MECHANISMS FOR PROVIDING OPTICAL BYPASS IN WDM-BASED NETWORKS

EYTAN MODIANO
ARADHANA NARULA-TAM
MIT Lincoln Laboratory
Lexington, MA 02420

ABSTRACT

The recent emergence of Wavelength Division Multiplexing (WDM) technology has led to a tremendous increase in the available transmission capacity in wide-area networks. Consequently, these networks may no longer be limited by transmission bandwidth, but rather by the processing capability of electronic switches, routers and multiplexers in the network. This realization has led to a new wave of research aimed at overcoming the electronic bottleneck by providing optical bypass at the WDM layer.

In this paper we survey some of the recent work in this area, including mechanisms for providing optical bypass in both circuit and packet switched networks. We pay particular attention to techniques that take advantage of configurable WDM technology to dynamically reconfigure the electronic topology of the network in response to changes in traffic conditions. We show that for both circuit and packet switched traffic, having a configurable WDM topology can lead to a significant increase in network capacity.

1 INTRODUCTION

Wavelength Division Multiplexing (WDM) is emerging as a dominant technology for use in the backbone network. With WDM, the capacity of a fiber is significantly increased by allowing simultaneous transmission on multiple wavelengths (channels), each operating at the maximum electronic rate. Systems with between 40 and 80 wavelengths are presently being deployed for point-to-point transmission. With tens of wavelengths per fiber and transmission rates of up to 10 Gbps per wavelength, capacities that approach a terabit per second can be achieved. While such enormous capacity is very exciting, it also places a tremendous burden on the electronic switches and routers at each node that must somehow process all of this information. Fortunately, it is not necessary to electronically process all traffic entering and leaving each node. For example, much of the traffic passing through a node is neither sourced at that node nor destined to that node. To reduce the amount of traffic that must be electronically processed at intermediate nodes, future WDM systems will employ WDM Add/Drop multiplexers (WADMs) that allow each wavelength to either be dropped and electronically processed at the node or to optically bypass the node electronics, as shown in Figure 1.

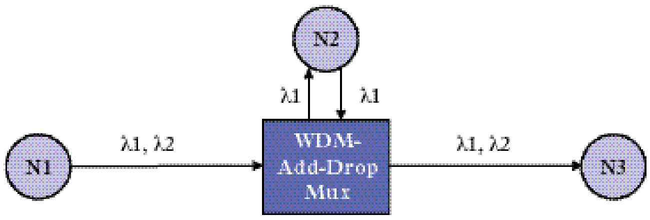

FIGURE 1: *Using a WADM to provide optical bypass.*

This capability of optically bypassing a node's electronic layer on a wavelength by wavelength basis can be exploited to reduce overall electronic processing using a variety of methods, each of which is appropriate for different types of traffic scenarios. For the case of low rate static circuit traffic, traffic grooming may be utilized to selectively multiplex multiple low rate circuits on wavelengths such that the number of wavelengths that must be processed at each node is minimized. For dynamically switched high rate circuits, the logical topology may be reconfigured to avoid electronic processing at all intermediate nodes. Similarly, logical topology reconfiguration may be utilized to reduce the load on the electronic routers in a packet-switched network via dynamic load balancing.Lastly, for large volume transmissions, it is possible to bypass all network electronics and provide an all-optical connection through optical flow switching. In this article we present an overview of these techniques for reducing electronic layer processing through optical bypass.

Much of today's network infrastructure is in the form of "static" SONET rings. A typical SONET ring multiplexes a number of low rate circuits onto a high rate SONET ring. For example, an OC-48 (2.5 Gbps) SONET ring can be used to multiplex 16 OC-3 circuits. At each node on the ring a SONET Add/Drop Multiplexer (SADM) is used to multiplex the low rate circuits onto the ring. With WDM, multiple SONET rings can be supported on a single fiber, requiring a large number of SADMs to be present at each node.However, if circuits can be assigned to wavelengths in such a way that some wavelengths do not carry any traffic to a node,then that node can be bypassed without needing a SADM on that wavelength. In section 2 of this paper we discuss this static circuit assignment problem and show that a significant savings in SADMs can be achieved through the use of clever circuit assignment algorithms.

While much of today's circuit traffic is statically allocated, in the future there will be an increased demand for dynamically assigned (switched) high rate circuits. For high rate circuit switched traffic, where each circuit requires on the order of a full wavelength, configurable WADMs may be used to selectively add/drop only wavelengths carrying traffic that is either sourced or destined to a particular node. In this way, the configurable WDM optical layer is used to dynamically reconfigure the "logical" (i.e., electronic) topology of the network in response to new call arrivals. This can lead to a significant increase in the capacity that the network can support.With a fixed WDM layer, calls may be blocked due to a lack of electronic resources at intermediate nodes between the source and the destination. However, a configurable optical layer can be used to dynamically bypass the electronic layer on heavily loaded intermediate nodes and hence increase the traffic load that can be supported in the network.The potential gains in circuit capacity achievable with a configurable WDM physical layer are discussed in section 3.

