Dynamic LSP Routing in IP/MPLS over WDM Networks

Sunggy Koo, Gokhan Sahin, and Suresh Subramaniam

***Abstract*— We consider an IP/MPLS over WDM network, in which label switched routers (LSRs) in the IP/MPLS layer are interconnected through optical cross-connects (OXCs) in the optical core network (WDM layer) providing an end-to-end wavelength routing capability. In this paper, we study a dynamic label switched path (LSP) routing problem for the three different network models of the IP/MPLS over WDM network, namely, Overlay, Augmented, and Peer models. For the overlay model, we propose two algorithms: ECF_OVLY and MLH_OVLY. In ECF_OVLY, a network always tries to use existing capacity first, whereas in MLH_OVLY, a network finds a path with the minimum number of logical hops for an LSP request. We also propose, for the augmented model, two simple and efficient dynamic LSP provisioning algorithms, called MCPI_AUG and DCPI_AUG, utilizing different type/amount of summarized capacity information from the WDM layer, namely Minimum CaPacity Information (MCPI) and Detailed CaPacity Information (DCPI). We compare the proposed algorithms with the existing algorithms available for the overlay [1] and peer [2], [3] models. The algorithms are compared and evaluated using two key performance measures: LSP blocking probability and network (lightpath) utilization. Simulation results show that at low loads with a limited number of ports available in the network, DCPI_AUG achieves an order of magnitude better blocking performance than the algorithm in [2] and outperforms the one in [3] by more than three times. It also achieves higher network utilization than the one in [2] by more than 10 % and the one in [3] by 2-7 % depending on the traffic load. Considering the small amount of information that is exchanged between the layers in the augmented model, these results suggest that the augmented model can be a practically good compromise between the overlay and peer models.**

***Index Terms*— Optical networks, IP-over-WDM, GMPLS, peer model, overlay model, augmented model, routing algorithms, blocking probability, utilization, port-limited network.**

I. Introduction

As communication networks have expanded to accommodate rapidly growing data traffic, the traditional network architecture with a full range of protocol stack (e.g., IP/ATM/SONET/WDM) has suffered from heavy network overheads and increasing management complexity. It is widely believed that a simpler network architecture of IP/MPLS directly over a WDM transport layer will be the most prominent network solution in the future. Emerging technologies and standardization activities on Multi-Protocol Label Switching (MPLS) [4], Generalized MPLS (GMPLS) [5], and User-Network Interface (UNI) [6] have been playing a key role in moving toward implementation of the IP/MPLS over WDM architecture. In this architecture, high-speed routers equipped with MPLS functions, called *label switched routers* (LSRs) are interconnected by intelligent optical core networks that provide dynamic point-to-point connectivity in the form of lightpaths. The resulting end-to-end path, which may traverse more than one lightpath, is referred to as a *label switched path* (LSP). In the IP/MPLS over WDM network, the *logical* topology seen by the IP/MPLS layer is the topology of the LSRs with logical links that are the lightpaths dynamically provided by the WDM layer with granularity of a whole wavelength. Note that a *logical edge* between two LSRs could comprise multiple *logical links* established on different wavelengths.[1]

In an IP/MPLS over WDM network, there are several architectural alternatives including overlay, augmented, and peer models [7], [8], [9]. As summarized in Table I, one of the key differences among the models is how much and what kind of network information can be exchanged between the two layers. In the peer model, the topology and other network information (e.g., routing and link state) are shared among all network elements across the layers by a unified signaling protocol and control plane. Such a model may be appropriate when the transport and service networks are operated by a single entity, whereas the overlay and augmented models are more suitable for the case with different management/ownership of each layer. In the overlay network model, there is no specific network information exchanged between the layers, since the routing in each layer is done separately with each layer's own signaling and control plane. The augmented model provides a compromise between the two extreme cases by allowing the exchange of some network information between the layers, such as reachability and/or summary of link state information (e.g., residual capacity), depending on a necessary and specific agreement between the two layers.

There are two different assumptions on the nature of network traffic that are often made in developing provisioning and routing algorithms: static and dynamic traffic. With static traffic, it is assumed that the traffic demands between all ingress-egress pairs are known at the time of initial provisioning. Then,

Manuscript received May 24, 2005; revised July 10, 2006. An earlier version of this paper appeared in the Proceedings of Infocom 2004. The work was supported in part by DARPA Grant N66001-00-18949 (co-funded by NSA) and by NSF Grants ANI-9973111, CNS-0434956, CNS-0519911.

S. Koo is with Spirent Communications, Calabasas, CA (email: Sunggy.Koo@spirentcom.com).

G. Sahin is with Miami University, Oxford, OH (email: sahing@muohio.edu).

S. Subramaniam is with the ECE Dept., George Washington University, Washington, DC 20052 (email: suresh@gwu.edu).

This work was done in part when all authors were with the George Washington University.

Digital Object Identifier 10.1109/JSAC-OCN.2006.021905.

[1] In this paper, we use the terms *logical edge* and *logical link* interchangeably, unless it is necessary to distinguish one from the other.

0733-8716/06$20.00 © 2006 IEEE

TABLE I
NETWORK MODELS IN IP/MPLS OVER WDM NETWORK.

	Overlay	Augmented	Peer
Routing	Separated	Separated	Integrated
Network information exchanged	No information	Part or Summary	Full information
Signaling and control plane	Separated	Separated	Unified

one can formulate an optimization problem in order to make the most efficient use of network resources, as in [10] for example. The increasing agility of optical components, however, suggests that optical network traffic in the future is likely to be dynamic in nature with lightpaths signaled on demand. Typically, the bandwidth demand of LSP requests arriving at and departing from a network dynamically is smaller than that of a lightpath on a wavelength. Multiple sub-wavelength LSPs can possibly be groomed into a single lightpath at an edge LSR. Under the dynamic IP/MPLS traffic assumption, a network has to make a routing decision upon every new LSP request arrival without any *a priori* knowledge about future requests. Since any network reconfiguration would inevitably disrupt the existing LSP traffic, it may not be desirable or practical to reconfigure existing lightpaths and reroute existing LSPs upon every new LSP request arrival. A more practical scenario would be to let the logical network evolve in response to LSP traffic demands, and re-optimize the entire network for more efficient use of network resources after some operational period.

In this paper, we study the dynamic LSP routing problem for the three different network models in IP/MPLS over WDM networks. Loosely speaking, this problem refers to the selection of an appropriate logical route for an arriving LSP request, including the opening of any new lightpaths at the optical layer. There are a few algorithms available for this problem that we are aware of – one for the overlay model [1] and two for the peer model [2], [3]. In this paper, we propose two new algorithms each for the overlay and augmented models: ECF_OVLY and MLH_OVLY for the overlay model, and MCPI_AUG and DCPI_AUG for the augmented model. As far as we are aware, no algorithms have been proposed and evaluated for the augmented model thus far. An observation worthy of note here is that all the previous algorithms assumed that there is no limit on the number of optical ports per LSR. As we will see later, the number of ports per LSR is a key determinant of the performance of the algorithms and the performances of the algorithms for the peer model deteriorate rapidly as the number of ports becomes more limited. Since a large majority of traffic at an OXC is likely to be pass-through traffic and optical ports tend to be expensive, the number of optical ports at an LSR is likely to be small, and therefore, the above result has serious practical implications. In our simulations, our DCPI_AUG algorithm which uses a relatively small amount of information about the WDM network gives better LSP blocking performance than [2], [3] in some cases when there are only a limited number of ports available in the network. In terms of network utilization (to be defined precisely later), DCPI_AUG outperforms the peer model algorithms in [2], [3] in all cases studied in this paper.

The rest of the paper is organized as follows. In Section II, we describe the problem with an illustrative example. We explain the differences in the problem for each model in terms of the available information regarding the WDM layer for LSP provisioning in the IP/MPLS layer. Our motivation and major contributions of the paper are summarized in Section III. Section IV provides the details of the previous work for the overlay [1] and peer [2], [3] models. Section V describes our proposed algorithms – two each for the overlay and augmented models. In Section VI, we present the performance evaluation results and compare the proposed algorithms with the existing algorithms. Section VII concludes the paper.

