

University of Nebraska - Lincoln

DigitalCommons@University of Nebraska - Lincoln

CSE Conference and Workshop Papers

Computer Science and Engineering, Department
of

2003

Inter-Domain Dynamic Routing in Multi-Layer Optical Transport Networks

Xi Yang

University of Nebraska-Lincoln, xyang@cse.unl.edu

Byrav Ramamurthy

University of Nebraska-Lincoln, bramamurthy2@unl.edu

Follow this and additional works at: <https://digitalcommons.unl.edu/cseconfwork>



Part of the [Computer Sciences Commons](#)

Yang, Xi and Ramamurthy, Byrav, "Inter-Domain Dynamic Routing in Multi-Layer Optical Transport Networks" (2003). *CSE Conference and Workshop Papers*. 106.

<https://digitalcommons.unl.edu/cseconfwork/106>

This Article is brought to you for free and open access by the Computer Science and Engineering, Department of at DigitalCommons@University of Nebraska - Lincoln. It has been accepted for inclusion in CSE Conference and Workshop Papers by an authorized administrator of DigitalCommons@University of Nebraska - Lincoln.

Inter-Domain Dynamic Routing in Multi-Layer Optical Transport Networks

Xi Yang and Byrav Ramamurthy

Department of Computer Science and Engineering, University of Nebraska-Lincoln
115 Ferguson Hall, UNL, Lincoln, NE 68588-0115 Email: {xyang, byrav}@cse.unl.edu

Abstract – Next-generation optical transport networks will automatically and dynamically provision end-to-end connections. In this paper, we study the problem of inter-domain dynamic routing under a multi-layer multi-domain network model, which allows the end-to-end connections to be set up not only across multiple routing domains but also through two transport layers: the optical layer and the digital layer. In this model, a connection can traverse the domain boundary either through optical bypass or through optical-electrical-optical (O/E/O) processing. We propose an inter-domain dynamic routing scheme with modest time complexity to address the problem from an algorithmic perspective.

1. INTRODUCTION

Optical transport network technologies have evolved from the earliest point-to-point transport technology which was designed for increasing link capacity. Although today's transport networks are more advanced and capable of switching on the digital path level, e.g. STM-xx or wavelength level, the routing of traffic still relies on the upper layer, e.g. the IP layer. The transport network is viewed as a static under-layer that can only be reconfigured by carriers or service providers at a slow pace. With the advent of new network applications and services, end users are demanding high-bandwidth end-to-end connections that have to be provisioned automatically and dynamically. This gives rise to the problem of dynamic routing in optical transport networks.

As a standard for next-generation optical transport networks, the Automatic Switched Transport Network (ASTN)/Automatic Switched Optical Network (ASON) architecture was proposed by ITU-T to meet the demand for dynamic provisioning. According to the ITU-T recommendations [1][2], ASTN/ASON has two transport layers: the optical layer and the digital layer. Because long-haul optical-layer transmission incurs severe physical impairments, even a wavelength-level end-to-end connection has to resort to O/E/O at some intermediate nodes [3]. O/E/O is also used to convert a wavelength when wavelength continuity cannot be satisfied due to contention. Therefore, an end-to-end connection can be set up through multiple consecutive digital paths that terminate at the O/E/O points. A common example of a digital path is an SONET/SDH path.

An optical transport network also contains multiple domains because different carriers or service providers may want to manage their own parts of the network. Another reason is that different parts of the transport network may use different technologies. The Internet today is partitioned into a great number of routing domains called Autonomous Systems (AS). In the Internet, a variety of Interior Gateway Protocols (IGP) are used to route IP traffic inside each domain, while

Border Gateway Protocol (BGP) is used to route traffic across multiple domains. This solution is very flexible because inter-domain routing is separated from diverse local routing schemes and policies. Automatic connectivity verification and dynamic update of routing tables make this solution scalable with respect to change of network topology. With the same concerns on flexibility and scalability, an inter-domain dynamic routing solution like that used in the Internet can be adopted in the multi-domain optical transport networks. In [4], the requirements and guidelines for interconnection of optical networks with multiple domains were described. The multi-segment wavelength routing problem addressed in [5] is very similar to the optical-layer inter-domain routing problem. In [6], a BGP/GMPLS solution was proposed (GMPLS represents Generalized Multi-Protocol Label Switching [7]). In this solution, the GMPLS extension of BGP was addressed from a protocol perspective, dealing with address, message, format etc.