The configurability of the WDM layer can also be used to reduce electronic processing bottlenecks in packet-switched networks (e.g., IP or ATM networks). In a packet-switched network, traffic between a source and destination node typically goes through a number of intermediate electronic switches and routers. As the traffic in the network changes,so does the load on the routers. With a configurable WDM topology, it is possible to change the way in which the routers are connected in order to balance the load on the lightpaths and electronic routers. Since network delay is generally dominated by the maximally loaded link/router, reducing the load on this link can significantly improve network delay characteristics. This reconfiguration of the logical topology can be executed in response to changes in traffic conditions. In section 4 we discuss load balancing algorithms and the potential reduction in network load.

Finally, we consider the case of large traffic flows resulting from extremely large transmissions such as multi-gigabit file transfers, high bit rate streams, etc. Transactions of this size have the potential of overwhelming the network's electronic layer and causing delays for all network traffic. To prevent such network congestion,Optical Flow Switching (OFS) may be employed whereby an all-optical (WDM) connection is established between end-users and all of the electronics in the network are bypassed. In section 5 we discuss issues and algorithms for OFS.

2 STATIC CIRCUIT ASSIGNMENT

Much of today's physical layer network infrastructure is built around Synchronous Optical Network (SONET) rings. Typically, a SONET ring is constructed using fiber (one or two fiber pairs are typically used in order to provide protection) to connect SONET Add/Drop Multiplexers (SADMs). Each SONET ADM has the ability to aggregate lower rate SONET signals into a single high rate SONET stream. For example, four OC-3 circuits can be multiplexed together into an OC-12 circuit and 16 OC-3's can be multiplexed into an OC-48. The recent emergence of Wavelength Division Multiplexing (WDM) technology has provided the ability to support multiple SONET rings on a single fiber pair. Consider, for example, the SONET ring network shown in Figure 2, where each wavelength is used to form an OC-48 SONET ring. With WDM technology providing dozens of wavelengths on a fiber, dozens of OC-48 rings can be supported per fiber pair using Wavelength Add/Drop Multiplexers (WADMs) to separate the multiple SONET rings.

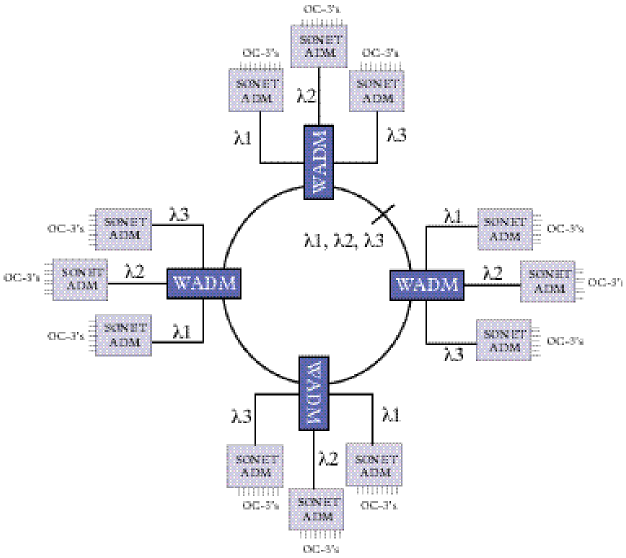

FIGURE 2: *SONET/WDM rings.*

This tremendous increase in network capacity, of course, comes at the expense of additional electronic multiplexing equipment. With the emergence of WDM technology, the dominant cost component in networks is no longer the cost of fiber but rather the cost of electronics.

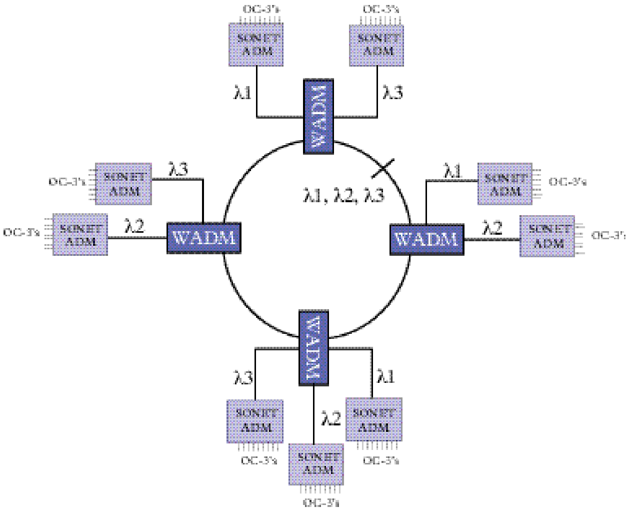

FIGURE 3: *Using WADMs to reduce the number of SONET ADMs.*

The SONET/WDM architecture shown in Figure 2 is potentially wasteful of SONET ADMs because every wavelength (ring) requires a SONET ADM at every node.As mentioned previously, not all traffic needs to be electronically processed at each node. A WADM at a given node is capable of dropping and adding any number of wavelengths at that node. In order for a node to transmit or receive traffic on a wavelength, the wavelength must be added or dropped at that node and a SONET ADM must be used.Therefore, it is no longer necessary to have a SONET ADM for every wavelength at every node, but rather only for those wavelengths that are used at that node. Therefore, in order to limit the number of SONET ADMs required, the traffic may be groomed in such a way that all of the traffic, to and from a node, is carried on the minimum number of wavelengths.