II. DYNAMIC LSP PROVISIONING

A. Goals and Constraints

Consider an IP/MPLS over WDM network to which LSP requests arrive (and from which they depart) dynamically. The WDM network is capable of providing lightpaths between LSR pairs upon request, if sufficient resources (i.e., wavelengths) are available to satisfy the request. Wavelength conversion is assumed to be unavailable in this paper, though the work can be extended in a straightforward manner to networks with wavelength conversion. A lightpath may be requested between a pair of LSRs if there is an unused optical port available at both LSRs.

Upon the arrival of an LSP request for sub-wavelength bandwidth, the network has to decide how to accommodate/route it. The LSP request may be satisfied by routing the LSP over the existing logical network if it has sufficient capacity available, or one or more lightpaths may be signaled and set up, and used in combination with existing logical capacity to accommodate the new LSP request.

Such routing of an LSP request must be done without any a priori knowledge of future requests, such that two goals can be achieved at the same time, namely, (a) to accept as many LSP requests as possible over a period of time, and (b) utilize the lightpaths as efficiently as possible. The number of wavelengths per physical link in the WDM layer and the number of ports per LSR in the IP/MPLS layer are two major constraints on network resources to be considered to make a *good* routing decision. Opening a new lightpath between two LSRs in the IP/MPLS layer costs one port at both ingress and egress LSRs and one wavelength on the physical link(s) along the path in the WDM layer. Without wavelength conversion capability, the wavelength on all the physical links along a lightpath has to be the same. An LSP request will be blocked if there is not enough capacity available on the current logical topology, and there is no wavelength and/or port available to open a new lightpath.

Depending on the network model, the LSRs receiving LSP requests in the IP/MPLS layer may have different information about the current status of the entire network. We look at each of the three network models below.

B. Overlay Model

In the overlay model, each LSR keeps only IP/MPLS layer information, such as residual capacity on all of the existing logical links and the number of unused ports in the LSRs. The IP/MPLS layer only receives a response of whether a requested lightpath can be set up or not, from the WDM layer. Therefore, in the overlay model, a network has to decide whether it should use the existing logical links or open new lightpath(s) for a new arriving request. If it chooses to use the existing logical topology, then how to route the request over the existing logical topology has to be decided. If the network would open new lightpath(s), it has to decide the logical edge(s) (LSR pair(s)) on which to open the lightpaths, without any network information from the WDM layer.

C. Peer Model

In the peer model, on the other hand, each LSR keeps information about the topology and status of physical links (e.g., availability of each wavelength) in the WDM layer as well as logical links in the IP/MPLS layer. As often visualized, a network in the peer model can be seen as one graph with both LSRs and OXCs interconnected with physical and logical edges as shown in Fig. 2. In this case, an integrated routing can be done with a unified control plane, where an integrated routing scheme for both layers decides routing over existing logical links, and routing and wavelength assignment (RWA) in the WDM layer at the same time. Note that this is a fundamental difference in the dynamic LSP routing problem between the peer model and the other two models. In the overlay and augmented models, the RWA in the WDM layer is beyond the scope of the LSP routing problem in the IP/MPLS layer.[2]

D. Augmented Model

For the augmented model, however, the available information at LSRs in the IP/MPLS layer depends on a specific agreement between the two network entities. In this paper, we assume that the IP/MPLS layer may utilize a small amount of capacity information passed from the WDM layer. Specifically, we define two different types of the capacity information from the WDM layer; Minimum CaPacity Information (MCPI) and Detailed CaPacity Information (DCPI). MCPI is defined as the WDM layer's capability to provide at least one new lightpath between every LSR pair. For more detailed capacity information, DCPI is defined as the number of new lightpaths that the WDM layer can further provide between every LSR pair in the current state of the WDM network (see Section V-B for formal definitions). In any case, a network in the augmented model has to make the same provisioning decision as in the overlay model, except that it has more information about the status of the WDM layer than in the overlay model.

[2]In this paper, we assume a fixed minimum-hop routing and the first-fit wavelength assignment in the WDM layer for the overlay and augmented models. Our focus here is the LSP routing problem at the IP/MPLS layer.

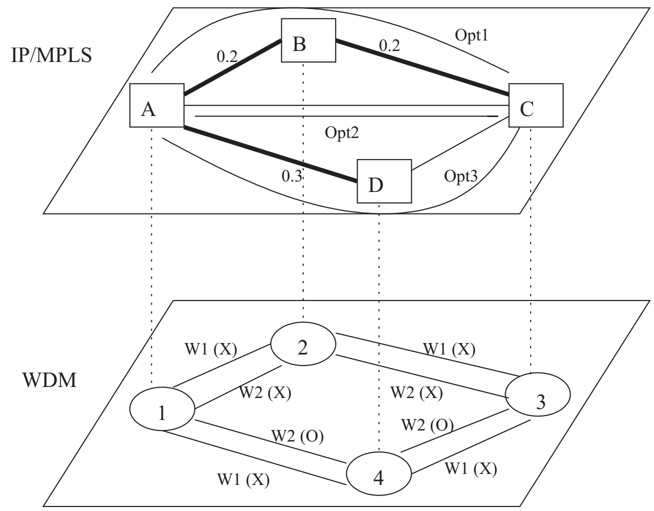

Fig. 1. An illustrative example of dynamic LSP provisioning.

E. An Illustrative Example

Let us illustrate the LSP routing problem using a simple illustrative example. Fig. 1 shows 4 LSRs in the IP/MPLS layer. In the WDM layer, 4 OXCs are connected by physical links with 2 wavelengths ($w1$, $w2$). In this example, each LSR in the IP/MPLS layer is connected to a corresponding OXC in the WDM layer. Assume that currently the network has the following residual capacity on the logical edges: A-B (0.2), B-C (0.2), and A-D (0.3)[3], which were already established using some of the wavelengths on the physical links (marked X).[4] Now, let us say that a new LSP request for 0.2 units of bandwidth between LSRs A and C arrives. For this request, the IP/MPLS layer has three routing options: (1) routing over the existing logical links A-B and B-C without opening any new lightpath, (2) routing with a new lightpath between A and C using $w2$ available on both physical links 1-4 and 4-3, and (3) routing over the existing logical link A-D and a new lightpath between D and C. Depending on the choice, more or fewer LSP requests may be accommodated in the future.

III. Motivation And Contributions

In many instances in practice, there exists a separate management/ownership for each network layer that prefers to keep its network information (e.g., topology and/or routing) from other layers. It is also true in the case of the peer model that a significant number of control messages have to be flooded across the network layers frequently to keep all the information updated. On the other hand, in the overlay model, the entire network could be managed inefficiently due to the lack of information exchanged between the two layers. Even though it is believed that a suitable augmented model could benefit from the advantages of both overlay and peer models, there has been no routing algorithm proposed for this model, and there is little understanding of what

[3]The unit of LSP bandwidth request is normalized to the bandwidth of a lightpath for all examples in this paper.

[4]Note that logical links that are utilizing W1 on WDM link 1-2, 2-3, and 4-3 are not shown in the figure because they are assumed to be in full use for other traffic.

kind of information would be most helpful in making LSP routing decisions. A major contribution of this paper is the presentation of simple dynamic provisioning algorithms based on two different types of WDM network information available for the augmented model. The paper also identifies different types/amount of network information exchanged between the layers and investigates their effects on performances of the algorithms for the augmented model.

As mentioned earlier, one may expect that the number of optical ports at an LSR to be limited, and therefore, it may not be possible to open a lightpath between a pair of LSRs even if there are enough wavelengths on the physical links connecting the two LSRs. All previous routing algorithms in the literature assumed that the network was wavelength-limited, and as indicated earlier, the performances of some of those algorithms degrade rapidly as they become more port-constrained. Another main contribution of this paper is a thorough investigation of the effects of a limited number of ports on the performance of the dynamic LSP routing algorithms discussed in this paper.