In this paper, we address the problem of inter-domain routing in optical transport networks from an algorithmic perspective. Our network model facilitates the inter-domain routing not only among multiple domains but also through multiple transport layers. In such a network model, some special concerns will be taken into consideration. The first is on wavelength continuity at the optical layer. In the Internet, BGP can decide the reachability of an entire domain or AS when one of its border routers is reachable. However, in an optical transport network, due to the constraint on wavelength continuity, the nodes inside the same domain do not share the same reachability from other domains, which adds to the complexity of routing computation. Second, compatible types of digital format should be supported by two consecutive digital paths at the digital layer. (Two types of digital format are compatible if they can be converted to each other at an O/E/O point.) Third, end-to-end transmission quality should be guaranteed. This imposes a constraint on the summation of bit error rate (BER) of all digital paths on the end-to-end connection. BER has been proposed as a link property in the IETF drafts on the SONET/SDH and WDM enhancements [8][9] of the Link Management Protocol (LMP) [10].

The rest of this paper is organized as follows. In Section 2, we describe the model of the multi-layer multi-domain optical transport networks. In Section 3, we propose an inter-domain dynamic routing scheme. In Section 4, experimental results are presented. We conclude this paper in Section 5.

2. NETWORK MODEL

The network model is a combination of multiple vertical layers and multiple horizontal domains, as shown in Figure 1.

Multi-Layer: This model follows the two-layer ITU-T ASTN/ASON architecture. An end-to-end connection can be

* This work was supported by NSF grants (ANI-0074121 and EPS-0091900).

set up through multiple digital paths at the digital layer, while each digital path can be set up through multiple optical links at the optical layer. The optical layer consists of wavelength switching capable optical cross-connects (OXC) and optical links that employ the WDM technology. On top of the optical layer, the digital layer can be either the existing digital path layer, e.g., the SONET/SDH layer, or the proposed digital wrapper layer [11]. At this layer, digital paths are provisioned between electronic interfaces. The electronic interfaces on both ends of a digital path support the same type of digital format, which is a combination of digital rate and frame format as defined in [6] and [11]. In addition, to add/drop digital paths to/from the optical layer, E/O and O/E resources, i.e., transmitters and receivers, will be used.

Multi-Domain: In this paper, we define a special category of boundary nodes called *domain gateways*. A domain gateway connects two adjacent domains. It can support both O/E/O and optical bypass at the boundary. Through O/E/O, a domain gateway is capable of conversion between optical signals of different wavelengths and/or different types of digital format. The entire optical transport network is partitioned into multiple domains by such domain gateways. Each remainder node becomes an *interior node* in one of the domains. The interior nodes may also support both O/E/O and optical bypass.

Routing: Routes crossing the domain boundary are computed at the domain gateways, which behave like border routers in the Internet. From the viewpoint of inter-domain routing, a local route within the boundary of each domain is a *routing hop*. Local routes are obtained by local routing schemes (LRS). The domain gateways maintain both local and next-hop routing information and route the end-to-end connections on a hop-by-hop basis.

Because a routing hop does not necessarily coincide with a digital path, the combination of *Multi-Layer* and *Multi-Domain* results in complicated inter-domain routing computation. In other words, a digital path can either terminate at the boundary of a domain or traverse the boundary through optical bypass. When an end-to-end connection is routed through multiple domains, each of the underlying digital paths can be set up between different types of nodes as in the following cases (see Figure 1).

- Between two interior nodes in two different domains.
- Between two domain gateways in two different domains.
- Between a domain gateway and an interior node in two different domains.
- Between two interior nodes within the same domain.
- Between two domain gateways within the same domain.
- Between a domain gateway and an interior node within the same domain.

Because different resources are allocated under different constraints are applied in different cases, these cases must be separately considered for inter-domain routing computation.

Cases *a*, *b* and *c*: A digital path spans over two consecutive routing hops, i.e., across two adjacent domains, through optical bypass. Wavelength continuity must be guaranteed for inter-domain routing at the domain boundary (in the absence of all-optical wavelength conversion).