As a simple, illustrative example consider a ring network with four nodes. Suppose that each wavelength is used to support an OC-48 ring, and that the traffic requirement is for 8 OC-3 circuits between each pair of nodes. In this example we have 6 node pairs and the total traffic load is equal to 48 OC-3's or equivalently 3 OC-48 rings. The question is how to assign the traffic to these 3 OC-48 rings in a way that minimizes the total number of SONET ADMs required.Consider, for example, the following two circuit assignments of traffic:

Assignment #1	**Assignment #2**
λ1: 1⌧2, 3⌧4	λ1: 1⌧2, 1⌧3
λ2: 1⌧3, 2⌧4	λ2: 2⌧3, 2⌧4
λ3: 1⌧4, 2⌧3	λ3: 1⌧4, 3⌧4

With the first assignment, each node has some traffic on every wavelength. For example wavelength 1 carries the traffic between nodes 1 and 2 and the traffic between nodes 3 and 4. Therefore,each node would require an SADM on every wavelength for a total of 12 SADMs. With the second assignment each wavelength contains traffic from only 3 nodes and hence only 9 SADMs are needed. Notice that both assignments carry the same amount of total traffic (8 OC-3's between each pair of nodes).The corresponding SADM allocations for both assignments are shown in Figures 2 and 3, respectively.

Recently this problem of grooming traffic onto wavelengths in order to minimize the number of electronic multiplexers in the network has received much attention [2,3,4,25,26].While the general traffic grooming problem has been shown to be NP-complete [3], efficient grooming algorithms have been developed for special cases. For example,[2] considers traffic grooming for a bidirectional ring with uniform traffic and [3, 4] consider the same problem for a uni-directional ring. All show that significant savings in the number of SADMs can be achieved through efficient traffic grooming. For example, shown in Figure 4 is the number of SADMs required when using the traffic grooming algorithm developed in [4]. This number is compared to the number of SADMs required when no grooming is used (i.e., all wavelengths are dropped at all nodes).It is also compared to a lower bound on the number of SADMs. As can be seen from the figure, the algorithms developed in [4] are nearly optimal for the uniform traffic case, and achieve significant SADM savings.

In [5] the traffic grooming problem is generalized to encompass more general traffic models. The traffic is no longer restricted to be uniform and nodes are allowed to have dynamically changing connections. The model in [5] allows each node to source up to t circuits. These t circuits can be destined to any of the nodes in the network without restriction, and the destinations of the circuits can be dynamically changed. Algorithms are developed to minimize the number of wavelengths that must be processed at each node. These algorithms provide methods for achieving significant reductions in SADMs under a variety of traffic requirements.

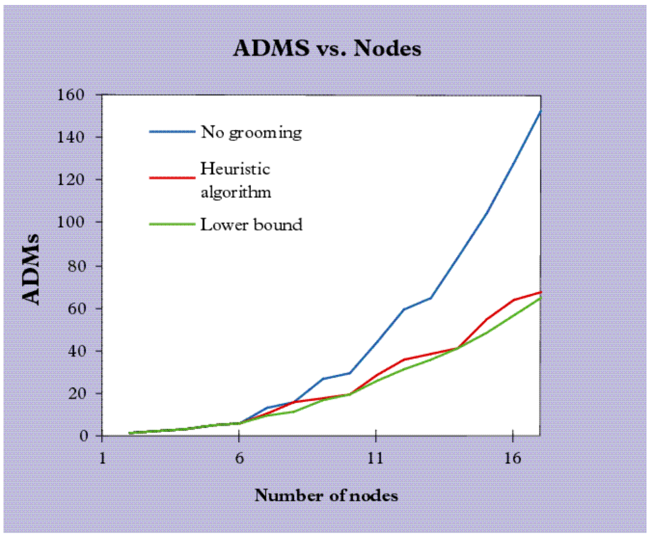

FIGURE 4: *Performance of grooming in a WDM/SONET ring network.*

3 DYNAMIC CIRCUIT SWITCHING

In WDM networks, the physical topology, consisting of passive or configurable optical nodes interconnected by fiber links, is the optical layer topology. The logical or virtual topology, seen by the electronic layer, consists of lightpath connections between nodes established by tuning the transmitter of one node and the receiver of another node to the same wavelength. In this way, WDM networks provide a way to interconnect electronic switches with high bandwidth bit pipes without dedicating a fiber pair between each pair of switches. Furthermore, the configurable nature of WDM also allows these high bandwidth pipes to be dynamically reconfigured in response to changes in traffic conditions.