We have implemented the existing algorithms and our proposed algorithms, and present extensive simulation results comparing the performances of various algorithms under various network models. To the best of our knowledge, no such complete performance comparison of dynamic LSP routing algorithms across the different network models is available.

IV. Previous Work

A. Y_OVLY for Overlay Model

In [1], Ye et al. presented a simple integrated routing/protection scheme to dynamically allocate restorable bandwidth-guaranteed paths in IP over WDM networks. Since our focus in this paper is to set up primary LSPs, we take only the primary path routing part of their algorithm into consideration, which we refer to as Y_OVLY. In Y_OVLY, the network first tries to route the request over the residual capacity on existing logical links. If it finds the existing capacity in the logical layer insufficient for the arriving request, then it tries to add capacity by opening a new lightpath directly between the ingress and egress LSRs.

Let us consider the same example shown in Fig. 1. However, this time assume that the new LSP request requires 0.3 units of bandwidth, instead of 0.2 as in the previous case, and that $w1$, instead of $w2$, is available on the physical link between OXC1 and OXC4. As a result, the existing logical topology does not have enough capacity for a demand of 0.3 units between A and C, and a new lightpath on the direct logical edge between A and C cannot be established because no wavelength is available on both physical links 1-4 and 4-3. In this case, Y_OVLY will block the LSP request.

B. K_PEER and Z_PEER for Peer Model

In [2], Kodialam and Lakshman proposed an integrated dynamic IP and wavelength routing algorithm that utilizes the combined knowledge of resource and topology information in both layers, assuming a peer network model. We refer to their algorithm as K_PEER in this paper. In [3], Zheng and Mohan proposed a dynamic protection scheme in integrated IP/WDM

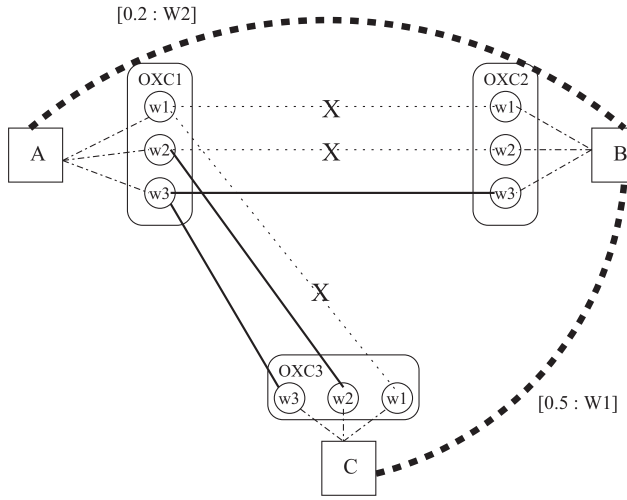

Fig. 2. An illustrative example of integrated graph.

networks. Since their scheme also considers routing primary as well as backup paths (LSPs), we consider only its primary path routing part, which is referred to as Z_PEER in this paper.

Both K_PEER and Z_PEER use an integrated graph, in which both the physical links and logical links (lightpaths shown as thick dashed lines) coexist as in Fig. 2. It is also a layered graph, wherein each wavelength is represented by a corresponding sub-node in the graph. When a new lightpath for a logical link is established using wavelength i, the corresponding physical link of the wavelength i is removed from the graph, and a logical link corresponding to the new lightpath is added between the LSRs. In Fig. 2, a logical link between A and B is created using $w2$, and say it has a residual bandwidth of 0.2 units. Likewise, 0.5 units of bandwidth are available on the logical link between B and C, using wavelength $w1$ on the physical links OXC3-OXC1 and OXC1-OXC2. Note that this lightpath between B and C does not use any port in LSR A because it by-passes that LSR.

In both algorithms, the basic procedure is to assign a cost to both the logical and physical edges first, and then find a minimum-cost path using a shortest path algorithm such as Dijkstra's algorithm. The resulting path may contain some logical links already established and/or new lightpath(s) to be opened for this LSP request. The main difference between the two algorithms lies in the way they assign cost to each edge in the graph. In K_PEER, all minimum-cut sets [11] for every ingress-egress LSR pair are first identified. All edges belonging to any minimum-cut set are defined to be *critical* for that ingress-egress LSR pair. Then, for each edge, the number of LSR pairs for which that edge is critical is noted, and is assigned as the cost of that edge. The idea behind this procedure is to find a route that would avoid using edges that are bottleneck edges for several ingress-egress LSR pairs. Note that the algorithm is rather complex and requires the computation of N^2 minimum-cut sets, where N is the number of LSRs.

In Z_PEER, the cost of an edge is assigned as the number of physical hops it uses. Therefore, the cost of a logical link is equal to the number of physical hops that it takes to establish that lightpath. For the example in Fig. 2, the cost of the logical edge between B and C is 2, whereas 1 is assigned as the cost for the edge between A and B. Z_PEER introduces a control parameter for the cost of a physical link. According to the numerical results in [3], the best value for the control parameter turned out to be 1, which implies that the cost of each physical edge is 1. This, in turn, means that the Z_PEER algorithm chooses the route with the minimum number of physical hops to satisfy the LSP request.

Note that a route that is chosen by these algorithms may consist of existing logical links as well as new physical links. The algorithms open a new lightpath between LSRs which are connected by the physical links that are part of the chosen route.

V. Proposed Algorithms

Before we proceed to present our algorithms for the overlay and augmented models, let us define the performance measures that are used to evaluate the algorithms. The main objective of all the algorithms studied in this paper is to minimize the probability of blocking an LSP request. Besides the blocking probability, we consider network (lightpath) utilization as another performance measure. To formally define utilization, let $B(t)$ be the total bandwidth of all the LSPs that are being served at time t. Correspondingly, let $C(t)$ be the total capacity of all the lightpaths that are operational at time t. Since the bandwidths are assumed to be normalized to lightpath capacity, $C(t)$ is just the number of lightpaths that exist at time t. Then, the network or lightpath utilization at time t, denoted $U(t)$ is defined as $U(t) \stackrel{\text{def}}{=} \frac{B(t)}{C(t)}$. The time-average utilization, U, (which we simply call utilization) is then defined as

$$U \stackrel{\text{def}}{=} \lim_{T\to\infty} \frac{1}{T} \int_0^T U(t)dt.$$

It is obviously desirable that U be as high as possible while the main objective of minimizing blocking probability is achieved.

All our algorithms (except MLH_OVLY) work by assigning costs to the edges of a graph and running a shortest-path algorithm such as Dijkstra's algorithm to obtain the route for the LSP. The graph that is used by the algorithms is a logical graph in which the nodes are the LSRs and the edges represent either potential new lightpaths that may be set up by the WDM layer or lightpaths that have already been established and that have adequate residual capacity to handle the arriving LSP request. In this paper, we assume that the entire traffic of a requested LSP has to be transmitted without splitting it onto multiple logical links (lightpaths). By default, the logical graph is a complete graph (with a logical edge between every pair of LSRs). However, if the IP/MPLS layer has information that a logical edge has inadequate residual capacity *and* no new lightpaths can be set up, then that logical edge is removed from the graph (or, equivalently, assigned a cost of infinity). After running the shortest-path algorithm, the resulting route is a combination of new lightpaths that need to be established and existing lightpaths that have enough capacity. The IP/MPLS layer must then signal to the WDM layer to establish any new lightpath(s) that are selected. In some cases, the WDM layer may not be able to set up the new lightpath(s) because of insufficient wavelengths (and the IP/MPLS layer was unaware of this because of the limited information available). In this case, the LSP request is blocked. (It is possible to rerun the shortest-path algorithm with the infeasible logical edge removed, but we do not pursue this direction in this paper.) Note that the graph used here is different from the integrated graph used by the peer model algorithms. That is a much more detailed graph that enables the joint selection of current lightpaths and new lightpaths including their wavelengths and routes. As mentioned earlier, routing and wavelength assignment of lightpaths is not our concern in this paper because that is not a function of the IP/MPLS layer.