Cases *b*, *c*, *e* and *f*: If the digital path is terminated at the boundary and a next digital path is started in a next domain, O/E, E/O and electronic resources are used and compatible types of digital format on these two digital paths must be guaranteed.

For all these cases, BER should be computed at the terminating point of each digital path. At the boundary of a domain, for each through end-to-end connection, the sum of the BER values of all the digital paths terminating in the corresponding routing hop is computed and used to constrain the transmission quality of this connection.

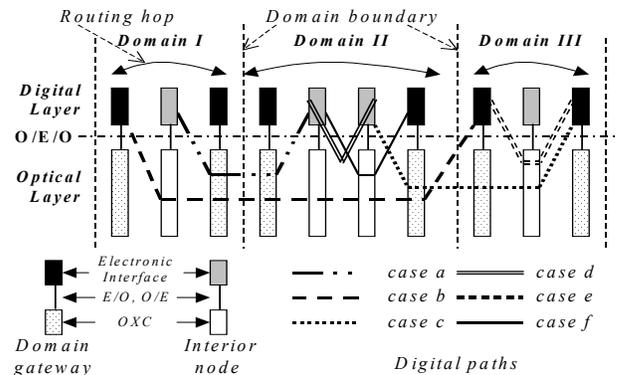


Figure 1: Multi-layer multi-domain optical transport network model.

3. INTER-DOMAIN DYNAMIC ROUTING SCHEME

In this section, we propose a dynamic routing scheme, which distributes inter-domain routing computation to domain gateways. The routing computation at the domain gateways is further divided into three functions. First, a domain gateway uses an LRS to compute *alternate local routes* between itself and each of the interior nodes as well as neighboring domain gateways. Second, a next-hop computation function is used to join the alternate local routes of this domain to the alternate routes of adjacent domains to form the *next-hop interfaces* leading to desired destinations. Finally, the hop-by-hop path selection function uses the obtained local and next-hop routing information to establish inter-domain end-to-end connections.

3.1 Computing Alternate Local Routes

Between a domain gateway and each of the interior nodes and neighboring domain gateways, a set of alternate local routes are computed in both directions. In each direction, up to k alternate local routes are computed by the LRS using different wavelengths and/or different types of digital format. We need alternate local routes because we cannot predetermine which local wavelengths or digital format types will be used for an arbitrary end-to-end connection to satisfy the constraints from previous and next routing hops. For every alternate local route, the domain gateway maintains complete routing information, including digital paths, types of digital format, optical links and wavelengths.

3.2 Computing Next Hops

A. Formulation

By U we denote the identifier of a domain gateway, which maintains the next-hop interfaces for every destination node it

can reach. We define the data structure of a next-hop interface as a data entry: $[T, D, r(U, T, D)]$, where T is the identifier of the destination node, D is the identifier of the next-hop node on the route, and $r(U, T, D) = [c_1, c_2, \dots]$ is an array that contains the next-hop routing information for a set of up to k alternate routes from the domain gateway U to the destination T via the next-hop node D . Each item of the next-hop routing information in $r(U, T, D)$ is defined as a sextuplet $\{w_U, w_D, t_U, t_D, h_{U-D}, b_{U-D}\}$. The first two elements represent that optical bypass is carried out at U and D using the wavelengths w_U and w_D respectively. If O/E and/or E/O instead of optical bypass is carried out at U and/or D , $w_U = null$ and/or $w_D = null$. t_U and t_D are the types of digital format after leaving U and before entering D . The last two elements represent costs of the route from U to T . h_{U-D} is the cost using *digital-layer distance* as metric, which is defined as the number of underlying digital paths. b_{U-D} is the cost using BER as metric, which equals to the sum of BER values collected from all O/E/O points between U and T . We call the sextuplet $\{w_U, w_D, t_U, t_D, h_{U-D}, b_{U-D}\}$ a *route component* and call the array $r(U, T, D)$ a *route vector*. Each route vector consists of a variable number of route components, which correspond to a variable number of alternate routes from the domain gateway U to the destination T via the next-hop node D .