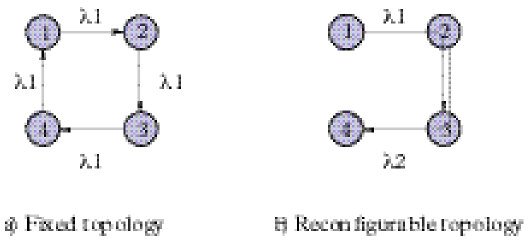

a) Fixed topology b) Reconfigurable topology

FIGURE 5: *Using WDM to reconfigure the electronic topology.*

Consider for example a network where each node is equipped with a single port (transmitter/receiver) and the fiber supports two wavelengths, l1 and l2. In a system with a fixed logical topology (also known as a switchless network), the transmitters and receivers are fixed tuned to their respective wavelengths at system inception. With configurable WADMs and tunable transmitters and receivers, the logical topology can be reconfigured to accommodate changing patterns in externally offered traffic. In Figure 5b, we show a four node example where the fixed topology is a clockwise uni-directional ring. If a full-wavelength call is in progress from node 1 to node 3, and a call request arrives between nodes 2 and 4, then that request must be blocked. In a reconfigurable system, both calls can be supported as shown in Figure 5. Reconfiguring the logical topology allows traffic to optically bypass intermediate node electronics, thereby minimizing the number of electronic ports necessary. In the example above, by reconfiguring the logical topology, the traffic from node 1 to node 3 optically bypasses node 2, thereby allowing the port at node 2 to carry traffic from node 2 to node 4. Similarly, the traffic from node 2 to node 4 optically bypasses node 3 whose port is busy with the connection from node 1 to node 3.

In a circuit switched network, reconfiguration increases the traffic load that the network can support for a given blocking probability. The benefits of reconfiguration can be quantified by comparing the blocking probabilities of fixed topology and configurable topology networks. Many researchers have studied blocking probabilities for WDM networks with and without wavelength changers [6,7,8,27,28]. This prior work, however, assumed that wavelengths were the most precious network resource. Therefore, in analyzing blocking probability previous researchers assumed that a call between a source and destination would be accepted if a wavelength was available between the source and destination, and blocked otherwise. This ignored the possibility of calls being blocked due to the lack of electronic resources. However, when considering a multi-hop circuit switched network, calls can be blocked even when wavelengths are available. A call may be blocked when ports on the source or destination nodes are occupied or when an intermediate node has no ports available, as demonstrated in the example. By reconfiguring the logical topology such that traffic optically bypasses intermediate nodes whenever possible, the number of ports required to sustain each call is minimized. In order to analyze the blocking probability in a system where both wavelengths and electronic ports are limited, a model for blocking probability that takes into account both restrictions is needed.

In [9] an approximate model for blocking probability in such a network was developed. The model developed in [9] assumes that calls arrive according to a Poisson process, utilize a full wavelength, and have exponential holding times. The model also assumes that, once placed, calls cannot be rearranged or rerouted. Prohibiting rerouting of existing calls eliminates the possibility of calls being adversely affected by reconfiguration. This model is used in Figures 6 and 7 to compare the performance of fixed topology and configurable topology ring networks with 10 and 100 nodes. Figures 6 and 7 show the reconfiguration capacity gain, defined as the ratio of the traffic load that can be supported by a reconfigurable ring network to that of a fixed topology bidirectional ring network, at a blocking probability of 0.01. These graphs illustrate that the capacity gain due to reconfiguration is most significant when

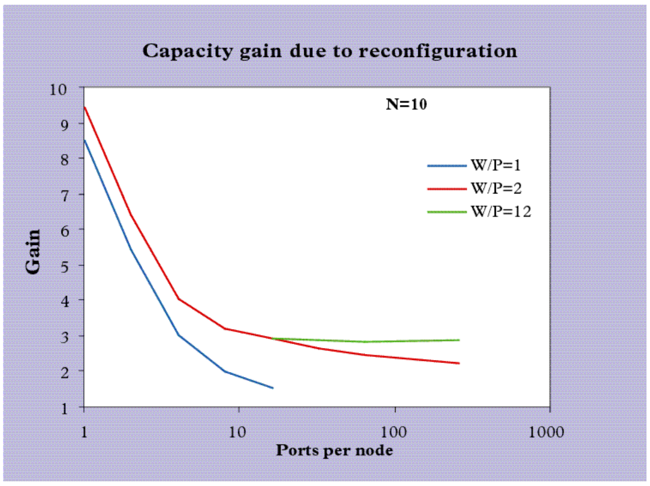

FIGURE 6: *Capacity gain due to reconfiguration in a 10 node bidirectional ring network.*

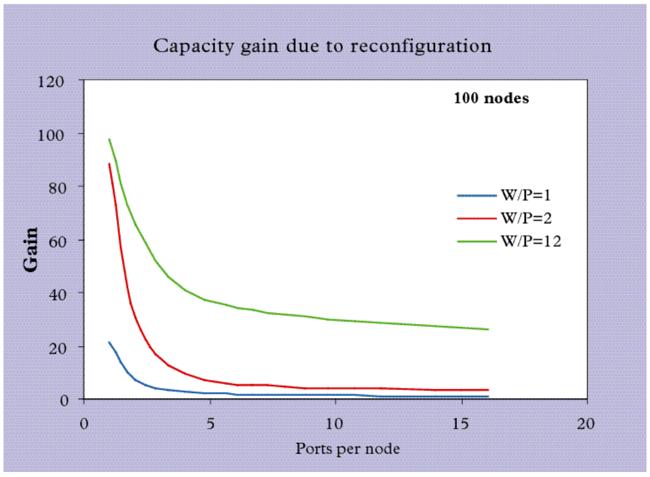

FIGURE 7: *Capacity gain due to reconfiguration in a 100 node bidirectional ring network.*

the ratio of wavelengths to ports per node (W/P) is large and the number of ports per node is small.