The various algorithms differ in how costs are assigned to the logical edges using the information available. We now present our algorithms.

A. Algorithms for Overlay Model

We make the same assumption as in [1], where only one lightpath is allowed to be established per LSP request.[5] We do so for two reasons – one is to enable us to make a fair comparison among the algorithms for overlay model, and the other is to ensure that the network utilization does not decrease very much. Note that the bandwidth of an LSP request is at most 1, and thus $B(t)$ can increase by at most 1, whereas each new lightpath that is opened increases $C(t)$ by 1.

1) ECF_OVLY: The objective of ECF_OVLY is to always try to use the existing capacity first (ECF). If there is not enough capacity, then it tries to open a new lightpath. The main idea behind this objective is to avoid the situation where lightpaths are under-utilized for a long time. The cost of the edges with enough residual capacity for the arriving request is assigned to be 1, and all other edges are assigned a large cost M that is guaranteed to be larger than the maximum number of logical hops over all possible paths in the network.

The resulting minimum-cost path will have the edges with enough residual capacity whenever possible. If the path contains more than one edge of cost M, then this implies that at least two new lightpaths are necessary (by the choice of M). In this case, we block the request because we are restricted to allowing no more than one new lightpath per LSP request. On the other hand, if there is exactly one edge with cost of M on the selected path, then the network tries to open a new lightpath on that edge. Unlike in the algorithms for the peer model, however, the IP/MPLS layer does not have any information about the status of the WDM layer. Therefore, it is possible that the WDM layer may not be able to provide the new lightpath requested by the IP/MPLS layer. In this case, the LSP request will be blocked.

2) MLH_OVLY: The objective of MLH_OVLY is to minimize the total number of logical hops that an LSP has to traverse. With the limited information about the residual capacity on the logical links only, the IP/MPLS layer in MLH_OVLY tries to save network resources for potential requests in the future by minimizing the number of logical hops.

[5] In general, this kind of restriction on the number of lightpaths to be opened per LSP request is not necessary.

In MLH_OVLY, the network first tries to route an arriving LSP request using a single hop on the direct logical edge between the ingress and egress LSRs, which means either using the residual capacity, or opening a new lightpath (logical link) on that direct edge. If the effort to accommodate the request on a single hop fails, the network checks all other LSR pairs (logical edges) between which it may open a new lightpath, to find an end-to-end path using that new lightpath and existing logical links with enough capacity. If there are multiple candidate LSR pairs for a new lightpath, it chooses to open a new lightpath between an LSR pair such that it could route the LSP request using the minimum number of logical hops.[6]

Let us consider the same situation in Section IV-A, where Y_OVLY would have blocked the request because it could not open a direct lightpath between A and C. In MLH_OVLY, however, it would open a new lightpath between D and C using $w2$, then route the request for 0.3 units of bandwidth between A and C successfully, with the existing logical link (A-D) and the new logical link (D-C) as the two logical hops.

B. Algorithms for Augmented Model

In this paper, we assume that only a summary of capacity information from the WDM layer is shared with the IP/MPLS layer in the augmented model. We consider the two different scenarios depending on the type/amount of capacity information regarding available to the IP/MPLS layer: MCPI and DCPI. In the MCPI case, only whether or not a new lightpath can be opened between between LSRs i and j is assumed to be known by the IP/MPLS layer. The minimum capacity information from the WDM layer is expressed as B_{ij} that takes on the value of one if at least one new lightpath can be established between LSRs i and j, for all i and j, and zero otherwise. In the DCPI case, we assume that the WDM layer passes L_{ij}, the number of lightpaths that can be established between LSRs i and j, for all i and j, to the IP/MPLS layer.

Without wavelength conversion capability,[7] L_{ij} is the number of common wavelengths that are available over every physical link on the path found by the routing algorithm in the WDM layer. These types of capacity information can be passed in various forms through a specifically designed UNI, or by the MPLS/GMPLS signaling mechanism such as RSVP-TE [12]. For a network with N ingress/egress LSRs, there are N^2 number of B_{ij}'s or L_{ij}'s to be passed to the IP/MPLS layer from the WDM layer. We believe that this is a relatively small amount of information compared to the information that is flooded across the network layers in the peer model. The information that must be actually passed to the IP/MPLS layer may even be less than this, as the WDM network state changes only when a lightpath is set up or torn down, and even then, not all L_{ij}'s or B_{ij}'s may change.

We illustrate with a simple example the information that must be updated and exchanged between the two layers. Consider a 3-node WDM network topology consisting of nodes A, B, and C, and links AB and BC. We consider a semi-dynamic traffic model, where LSP connection requests arrive sequentially over time, but each LSP request that is established remains in the network permanently. The capacity of each wavelength is normalized to 1, and there are 100 wavelengths on each link. Now consider the following sequence of four LSP requests and their bandwidths: AB: 0.1, AC: 0.2, BC: 0.1, AC: 0.2, and assume that this sequence of requests is repeated periodically. The table below (Table II) summarizes the relevant metrics for different algorithms after each LSP request, i.e., link capacities for the peer algorithms, and L_{ij}'s for DCPI_AUG. The B_{ij}'s for MCPI_AUG are all 1 and are not shown in the table. The metrics that are updated are in bold. For the peer model, it is easy to see that every connection request triggers at least one update, and there will be a total of 12 capacity changes for every 8 LSP requests. On the other hand, it will take 800 LSP requests for a total of 3 B_{ij} changes in the MCPI_AUG algorithm (no changes will occur until all of the wavelengths on a link are full). For DCPI_AUG, there will be a total of 3 L_{ij} changes for every 8 LSP requests. It is also clear that the reduction in the number of the updates in MCPI_AUG increases with the grooming factor (i.e., the maximum number of LSPs that a wavelength can support) and the number of wavelengths, and that for DCPI_AUG increases with the grooming factor. This concludes the illustrative example.

TABLE II
AN EXAMPLE TO ILLUSTRATE THE AMOUNT OF INFORMATION EXCHANGED IN THE PEER AND AUGMENTED MODEL ALGORITHMS.

Algorithm	Peer		DCPI_AUG		
Metric	Cap. (AB)	Cap. (BC)	L_{AB}	L_{BC}	L_{AC}
Initial	100	100	100	100	100
LSP 1 (AB)	**99.9**	100	**99**	100	**99**
LSP 2 (AC)	**99.7**	**99.8**	99	**99**	99
LSP 3 (BC)	99.7	**99.7**	99	99	99
LSP 4 (AC)	**99.5**	**99.5**	99	99	99
LSP 5 (AB)	**99.4**	99.5	99	99	99
LSP 6 (AC)	**99.2**	**99.3**	99	99	99
LSP 7 (BC)	99.2	**99.2**	99	99	99
LSP 8 (AC)	**99.0**	**99.0**	99	99	99
LSP 9 (AB)	**98.9**	99.0	**98**	99	**98**
LSP 10 (AC)	**98.7**	**98.8**	98	**98**	98
. . .					

The IP/MPLS layer has its own limitation on the number of optical ports that are available on an LSR. Let P_i be the number of unused ports available in LSR i. In the MCPI case, the minimum capacity information, D_{ij}, i.e., if it is possible to set up a lightpath between LSRs i and j, can be found as follows:

$$D_{ij} = \min\{B_{ij}, P_i, P_j\}. \quad (1)$$

In the DCPI case, C_{ij}, the number of lightpaths that can *actually* be established between LSRs i and j, considering both the number of wavelengths available in the WDM layer and the unused ports in the IP/MPLS layer, can be found as follows:

$$C_{ij} = \min\{L_{ij}, P_i, P_j\}. \quad (2)$$

Now, let r_{ij}^w be the residual capacity on the existing logical link established using wavelength w between LSR i and j. Then,

[6]In MLH_OVLY, a path with fewer hops consisting of a new lightpath and existing logical link(s) is preferred over one with more hops using only existing capacity.