A local route vector $r(U, D, D)$ of U contains the routing information of a set of local alternate routes from the domain gateway U to a neighboring domain gateway D , while the neighboring gateway has a route vector $r(D, T, E)$ that contains the routing information of a set of alternate routes leading to the destination T via its next-hop node E . Therefore, the problem of finding the next-hop interface for a desired destination T can be formulated as joining the route vector $r(U, D, D)$ to the route vector $r(D, T, E)$ to form a new route vector $r(U, T, D)$, which will contain routing information of a set of alternate routes from U to T via the next-hop node D .

B. Algorithms

By joining $r(U, D, D)$ to $r(D, T, E)$ to form a new route vector $r(U, T, D)$, a set of alternate local routes from U to D , represented by the route components in $r(U, D, D)$, is joined to a set of alternate routes from D to T , represented by the route components in $r(D, T, E)$, to obtain a set of alternate routes, represented by the joined route components in $r(U, T, D)$. Note that an arbitrary component in $r(U, D, D)$ cannot be joined to an arbitrary component in $r(D, T, E)$ unless the constraints on wavelength continuity and compatibility of types of digital format are satisfied. In addition, the two route costs, digital-layer distance and BER, in joined components should be added up and used to constrain the end-to-end quality of transmission. We use the following procedure to join two route components.

Route Component Joining (RCJ) Procedure:

Input: A route component $c_i = \{w_U, w_D, t_U, t_D, h_{U-D}, b_{U-D}\}$ and a route component $c_j = \{w_D, w_E, t_D, t_E, h_{D-T}, b_{D-T}\}$, where $c_i \in r(U, D, D)$ and $c_j \in r(D, T, E)$.

Output: A route component $c_l = \{w_U, w_D, t_U, t_D, h_{U-T}, b_{U-T}\}$ in $r(U, T, D)$ or null.

Procedure:

1. If t_{D1} of c_i is not compatible with t_{D2} of c_j , return *null*.
2. If w_{D1} of c_i and w_{D2} of c_j are equal (*null* or not *null*), $h_{U-T} = h_{U-D} + h_{D-T}$, $b_{U-T} = b_{U-D} + b_{D-T}$ and return $c_l = \{w_U, w_D, t_U, t_D, h_{U-T}, b_{U-T}\}$.
3. Otherwise, return *null*.

Joining of the two route vectors $r(U, D, D)$ and $r(D, T, E)$ is realized by selectively joining their components. We use the following algorithm to join the route vectors to obtain next-hop routing information for destinations in other domains.

Route Vector Joining (RVJ) Algorithm:

Input: A local route vector $r(U, D, D) = [c_1 \dots c_i \dots]$ and a received route vector $r(D, T, E) = [c_1' \dots c_j \dots]$.

Output: A new route vector $r(U, T, D) = [c_1 \dots c_l \dots]$.

Algorithm:

1. For each component c_i in $r(U, D, D)$
2. For each component c_j in $r(D, T, E)$
3. Use the RCJ procedure to join c_i and c_j .
4. Let $c_l = \text{RCJ}(c_i, c_j)$.
5. If c_l is NOT *null*, put c_l into $r(D, T, E)$.
6. End For
7. End For

The computational complexity of CVJ is $O(k^2)$, where k is the maximum number of route components in each route vector.

In practice, a route vector for a specific destination may be updated from time to time due to dynamic change of network topology and link states. When the domain gateway U receives a new route vector $r(B, T, F)$ from a neighboring domain gateway B , it will use the RVJ algorithm to join another local route vector $r(U, B, B)$ to $r(B, T, F)$ to form a new route vector $r(U, T, B)$, which represents another set of alternate routes from the domain gateway U to the destination T via an alternate next-hop node B . Then, the inter-domain routing scheme needs to compare $r(U, T, B)$ with $r(U, T, D)$ and decide which is to be used as the next-hop interface. The comparison is based on one of the two cost metrics, digital-layer distance and BER. We use the following algorithm to compare and update the route vectors.

Route Vector Comparison and Update (RVCU) Algorithm:

Input: $r(U, T, D)$, $r(U, T, B)$, and the cost metric (e.g. BER).

Output: A decision on which route vector is to be chosen.