The results from [9] indicate that when the number of wavelengths is much larger than the number of available ports per node, a capacity gain on the order of N/2 can be obtained in a unidirectional ring, where N is the number of nodes in the network. In a bidirectional ring a factor of N/4 increase in capacity can be obtained. Intuitively, this can be explained by observing that when the number of wavelengths is much greater than the number of ports per node, the system is port limited. That is, calls are blocked only when there are no ports available to complete the call. With a fixed topology system configured as a unidirectional ring, each call takes an average of N/2 hops and hence N/2 ports (transmitters and receivers). With a configurable topology, each call only uses only one transmitter and one receiver. Therefore, a configurable system requires N/2 fewer transceivers per call, and can therefore support N/2 times the traffic of a fixed topology system. Similarly, in a bidirectional ring, each call takes an average of N/4 hops leading to an N/4 capacity gain when the topology is configurable. When the number of wavelengths approaches the number of ports per node, the benefits of reconfiguration are significantly diminished since both the configurable and fixed topology systems are wavelength limited, and many calls are blocked due to insufficient bandwidth availability. In fact, in a unidirectional ring, when the number of wavelengths is the same as the number of ports (W=P), there is no gain due to reconfiguration because all calls must be routed in the same direction and all of the wavelengths can be processed at all of the nodes. In a bidirectional ring, as shown in Figures 6 and 7, some gain is still obtained due to reconfiguration even when W=P. This is because in the bidirectional topology each fiber can support W wavelengths and so a number of different logical topologies can be realized.

4 PACKET-SWITCHED SYSTEMS

The concept of reconfiguring the virtual topology of the network in response to changes in traffic conditions can also be applied to packet-switched networks (e.g., IP networks) for the purpose of balancing the load on the electronic routers in the network. This is in contrast to the previous sections that considered circuit switched networks. Consider, for example, the network shown in Figure 8. Shown on the left (a) is the traffic matrix denoting the traffic requirement between the various nodes (i.e., one unit of traffic between node 1 and 4, etc.). When this matrix is routed on the clockwise ring shown in Figure 8b, the resulting load on each link of the ring is 3 units of traffic. Using configurable WADMs and tunable transmitters and receivers, the logical topology of the ring can be rearranged, as shown in (c), so that a counter-clockwise ring results and the load on each link is only one unit of traffic. This reduces the load on each of the electronic routers by a factor of 3.

Since network delays are generally limited by the link/router carrying the maximum load, it is beneficial to reconfigure the logical topology to minimize the maximally loaded link2. The problem of determining the optimal logical topology, however, is NP-complete [12]. Therefore, most prior work on topology design for WDM-based packet networks proposes heuristics for solving the optimal virtual topology design problem [15,16].

Once the new configuration is determined, the lightpaths from the old topology may be replaced by lightpaths for the new configuration. However, this process of reconfiguring to the new network topology can be disruptive to existing traffic. Traffic must be re-routed or buffered at each node while the transmitters and receivers are re-tuned. To minimize network disruption, methods of migrating to the optimal topology using a sequence of branch-exchanges (i.e., exchanging the source and destinations of two links in the network) have been developed [17]. Furthermore, it has been suggested that reconfiguration occur very infrequently, for example once or twice a day. Policies for reconfiguring when the benefits of reconfiguration outweigh the costs were examined in [14].

Reconfiguring the logical topology infrequently, however, may lead to logical topologies that are obsolete for much of the time they are in use. If traffic changes rapidly, it may not be possible, or desirable, to completely reconfigure the logical topology of the network with every perceived traffic change. Instead, a more gradual change in the topology may be desired. Migrating to the optimal topology using branch-exchange sequences as described in [15] minimizes the network disruption at each step. However, if traffic is rapidly changing, by the time that the new topology is implemented, the traffic may have changed. Furthermore, intermediate steps may leave the network temporarily disconnected.

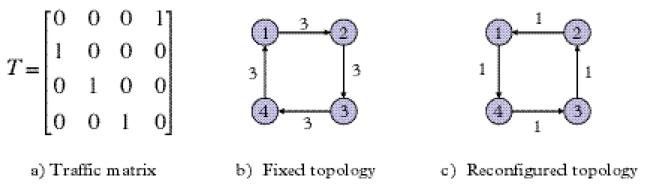

FIGURE 8: *Reconfiguring the logical topology of the network to reduce the traffic load.*

In [18] a reconfiguration strategy is developed that updates the logical topology via small changes at regular time intervals. The small changes in the topology limit the disruption to the network, allowing reconfiguration to be employed more regularly and consequently tracking the changes in traffic patterns. For a network in which each node is equipped with a single transmitter and receiver, the logical topology is iteratively improved using a three-branch exchange mechanism. A three-branch exchange selects three links in the ring, for example links a, b, and d, shown in figure 9, and reorders the connections so that the source of link a is connected to the destination of link b, the source of link b is connected to the destination of link d, and the source of link d is connected to the destination of link a. Therefore, each three-branch exchange maintains network connectivity while disrupting only three links in the network.