[7]Recall that we assume that the WDM transport network does not have wavelength conversion capability for all network models discussed.

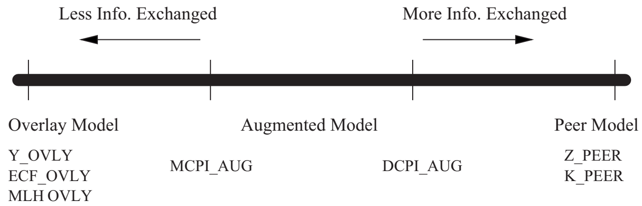

Fig. 3. Spectrum of algorithms for network models.

R_{ij}, the total amount of residual capacity on the logical edge ij, is:

$$R_{ij} = \sum_w r_{ij}^w. \tag{3}$$

In this paper, we propose the algorithms MCPI_AUG and DCPI_AUG for each case respectively.

Fig. 3 shows the spectrum of the algorithms studied in this paper according to the type/amount of network information exchanged between the two layers.

1) MCPI_AUG: As in ECF_OVLY, the basic procedure of MCPI_AUG is to assign a cost to all the logical edges first, and then find a minimum-cost path using a shortest path algorithm. In MCPI_AUG, ϕ_{ij}, the cost of a logical edge between LSR i and j, is defined as follows:

$$\phi_{ij} = \begin{cases} \infty & \text{if} \quad D_{ij} = 0, \quad \text{and } r_{ij}^w < b \quad \forall\, w \\ M & \text{if} \quad D_{ij} > 0, \quad \text{and } r_{ij}^w < b \quad \forall\, w \\ 1 & \text{if} \quad r_{ij}^w \geq b \quad \text{for} \quad \text{any} \quad w, \end{cases} \tag{4}$$

where M is a large number that is guaranteed to be larger than the maximum number of logical hops over all possible paths in the network. Once the costs of all logical edges are decided, one may find a minimum-cost path between an ingress and egress LSR. By assigning edge costs with the given small amount of network information from the WDM layer as defined above, MCPI_AUG always tries to use the edges with enough residual capacity first in order to keep network utilization as high as possible. If there is no existing logical link that has enough capacity for the logical edge on the path, then it opens a new lightpath (logical link) on that logical edge. There could be more than one logical link with enough capacity on a logical edge. In that case, MCPI_AUG picks the one with minimum capacity among them.

Since a single LSP's traffic may not be split over multiple lightpaths, even for some edges for which $R_{ij} \geq b$, it could be true that $\phi_{ij} = \infty$ if $r_{ij}^w < b$, $\forall\, w$.

Let us illustrate the working of the algorithm with an example. Fig. 4 shows the current capacity on logical edges, including the number of potential lightpaths (C_{ij}) and residual capacity of each logical link already established (r_{ij}^w).[8] On the logical edge between LSRs B and C, for example, there are two existing logical links with 0.4 and 0.5 units of residual bandwidth, respectively. $C_{BC} = 1$ means that another new lightpath can be opened on the logical edge B-C. Between C and E, there is one logical link opened, but it is fully used so that there is no residual capacity left on that logical link.

Assume that there is a new LSP request between A and E for 0.2 units of bandwidth. Fig. 5 shows the cost of each

[8]Logical edges that have no capacity left ($C_{ij} = 0$, $r_{ij}^w = 0$, $\forall\, w$) are not shown in the figures.

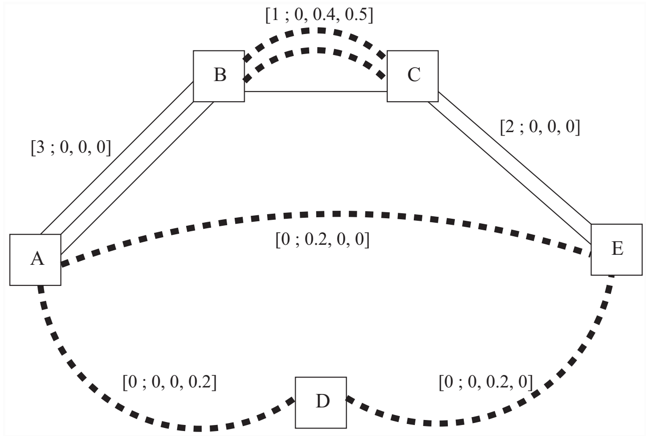

Fig. 4. Current capacity on logical edges [C_{ij} ; r_{ij}^w].

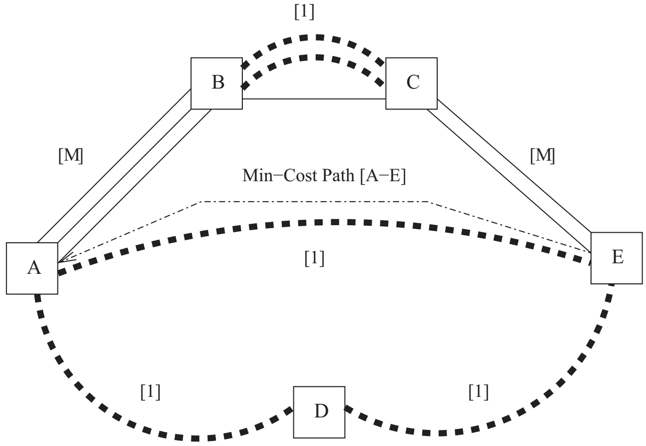

Fig. 5. For MCPI_AUG, cost of logical edges [ϕ_{ij}].

logical edge assigned by MCPI_AUG on the same example as in Fig. 4. Note that the C_{ij}'s are not available to MCPI_AUG, only the $D_{ij}'s$ are. ($D_{ij} = 1$ if $C_{ij} > 0$, and $D_{ij} = 0$ if $C_{ij} = 0$.) Therefore, the cost of the edges are decided based on D_{ij}, as shown in Figure 5. The figure also shows that the minimum-cost path found between A and E is using existing capacity on the edge A-E.

2) DCPI_AUG: The only difference between the MCPI_AUG and DCPI_AUG[9] is in the way costs are assigned to logical edges using the given information. The main idea of DCPI_AUG is to assign to each logical edge ij a cost that is inversely proportional to the total *potential* capacity between i and j, which is the sum of the existing logical capacity (R_{ij}) and the potential capacity available from the WDM layer (C_{ij}). In DCPI_AUG, ϕ_{ij}, the cost of a logical edge between LSR i and j, is defined as follows:

$$\phi_{ij} = \begin{cases} \infty & \text{if} \quad C_{ij} = 0, \quad \text{and } r_{ij}^w < b \quad \forall\, w \\ \frac{1}{C_{ij} + R_{ij}} & \text{otherwise.} \end{cases} \tag{5}$$

As shown in Fig. 4 and Fig. 6, the idea in DCPI_AUG is to try to avoid using those logical edges with smaller total potential capacity. By doing so, it spreads out the traffic, thereby reducing potential bottleneck edges on the network

[9]DCPI_AUG was called as CAPA_AUG in [13].

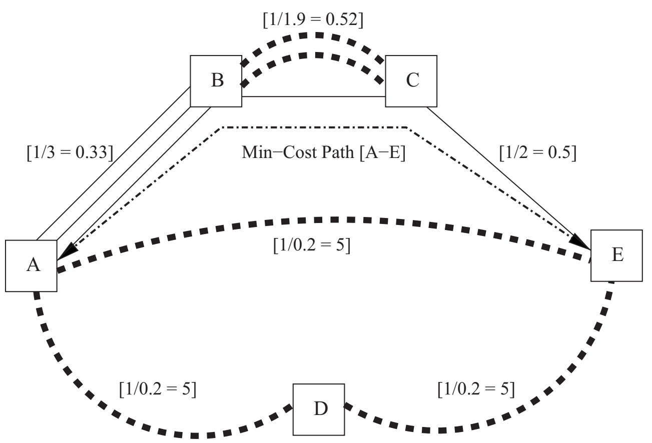

Fig. 6. For DCPI_AUG, cost of logical edges [1 / total capacity = ϕ_{ij}].