Algorithm:

1. If using BER as the cost metric, find the route component with the lowest BER cost b_1 in $r(U, T, D)$ and the route component with the lowest BER cost b_2 in $r(U, T, B)$; if using digital-layer distance as the cost metric, find the route component with the lowest h_1 in $r(U, T, D)$ and the route component with the lowest h_2 in $r(U, T, B)$.
2. If using BER as the cost metric and $b_1 < b_2$, choose $r(U, T, D)$; if $b_1 > b_2$ choose $r(U, T, B)$; otherwise, $b_1 = b_2$, go to step 4.
3. If using digital-layer distance as the cost metric and $h_1 < h_2$, choose $r(U, T, D)$; if $h_1 > h_2$, choose $r(U, T, B)$; otherwise, $h_1 = h_2$, go to step 4.

4. If either of the two route vectors has run out of route components, choose the other; if both route vectors have all their route components compared, choose $r(U, T, D)$; otherwise, get the components with the next lowest costs in respective route vectors and repeat 2 or 3.

The computational complexity of RVCU is $O(k)$. By exchanging the route vectors between neighboring domain gateways, global information of network topology and link states can be dynamically disseminated to all domains. Each domain gateway uses the RVJ algorithm to join those route vectors to obtain the next-hop routing information for an arbitrary destination. Also, the domain gateway uses the RVCU algorithm to update the next-hop interfaces to respond to the change of network topology and link states.

3.3 Hop-by-Hop Path Selection

An end-to-end connection request from a source node S consists of the destination node identifier T , the initial type of digital format x and the initial wavelength set $Set_W(S)$. Through the next-hop interfaces, the connection request can be forwarded to the correct next-hop domain gateways until the destination is reached. Through the alternate local routes maintained by the domain gateways, the complete routing information of each routing hop can be extracted and form the complete end-to-end routing information. The procedure of hop-by-hop selection is described as follows.

1. If S and T are in the same domain, or S and T are neighboring domain gateways, extract routing information of all alternate local routes from S to T and go to step 5.

2. If S is a domain gateway, search for the next-hop interface $[T, D, r(S, T, D)]$. If no such next-hop interface exists, reject the request; otherwise, extract complete routing information of all alternate local routes from S to the next-hop node D , each with an initial type of digital format x and an initial wavelength in $Set_W(S)$. The local routing information and the route vector $r(S, T, D)$ are sent together with the connection request to the next-hop node D . Then go to step 4.

3. If S is not a domain gateway, S forwards the connection request to a domain gateway U in the same domain. U generates an initial route vector $r(S, T, U)$ for S with route components in form of $\{null, w_0, t, h_0, b_0\}$, where w_0 is the wavelength of the last optical link, h_0 is the number of digital paths, and b_0 is the sum of BER on the corresponding alternate local route from S to U that starts with a wavelength in $Set_W(S)$ and with the digital format type x . U searches for the next-hop interface $[T, D, r(U, T, D)]$. If no such next-hop interface exists, the connection request is rejected; otherwise, $r(S, T, U)$ is joined to $r(U, T, D)$ to form a new route vector $r(S, T, D)$. Those alternate local routes, whose corresponding route components cannot be joined, are discarded. The complete routing information of the joined alternate routes from S to D and the route vector $r(S, T, D)$ are sent together with the connection request to the next hop D .

4. When the next-hop domain gateway D receives a connection request forwarded from another domain gateway, it searches for the next-hop interface $[T, E, r(D, T, E)]$. If the next-hop interface does not exist, reject the request; otherwise,

join the received $r(S, T, D)$ to $r(D, T, E)$. Those local routes, whose corresponding route components cannot be joined, are discarded. The complete routing information of joined alternate routes from S to E and the route vector $r(D, T, E)$ are sent together with the connection request to the next hop E .

5. Step 4 is repeated until the destination T is reached. Finally, a complete sequence of optical links and digital paths, with corresponding wavelengths and types of digital format, are obtained for each alternate route from S to T . If more than one route is obtained, the inter-domain routing scheme selects the one with appropriate end-to-end transmission quality, e.g., BER is less than 10^{-12} , and with the minimum cost, and returns the complete routing information in a routing confirmation message.