FIGURE 9: *Three branch exchange sequence.*

The algorithm in [18] performs a single three-branch exchange at regular time intervals (e.g., one per second). At each step, the algorithm searches for the three-branch exchange that maximally reduces the maximum link load. In order to evaluate the performance of the algorithm, two time-varying traffic models were utilized. In the first traffic model, the traffic between each node pair is independent and identically distributed (i.i.d) according to a uniform distribution. In the second traffic model, a number of clusters are generated (at random locations) where the traffic amongst nodes within a cluster is significantly greater than the traffic amongst non-cluster nodes. This clustered or hot spot traffic model represents the situation where a large amount of traffic may be flowing from a source node to several destination nodes (file server model) and also the scenario in which several source nodes are sending large amounts of traffic to a single destination. Furthermore,when the network nodes are aggregation points, clustered traffic is quite natural. Under clustered traffic, reconfiguration is expected to produce significant performance gains because the logical topology of the network can be tailored to the structured traffic pattern. This is in contrast to the i.i.d traffic model in which traffic between each pair of nodes is uncorrelated. To model the time-varying nature of traffic, the traffic evolves to a new random traffic matrix every D seconds.

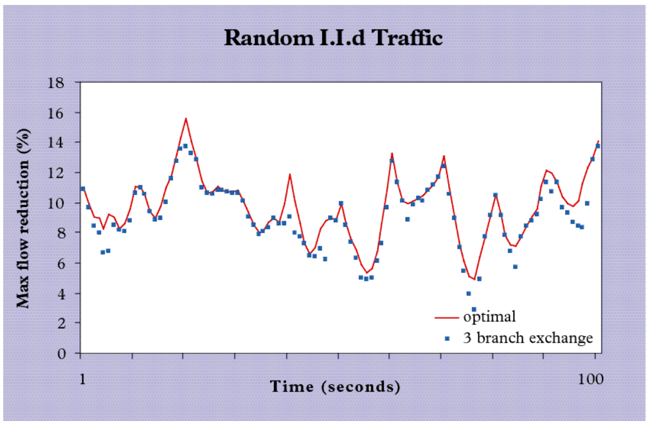

FIGURE 10: *Benefits of load balancing with random i.i.d traffic.*

Since the goal of reconfiguration is to reduce the maximum link load on the network, the performance of the algorithm can be measured in terms of the percent reduction in maximum load that it achieves over a fixed topology network. Figures 10 and 11 plot this reduction for i.i.d and clustered traffic, respectively, assuming that the traffic evolves to a new random traffic matrix every ten seconds and that a single three-branch exchange is implemented every second. This corresponds to ten reconfiguration steps between two independent random traffic realizations. The figures show the percent reduction in maximum link load as a function of time. Also, shown in the figures is the reduction in maximum load that can be achieved if the optimal topology is utilized at all times. As can be seen from the figures, the algorithm tracks the optimal topology very closely. However, it is also clear that the gain due to reconfiguration is rather modest when the traffic is i.i.d. In fact, for i.i.d traffic the average reduction in load using the algorithm is 10% (compared to a maximum gain of 11% when the optimal topology is used). However, when the traffic is clustered, an average reduction in load of 30% is achieved (compared to 31% optimal).

FIGURE 11: *Benefits of load balancing with clustered traffic.*

5 OPTICAL FLOW SWITCHING

All of the algorithms presented in this paper aim to reduce the electronic component requirements in the network by providing optical bypass. The ultimate form of optical bypass is Optical Flow Switching (OFS). The basic idea behind OFS is very similar to IP switching or Tag switching that is used to bypass IP routing in the internet. One of the main bottlenecks in the present Internet is routing at the IP layer and several methods have been proposed to alleviate this bottleneck by switching long duration flows at lower layers. Tag switching uses routing protocols to predefine routes within the network and assigns tags to the routes; packets are then switched based on these tags avoiding the need for routing table lookups [19]. IP switching dynamically sets up layer-2 (e.g., ATM) virtual circuits for connections that are perceived to be long [20].

This concept of lower-layer switching can be extended to switching large volume and/or long duration flows at the optical layer. That is, a lightpath can be established for large data transactions such as the transfer of large files or long duration and high bandwidth streams across the network. This optical flow will bypass all of the electronics in the network and be switched at the WDM layer as shown in Figure 12.

Unlike IP switching where an ATM virtual circuit is set up for a perceived flow, an optical flow switching protocol must be much more judicious in determining when to set up such flows. Lightpaths in a network are a scarce resource that cannot be arbitrarily assigned. In addition, without the use of wavelength converters,it is necessary to assign the same wavelength to a lightpath along its entire path, further restricting the available resources for lightpath establishment.

