transparent optical networks for the Next Generation Information Infrastructure,” *OFC’95*, San Diego, CA, February 26-March 3, 1995, p. 241.

[10] Elrefaie, A. F., “Multiwavelength survivable ring network architectures,” *ICC’93*, Geneva, Switzerland, Paper 48.7, May 23-26, 1993, pp. 1245-1251.