The hop-by-hop selection goes through at most M hops, where M denotes the maximum number of domains in the network. In each hop, it needs to search the routing table that contains up to N route vectors and execute the RVJ algorithm, which take $O(\log N)$ and $O(k)$ time respectively, where N denotes the number of nodes in the network. Therefore, the hop-by-hop selection has a computational complexity of $O(M \cdot (k + \log N))$, which is equal to $O(M \cdot k)$ if $k \geq \log N$.

Note that the route vectors in next-hop computation are initiated from the destinations and then joined by the local route vectors at intermediate domain gateways to form next-hop interfaces until reaching the sources. The procedure of hop-by-hop selection follows the reverse direction. The source nodes use the obtained next-hop interfaces to explore the routes leading to the destination nodes by joining the alternate local routes of a domain to the alternate local routes of the next domain until reaching the destination nodes. With the already obtained alternate local routes and next-hop interfaces, the hop-by-hop path selection can be carried out "on-the-fly."

4. EXPERIMENTAL RESULTS

In this section, we conduct simulation experiments in a 44-node 68-link mesh optical transport network (shown in Figure 2) to study the characteristics and performance of the proposed inter-domain dynamic routing scheme. Every node in the network consists of a 64×64 OXC and an access station following the node model described in [3]. Each optical link supports 16 wavelengths in both directions. The length of a link ranges from 50 km to 300 km, with amplifiers placed every 50 km. The domain gateways partition the network into 4 domains. Each of the remainder nodes becomes an interior node in one of the domains. Every interior node is assigned 4 pairs of transmitters and receivers, 4 STM-16 and 4 STM-64 electronic interfaces. Every domain gateway has twice as many numbers of transmitters, receivers and electronic interfaces as an interior node. We generate uniformly distributed traffic between every pair of nodes in the network. The traffic is converted into wavelength-level end-to-end connection requests. Such requests arrive in Poisson distribution with exponential holding times. Through changing the ratio of arrival interval to holding interval, the total amount of traffic is increased from one Erlang to a value high enough to cause severe blocking in the network. Two types of digital format,

STM-16 and STM-64, are randomly carried on these connections in a half-to-half proportion. Note that these two types of digital format are not compatible.

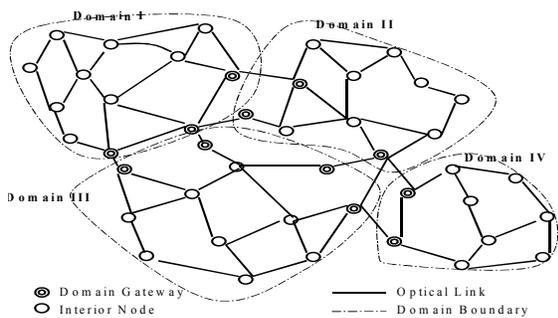


Figure 2: Topology of a 44-node 68-link mesh network.

The LRS uses the fixed alternate shortest path first routing algorithm and the first fit wavelength assignment algorithm. O/E/O points inside a domain are determined using the method described in [3]. The BER values are estimated online using the models and methods proposed in [12] and the system parameters in [3]. The end-to-end BER is constrained to less than 10^{-12} . In each experiment, we generate 100,000 connection requests and measure the blocking probability.

Figure 3 shows the blocking probability under different cost metrics. We have proposed to use digital-layer distance and BER as cost metrics. In addition, a mixed metric can be obtained by using digital-layer distance metric first and using BER metric only if the digital-layer distance of two route components is equal. The results indicate that digital-layer distance is more preferable than BER as a cost metric. In particular, using the mixed cost metric that first considers the digital-layer distance results in lowest blocking probability.

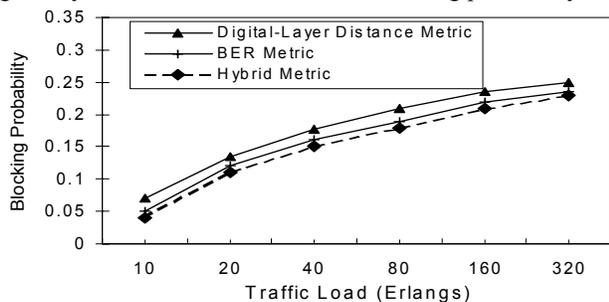


Figure 3: Blocking probability under different cost metrics.