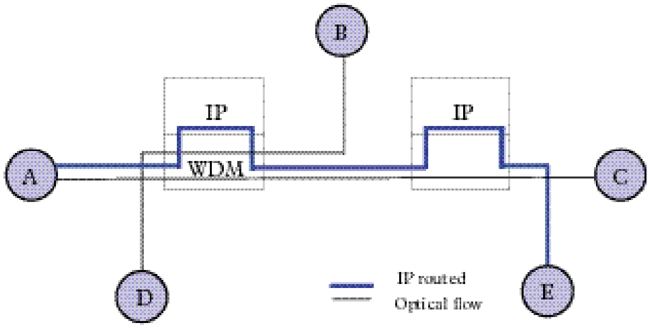

FIGURE 12: *Benefits of load balancing with clustered traffic.*

Consequently, for optical flow switching, it may be necessary for the application layer to inform the IP layer of the arrival and characteristics of large flows for switching. Also, the size of such a flow must be relatively large when compared to the flow setup time, which will likely exceed the network round trip delay time. In [21], a threshold mechanism is proposed to decide whether to switch a flow optically or electronically based on the size of the flow. The threshold is established so that the expected delay of the message is minimized, taking into account both the network congestion and flow setup times. The threshold can be dynamically changed based on network traffic conditions.

In addition to flow detection, an OFS algorithm must address issues such as connection establishment and routing and wavelength assignment. In [22] an Optical Burst Switching (OBS) mechanism for connection establishment is proposed. With OBS, a control packet is sent, on an out-of-band channel, to announce an upcoming burst. The control packet is then followed, after a short delay to allow time for processing the control message at every node, by a burst of data. Since a connection is not explicitly established before the burst is transmitted, it is possible that the control packet may fail to reserve resources at some node along its path and thus the burst may need to be dropped. Alternatively, one might consider a reservation based approach which establishes a connection before data transmission, thereby avoiding the need to drop bursts in mid-transmission.

To create an optical flow, an appropriate route and wavelength must be selected. To obtain optimal performance (minimize blocking probability), routing and wavelength assignment must be performed jointly [23]. Most previous work on routing and wavelength assignment assumes the existence of perfect network status information. That is, every node always knows the status of all wavelengths on all links in the network. In a practical network, this assumption is not realistic for two reasons. First, it may not be possible for each node to maintain the state information on the complete network and second, inevitably, this information will be somewhat outdated due to delays in the dissemination of status updates. The impact of having imperfect network state information is examined in [24], where a number of distributed approaches for optical flow switching are described.

6 CONCLUSION

Recent advances on WDM technology have significantly increased the usable transmission capacity of optical fiber. However, network transmission capacity is still limited by electronic processing capability. In fact,a dominant part of the cost of high-speed backbone networks is in the cost of the electronic switches, routers, and multiplexers. In order to reduce the size and cost of these electronic components a number of techniques have been developed that use WDM to bypass electronic switches and routers whenever possible. This paper surveys some of the recent work in this area, primarily focusing on dynamic techniques that take advantage of the configurability of WDM switches to reconfigure the electronic layer topology of the network in response to changes in traffic conditions.

Most of the concepts described in this paper are still at their infancy. The work described in this paper focuses mainly on WDM rings and assumes restrictive traffic patterns. In the future we hope to explore the application of these concepts to more general network topologies and traffic distributions.

7 REFERENCES

[1] S. Subramanian and R. Barry, “Wavelength Assignment in Fixed Routing WDM Networks”, Proc. of ICC ’97, Montreal, CA, June, 1997.

[2] J. Simmons, E. Goldstein, A. Saleh, “On The Value of Wavelength-Add/Drop in WDM Rings With Uniform Traffic,” OFC’98, San Jose, CA, February, 1998.

[3] Eytan Modiano and Angela Chiu, “Traffic Grooming Algorithms for Minimizing Electronic Costs in Unidirectional SONET/WDM Ring Networks,” CISS‘98, Princeton, NJ, March, 1998

[4] A. L. Chiu and E. H. Modiano, “Reducing Electronic Multiplexing Costs in Unidirectional SONET/WDM Ring Networks via Efficient Traffic Grooming,” Globecom ’98, Sydney, Australia, November, 1998.

[5] R. Berry and E. Modiano, “Grooming Dynamic Traffic in Unidirectional SONET Ring Networks,” to appear, OFC ’99, San Diego, CA, February 1999.

[6] R. Barry and P. humblet, “Models of Blocking Probability in All-Optical Networks with and without Wavelength Changers,” IEEE JSAC,Vol.14,pp.858-867, 1996.

[7] Y. Zhu, G. N. Rouskas and H. G. Perros, “Blocking in Wavelength Routing Networks, Part I: The Single Path Case,” Infocom ’99, NY, March, 1999.

[8] A. Birman, “Computing Approximate Blocking Probabilities in a Class of All-Optical Networks,” IEEE JSAC, June, 1996.

[9] B. Schein and E. Modiano, “Increasing Traffic Capacity in WDM Ring Networks Via Topology Reconfiguration,” CISS ’99,March, 1999.

[10] J. Bannister, L. Fratta, and M. Gerla, “Topological Design of WDM Networks,” Proc. Infocom ’90, Los Alamos, CA, 1990.

[11] R. Ramaswami and K. Sivarajan, “Design of Logical Topologies for Wavelength Routed Optical Networks,” IEEE JSAC, Vol. 14, No. 5, pp. 840-51, June, 1996.