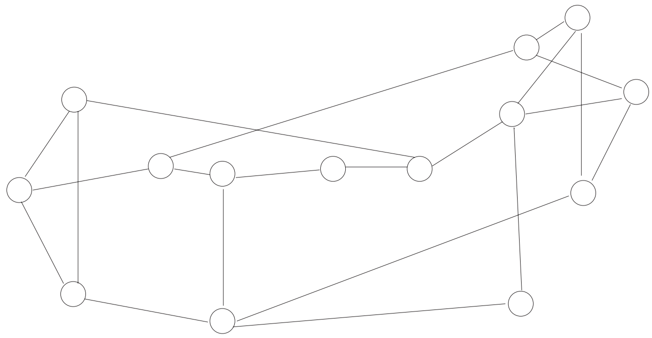

Fig. 7. NSF network topology.

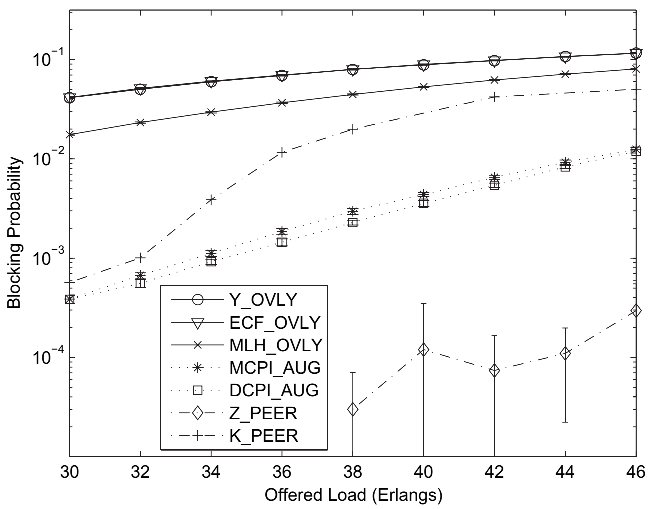

Fig. 8. Blocking probability vs. offered load with no port limit and 8 wavelengths.

for the future requests. We note that we attempt to reduce potential bottleneck edges at the expense of reducing network utilization in the DCPI_AUG algorithm.

Fig. 6 shows how to calculate the cost of each logical edge on the same example as in Fig. 4. It also shows the minimum cost path between A and E, which is A-B-C-E. For the logical edges A-B and C-E, DCPI_AUG opens a new lightpath (logical link). For the edge B-C, however, it chooses to use the logical link with 0.4 units of residual bandwidth.

Note that L_{ij}, and therefore C_{ij}, (respectively, B_{ij} and D_{ij}) depend on the RWA used in the WDM layer, and that it is not necessarily true that the WDM layer can establish L_{ij} (respectively, B_{ij}) lightpaths simultaneously for every LSR pair (i, j). It *is* true that there are enough wavelengths in the WDM layer to establish L_{ij} lightpaths between LSRs i and j, but the actual establishment of one or more lightpaths between i and j may decrease the number of lightpaths that can be established between some other pair of LSRs m and n. In other words, the entries of the matrix $L \stackrel{\text{def}}{=} [L_{ij}]$ (respectively, $B \stackrel{\text{def}}{=} [B_{ij}]$) are not independent of each other. Therefore, even if the DCPI_AUG (respectively, MCPI_AUG) algorithm produces a minimum-cost path that requires more than one new lightpath to be established, it may not be possible to open all of them at the same time, even if $C_{ij} > 0$ (respectively, $D_{ij} > 0$), for all the logical edges ij. In this case, DCPI_AUG (respectively, MCPI_AUG) will block the request.

In all the algorithms, a lightpath is released as soon as the last LSP that uses the lightpath departs.

VI. Performance Evaluation

We evaluate the performance of the proposed algorithms through extensive simulations. We use the NSF network as our WDM transport layer topology, which includes 14 OXCs without wavelength conversion capability and 21 bi-directional links as shown in Fig. 7. All our experiments were performed assuming 8 wavelengths per physical link. We assume that one LSR in the IP/MPLS layer is connected to each OXC in the WDM layer as shown in Fig. 1, and every LSR is an ingress/egress LSR. These are assumptions used for generating results and are not assumptions needed by the algorithms themselves. We assume that LSP requests arrive and depart according to a Poisson traffic model, and the ingress and egress LSRs for a request are uniformly distributed. We also assume that the bandwidth granularity of an LSP request is $\frac{1}{16}$, i.e., the normalized bandwidth of every LSP request is $\frac{i}{16}$ where i is an integer that is randomly (uniformly) chosen from the numbers $1, \ldots, 16$. We calculate the blocking probability (P_b) and the utilization (U) for each algorithm. All confidence intervals shown are 95% intervals. (The time complexity of K_PEER was too high to run a sufficient number of trials and get meaningful confidence intervals.) Note that in some cases the confidence intervals are smaller than the symbol used to show the data point, and may not be clearly visible.

A. No Port Limit

We first show the performance of the algorithms without any limit on the number of ports on the LSRs. We plot P_b and U in Fig. 8 and Fig. 9 against the network offered load measured in Erlangs. Without any limitation on the number of ports per LSR, Z_PEER shows the best performance, followed by the two AUG algorithms, then K_PEER, and then the OVLY algorithms. It is interesting to note that both AUG algorithms perform better than the K_PEER algorithm despite having access to more limited information than K_PEER. The overlay algorithms show the worst performance, as expected. We will see later that when the number of ports is limited, the relative

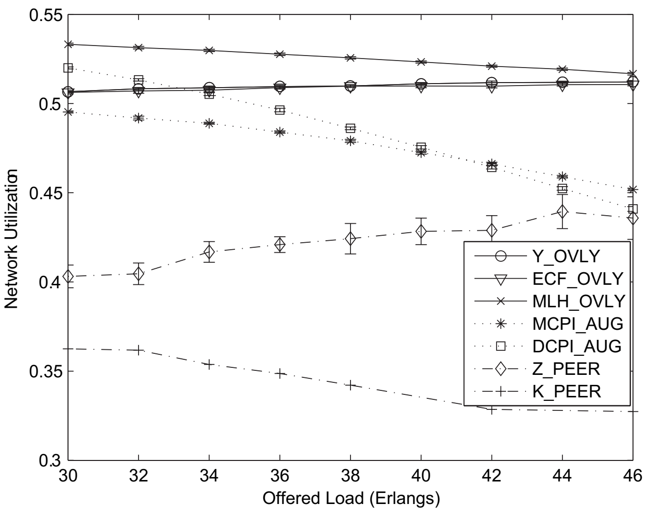

Fig. 9. Network utilization vs. offered load with no port limit and 8 wavelengths.

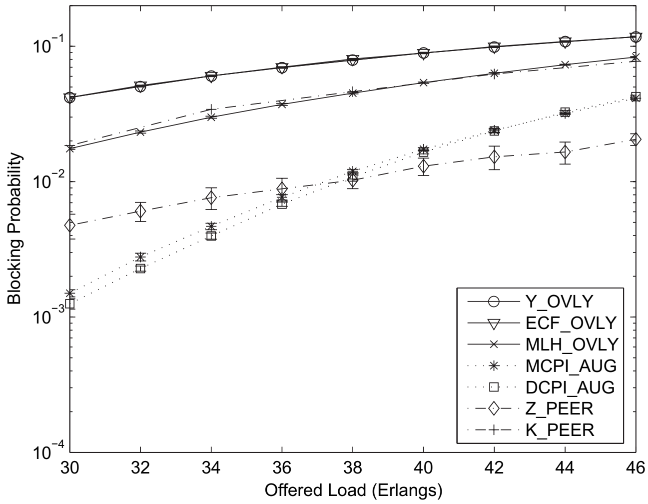

Fig. 10. Blocking probability vs. offered load with 8 wavelengths and 12 ports.

performance of K_PEER deteriorates further, and the AUG algorithms outperform even Z_PEER for some loads.