Figure 4 shows the blocking probability under the three different network models. The proposed *Multi-Layer* model allowing for both O/E/O and optical bypass at the domain boundary. We can reduce it into a *Digital-Layer* model by enforcing O/E/O at the boundary so that the constraint on wavelength continuity is removed and only digital-layer routing is considered. We can also obtain an *Optical-Layer* model by enforcing optical bypass at the boundary so that inter-domain routing is only carried out at the optical layer. The latter two models can be realized by slightly changing the route component joining conditions in the RCJ procedure. The mixed cost metric is used for routing computation. The results show that under all traffic loads the *Multi-Layer* model

outperforms other models that only consider one of the two transport layers. Our proposed inter-domain dynamic routing scheme contributes to such performance improvement by being aware of resources as well as constraints at both layers.

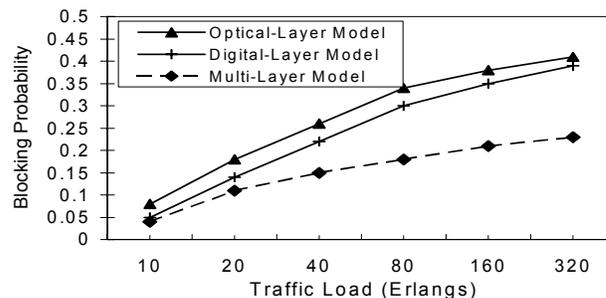


Figure 4: Blocking probability under different network models.

5. CONCLUSION

In this paper, we addressed the problem of inter-domain routing in the next-generation optical transport networks from an algorithmic perspective. The multi-layer multi-domain network model we used results in better performance than the single-layer models. The proposed inter-domain dynamic routing scheme is not only flexible and scalable like the scheme used in the Internet but also dedicated to special concerns/factors from both optical and digital transport layers. Experimental results show that our scheme can effectively set up end-to-end connections across multiple domains. We leave other issues such as avoidance of routing loops to the protocol design. From the protocol point of view, this scheme is not bundled to but can work with the BGP/GMPLS solution.

REFERENCES

1. ITU-T Rec. G.8070/Y.1301, "Requirements for the automatic switched transport network (ASTN)," Jun. 2001.
2. ITU-T Rec. G.8080/Y.1304, "Architecture of the automatic switched optical network (ASON)," Nov. 2001.
3. B. Ramamurthy, S. Yaragorla and X. Yang, "Translucent optical WDM networks for the next-generation backbone networks," in Proceedings of IEEE GLOBECOM 2001, San Antonio, TX, Nov. 2001.
4. J. Strand and Y. Xue, "Routing for optical networks with multiple routing domains," OIF Contribution OIF2001.046, Jan. 2001.
5. Y. Xu, A. Basu and Y. Xue., "A BGP/GMPLS solution for inter-domain optical networking," Internet Draft, Work in progress, draft-xu-bgp-gmpls-02.txt, Jun. 2002.
6. E. Mannie (ed.) et al., "Generalized multi-protocol label switching (GMPLS) architecture," Internet Draft, draft-ietf-ccamp-gmpls-architecture-03.txt, Aug. 2002.
7. Y. Zhu, A. Jukan and M. Ammar, "Multi-segment wavelength routing in large-scale optical networks," in Proceedings of IEEE ICC, Anchorage, AK, May 2003.
8. J. Lang, D. Papadimitriou et al., "SONET/SDH encoding for link management protocol (LMP) test messages," Internet Draft, Work in Progress, draft-ietf-ccamp-lmp-test-sonet-sdh-00.txt, Sep. 2002.
9. A. Fredette, J. Lang et al., "Link management protocol (LMP) for DWDM optical line systems," Internet Draft, Work in Progress, draft-ietf-ccamp-lmp-wdm-01.txt, Sep. 2002.
10. J. Lang, et al., "Link management protocol," Internet Draft, Work in Progress, draft-ietf-ccamp-lmp-07.txt, Nov. 2002.
11. ITU-T Rec. G.709, "Network node interface for the Optical Transport Network (OTN)," Feb. 2001.
12. B. Ramamurthy et al., "Impact of transmission impairments on the teletraffic performance of wavelength-routed optical networks," IEEE/OSA Journal of Lightwave Technology, vol. 17, no. 10, pp. 1713-1723, Oct. 1999.