[12] D. Bienstock and O. Gunluk, “ Computational Experience with a Difficult Mixed-Integer Multicommodity Flow Problem,” Mathematical Programming, Vol. 69, pp. 213-237, 1995.

[13] I. Baldine and G. N. Rouskas, “Dynamic Reconfiguration Policies for WDM Networks,” Infocom ’99,NY, March, 1999.

[14] G. N. Rouskas and M. H. Ammar, “Dynamic Reconfiguration in Multi-hop Lightwave Networks,” J. High Speed Networks, pp. 221-238, 1995.

[15] S. Banerjee and B. Mukherjee, “The Photonic Ring: Algorithms for Optimized Node Arrangements,” Fiber and Integrated Optics, Vol. 12, No. 2, pp. 133-17, 1993.

[16] J.-F. P. Labourdette and A.S. Acampora, “Logically Rearangable Lightwave Networks,” IEEE Trans. On Communications,August, 1991.

[17] A.S. Acampora, J.-F. P. Labourdette and G. W. Hart, "Branch-Exchange Sequences for Reconfiguration of Lightwave Networks," IEEE Transactions on Communications, Vol. 39, No. 10, pp. 2822-2832, Oct., 1994.

[18] A. Narula-Tam and E. Modiano, "Dynamic Load Balancing for WDM-based Packet Networks," Submitted to Infocom 2000.

[19] Y. Rekhter, et. al., ``Cisco Systems' Tag Switching Architecture Overview," IETF RFC 2105, Feb. 1997.

[20] P. Newman, T. Lyon, and G. Minshall, ``Flow Labeled IP: A Connection Approach to ATM." IEEE Infocom, 1996.

[21] D. Brant, "Threshold Schemes for Optical Flow Switching," Masters of Engineering thesis,MIT, Dept.of Electrical Engineering and Computer Science, 1999.

[22] M. Yoo and C. Qiao, "Just-enough-time (JET): A high speed protocol for bursty traffic in optical networks," IEEE/LEOS Summer Topical meeting, August, 1997.

[23] R. Ramaswami and K. Sivarajan, "Routing and Wavelength Assignment in All-Optical Networks," IEEE/ACM Transactions on Networking, Vol. 3, pp.489-500, 1995.

[24] R. Ganguly and E. Modiano, "Distributed Algorithms and Architectures for Optical flow Switching in WDM networks," Submitted to Infocom 2000.

[25] O. Gerstel, P. Lin and G. Sasaki, "Wavelength assignment in a WDM ring to minimize cost of embedded SONET rings," Infocom '98, San Francisco, CA, March, 1998.

[26] X. Zhang and C. Qiao, "An Effective and Comprehensive Solution to Traffic Glooming and Wavelength Assignment in SONET/WDM Rings", Conf. on All-optical Networking, SPIE Proc. Vol. 3531, Boston, MA, Sept. 1998.

[27] Zhang and C. Qiao "On scheduling all-to-all personalized connections and cost-effective designs in WDM rings", IEEE/ACM ToN, Vol. 7, No. 3, 1999.

[28] B. Mukherjee, "Optical Communication Networks," McGraw-Hill, NY, NY, 1997.

This work was supported by the Defense Advanced Research Project Agency (DARPA) under the Next Generation Internet initiative.

While the actual load on each of the fibers remains the same, the load on each of the "electronic" links and the amount of electronic processing is reduced.

Eytan Modiano
MIT Lincoln Laboratory
Lexington, MA 02420
Modiano@mit.edu

Eytan Modiano received his B.S.degree in Electrical Engineering and Computer Science from the University of Connecticut at Storrs in 1986 and his M.S. and PhD degrees, both in Electrical Engineering, from the University of Maryland,College Park,MD, in 1989 and 1992 respectively. He was a Naval Research Laboratory Fellow between 1987 and 1992 and a National Research Council Post Doctoral Fellow during 1992-1993 while he was conducting research on security and performance issues in distributed network protocols. Between 1993 and 1999 he was with the Communications Division at MIT Lincoln Laboratory where he worked on communication protocols for satellite,wireless, and optical networks and was the project leader for MIT Lincoln Laboratory's Next Generation Internet (NGI) project. In 1999, he joined the faculty of the Aeronautics and Astronautics Department, and the Laboratory for Information and Decision Systems (LIDS) at MIT, where he conducts research on communication networks and protocols with emphasis on satellite and hybrid networks, and high speed networks.

Aradhana Narula-Tam
MIT Lincoln Laboratory
Lexington, MA 02420
arad@ll.mit.edu

Aradhana Narula-Tam received the B.S.E. degree in electrical engineering from the University of Pennsylvania,Philadelphia, in 1990 and the S.M.and Ph.D. degrees also in electrical engineering from the Massachusetts Institute ofTechnology (MIT),Cambridge, in 1992 and 1997, respectively. From 1997—1998 she was in the Research and Advanced Development Group at Motorola Information Systems Group, Mansfield, MA. She is now in the Advanced Networks group at MIT Lincoln Laboratory. Her research interests include communication systems and optical networks.