Note that one of the reasons for having the worst blocking probabilities in Y_OVLY, ECF_OVLY, and MLH_OVLY, higher than those of the other algorithms by more than an order of magnitude, is the fact that only one new lightpath is allowed to be opened per request in Y_OVLY, ECF_OVLY, and MLH_OVLY, whereas there is no such restriction assumed by the other algorithms. It is interesting to note that the very same restriction leads to high network utilizations for Y_OVLY, ECF_OVLY, and MLH_OVLY, approximately 10-15% more than the algorithms for the peer model, as shown in Fig. 9. They are restricted to open one lightpath per request, blocking some requests that would have been accepted with more than one new lightpath, which, in turn, would lower the network utilization. Among the three algorithms for the overlay model, however, MLH_OVLY shows as much as twice better P_b and 2-3% better U than Y_OVLY and ECF_OVLY, depending on the traffic load. This comes from the additional step in MLH_OVLY of checking other logical edges for a new lightpath.

The proposed MCPI_AUG and DCPI_AUG algorithms in the augmented model provide much lower P_b (by more than an order of magnitude) than Y_OVLY, ECF_OVLY and MLH_OVLY, but does not perform as well as Z_PEER. However, they also achieve approximately 5-10% higher network utilization than Z_PEER at low loads as shown in Fig. 9. This is because of its effort to use existing logical capacity whenever it finds enough capacity for a logical edge on the chosen minimum-cost path.

Between the two algorithms for the augmented model, DCPI_AUG shows a slight improvement over MCPI_AUG. As mentioned earlier, this is because more information from the WDM layer is utilized in DCPI_AUG, but the extra updates and information exchanges may not be worth this slight blocking improvement.

We also note the following interesting phenomenon. As the traffic load increases, the utilization for the AUG algorithms decrease faster than for the other algorithms, while as expected the blocking probability increases for all algorithms. The reason is due to the inherent nature of the AUG algorithms. More lightpaths are opened by MCPI_AUG and DCPI_AUG (relative to the other algorithms) in an effort to keep the blocking probability low. As a result, network utilization drops somewhat faster than for other algorithms. Despite this, note that the blocking probabilities for the AUG algorithms are better than the P_b for the K_PEER algorithm which has a far lower utilization as well. The P_b for Z_PEER is better but the utilization for Z_PEER is correspondingly lower.

An important observation we made throughout this performance evaluation is that for the dynamic LSP routing problem, there exists a trade-off between blocking probability and network utilization. Y_OVLY, ECF_OVLY, and MLH_OVLY in the overlay model, for example, tend to open a new lightpath more conservatively than the other algorithms in the augmented and peer models. It is clearly shown that such a conservative approach leads to higher network utilization at the cost of higher blocking probability. This observation also implies that one can find an optimum point on the trade-off between P_b and U for given performance goals in a network model.

B. Limited Number of Ports

As mentioned earlier, we believe that the high cost of ports and the fact that a large majority of traffic is pass-through, would lead to a limit on the number of ports, in practice. In order to study the effects of a limited number of ports on the performance of the algorithms, we now plot P_b and U against network offered load for the case of 8 wavelengths per physical link and 12 ports[10] per LSR in Fig. 10 and Fig. 11. Fig. 10 shows a relatively larger increase in P_b for K_PEER and Z_PEER than for MCPI_AUG and DCPI_AUG, compared to the case with no port limit in Fig 8.

Notice that there is no change in the blocking performances of Y_OVLY, ECF_OVLY, and MLH_OVLY from Fig.

[10]We assume that a lightpath on *any* wavelength may be terminated at a given port.

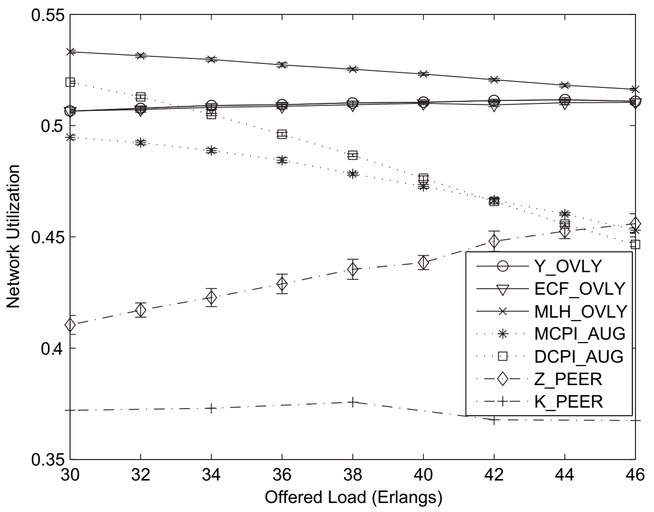

Fig. 11. Network utilization vs. offered load with 8 wavelengths and 12 ports.

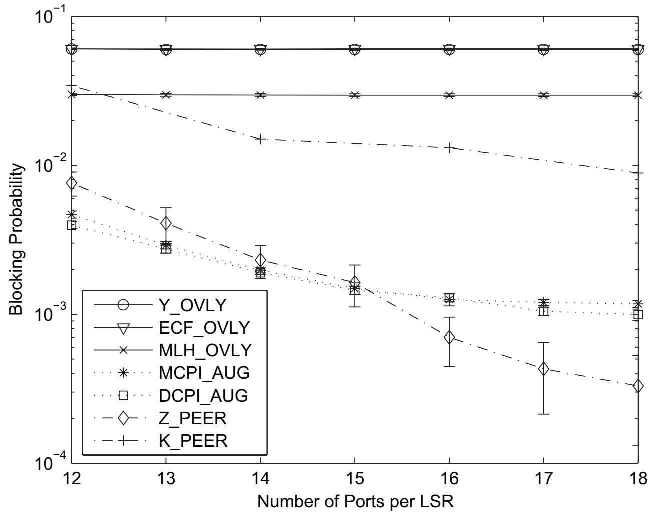

Fig. 12. Blocking probability vs. number of ports with 8 wavelengths, offered load = 34.

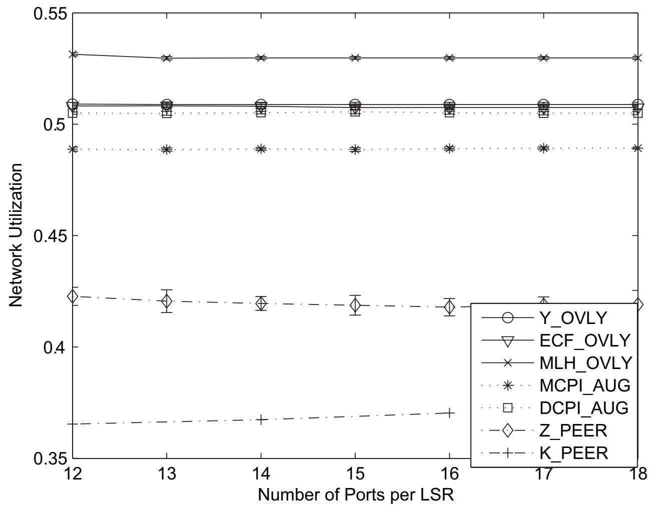

Fig. 13. Network utilization vs. number of ports with 8 wavelengths, offered load = 34.

8 and Fig. 9. This is because the performances of Y_OVLY, ECF_OVLY, and MLH_OVLY are limited by the number of wavelengths per link and are not port-limited.

A very important observation from Fig. 10 is that the AUG algorithms achieve a P_b that is about half that achieved by the Z_PEER algorithm at low loads, and continue to outperform the K_PEER algorithm at all loads. Furthermore, the utilization levels achievable by the AUG algorithms are higher than those achieved by Z_PEER and K_PEER for all loads, except at very high loads where Z_PEER has a slightly higher utilization (see Fig. 11) and a lower P_b as well.

Considering the complexity of network management information to be exchanged between the network layers in the different models, it is significant that simple algorithms such as MCPI_AUG and DCPI_AUG designed for the augmented model can outperform the peer model algorithms K_PEER and Z_PEER. While an algorithm for the peer model can always be made to outperform a corresponding one for an augmented model (as the peer model algorithm has more information at its disposal than any algorithm for the augmented model), the above observation points out the importance of using the available information in a "correct" manner. This also shows the potential of an augmented model as a practical solution that could benefit from the advantages of both the peer and overlay models.

C. *Effect of Number of Ports*

We next plot P_b and U against the number of ports per LSR in Fig. 12 and Fig. 13 to show how the performance improves as the number of ports increases. With 8 wavelengths on a physical link and a network load of 34 Erlangs, it is shown that Z_PEER's performance is the most sensitive to port limits. As the number of ports per LSR increases, P_b for Z_PEER drops faster than for any other algorithm, whereas in Y_OVLY, ECF_OVLY, and MLH_OVLY, there is no perceptible change.

As the number of ports increases, MCPI_AUG and DCPI_AUG also improve their P_b but not as much as Z_PEER or K_PEER. As we already showed in Fig. 8, if there is no limit on the number of ports, then Z_PEER outperforms MCPI_AUG and DCPI_AUG.

VII. Conclusions

In this paper, we studied the dynamic LSP routing problem for three different network models of IP/MPLS over WDM networks. While there have been algorithms proposed for the peer and overlay models, there has been no algorithm for the augmented model. We proposed two simple algorithms for the augmented model that achieve very good blocking performance and network utilization when compared to the peer and overlay model algorithms available in the literature. We also proposed two new algorithms for the overlay model, one of which performs better than the algorithm in [1], in blocking probability by 30-40% and network utilization by 3-5%.

While previous work assumed that the network is wavelength-limited in performance, we also considered a port-limited case which is likely to occur in practice due to the high cost of ports. We presented a comprehensive study of the performance of algorithms for the three network models for the first time, and showed the trade-off between network utilization and blocking probability.

An interesting observation we made was that our proposed augmented model algorithm outperformed both the peer model algorithms in terms of blocking probability over a wide range of network loads. We also showed that depending on the type/amount of network information exchanged between the two layers, the performance of algorithms utilizing the information may vary. Nevertheless, we showed that it is possible to achieve very good performance using very limited information in the augmented model, when compared to the large amount of information that must be flooded across the network in the peer model. While we have presented algorithms for the augmented model that shows good performance, more work remains to be done in quantifying the amount of network information that the various models use. Moreover, we have presented algorithms for setting up service paths only. Developing algorithms for routing LSPs when protection is taken into account would be an interesting future course of work.

ACKNOWLEDGMENTS

The authors would like to thank Mr. Onur Turkcu for his help in running some simulations.

REFERENCES

[1] Y. Ye, C. Assi, S. Dixit, and M. A. Ali, “A simple dynamic integrated provisioning/protection scheme in IP over WDM networks,” *IEEE Commun. Mag.*, vol. 39, no. 11, pp. 174–182, Nov. 2001.

[2] M. Kodialam and T. V. Lakshman, “Integrated dynamic IP and wavelength routing in IP over WDM networks,” in *Proc. INFOCOM 2001*, vol. 1, pp. 358–366.

[3] Q. Zheng and G. Mohan, “An efficient dynamic protection scheme in integrated IP/WDM networks,” in *Proc. ICC 2003*, vol. 2, pp. 1494–1498.

[4] E. Rosen, A. Viswanathan, and R. Callon, “Multiprotocol label switching architecture,” RFC 3031, Jan. 2001.

[5] P. Ashwood-Smith *et al.*, “Generalize multi-protocol label switching (GMPLS) architecture,” RFC 3945, Oct. 2004.

[6] “OFI-User Network Interface (UNI) 1.0 signaling specification,” Oct. 2001.

[7] S. Baroni, J. O. Eaves, M. Kumar, M. K. Qureshi, A. Rodriguez-Moral, and D. Sugerman, “Analysis and design of backbone architecture alternatives for IP optical networking,” *IEEE J. Sel. Areas Commun.*, vol. 18, no. 10, pp. 1980–1994, Oct. 2000.

[8] A. Rodriguez-Moral, P. Bonenfant, S. Baroni, and R. Wu, “Optical data networking: Protocols, technologies, and architectures for next generation optical transport networks and optical internetworks,” *IEEE/OSA J. Lightw. Technol.*, vol. 18, no. 12, pp. 1855–1870, Dec. 2000.

[9] J. Y. Wei, “Advances in the management and control of optical Internet,” *IEEE J. Sel. Areas Commun.*, vol. 20, no. 4, pp. 768–785, May 2002.

[10] S. Koo, G. Sahin, and S. Subramaniam, “Cost efficient LSP protection in IP/MPLS-over-WDM overlay networks,” in *Proc. ICC 2003*, vol. 2, pp. 1278–1282.

[11] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, 1993.

[12] L. Berger *et al.*, “Generalized MPLS signaling - RSVP-TE extensions,” RFC 3473, Jan. 2003.

[13] S. Koo, G. Sahin, and S. Subramaniam, “Dynamic LSP provisioning in overlay, augmented, and peer architectures for IP/MPLS over WDM networks,” in *Proc. INFOCOM 2004*.

Sunggy Koo is a Product Manager at Spirent Communications. He received the B.S. degree in international business from Korea University, Seoul, Korea in 1996, the M.S. degree in information systems technology and the Ph.D. degree in electrical and computer engineering, in 1998 and 2004 respectively, both from the George Washington University, Washington, DC. Dr. Koo has held various engineering and product marketing positions at Isocore and Spirent Communications. His research interests include IP-Optical routing, triple play networks, IPTV, MPLS/GMPLS, network QoS measurement, multicast, mobile IP, fixed-mobile convergence, IMS and WiMax.

Gokhan Sahin is an assistant professor of Electrical and Computer Engineering at Miami University. He received the B.S. degree from Bilkent University, Ankara, Turkey in 1996, and the M.S.E.E. and Ph.D. degrees from the University of Washington, Seattle, in 1997 and 2001, respectively, all in Electrical Engineering. He was a Wayne C. and Grace M. Stanley fellow at the University of Washington. Dr. Sahin has held various research positions in both academia and the industry, at AT&T Labs Research, Corvis Corporation, and GWU. Prior to joining Miami University, he was an assistant research professor with the Computer Science and Engineering Department at University of Nebraska. His research interests are in the design and analysis of computer and communication networks, and in particular in the areas of wireless and IP/optical networking, where he has published extensively on network survivability, GMPLS, and resource allocation problems.

Suresh Subramaniam (S ’95, M ’97) received the Ph.D. degree in electrical engineering from the University of Washington, Seattle, in 1997. He is an Associate Professor in the Department of Electrical and Computer Engineering at the George Washington University, Washington, DC. His research interests are in the architectural, algorithmic, and performance aspects of communication networks, with particular emphasis on optical and wireless ad hoc networks.

Dr. Subramaniam is a co-editor of the books *Optical WDM Networks - Principles and Practice* and *Emerging Optical Network Technologies: Architectures, Protocols, and Performance*. He has been on the program committees of several conferences including Infocom, ICC, Globecom, and Broadnets, and served as TPC Co-Chair for the optical networks symposium at Broadnets 2004. He is currently serving as General Co-Chair for Broadnets 2006, and as TPC Co-Chair for the optical networks symposia at Globecom 2006 and ICC 2007. He is currently on the editorial boards of *KICS Journal of Communications and Networks* and *IEEE Communications Surveys and Tutorials*. He is a co-recipient of Best Paper Awards at the ICC 2006 Symposium on Optical Systems and Networks, and at the 1997 SPIE Conference on All-Optical Communication Systems.