Architecture and Analysis for providing Virtual Private Networks (VPN) with QoS over Optical WDM Networks

Yang Qin

School of Electrical and Electronic Engineering

Nanyang Technological University

Singapore, 639789

Krishna Sivalingam*

School of Electrical Engineering and Computer Science

Washington State University, Pullman, WA 99164, USA

Bo Li

Department of Computer Science

Hong Kong University of Science and Technology

Clear Water Bay, Kowloon, Hong Kong

*Corresponding author: K. Sivalingam. Tel: +1 509-335-3220, Fax: +1 509-335-3818. A conference version of this paper was presented at SPIE/ACM/IEEE Opticomm conference at Dallas, TX in October 2000. This work was supported in part by a grant from Cisco Systems and by Intel Corporation. Part of the work was done while the first author was at Washington State University. E-Mail: eyqin@ntu.edu.sg, krishna@eecs.wsu.edu, bli@cs.ust.hk

Correspondence Information

Prof. Krishna Sivalingam

Boeing Associate Professor of Computer Science

School of Elect. Engg. & Computer Science

102 EME Building

Washington State University

Pullman, WA 99164-2752

Phone: 509 335 3220

Fax: 253-295-9458 (Please email to addr below after sending Fax)

Email: krishna@eecs.wsu.edu

Abstract

In this paper, we study the problem of employing virtual private networks (VPN) over wavelength division multiplexing (WDM) networks to satisfy diverse quality of service (QoS) requirements of different VPNs. A wavelength routed backbone network is considered. A VPN is specified by the desired logical topology and an *a priori* traffic matrix. The network provides three type of paths over which sessions are established: (i) *Dedicated* lightpath (DLP) – an all-optical path spanning intermediate optical cross connects, which is used by exactly one VPN; (ii) *Shared* lightpath (SLP) – an all-optical path shared by multiple VPNs. Access nodes (where E/O conversion takes place) at the border of the optical backbone provide the necessary electronic buffering when contention arises due to shared lightpaths; and (iii) *Multi-hop path* (MHP) – a hybrid path composed of a tandem of optical lightpaths with O/E and E/O conversion at the junction between the two lightpaths. Depending on the QoS requirements of the VPN, one or more of these lightpath types are used to carry the VPN's traffic. Three traffic types are defined – *Type 1* carried over DLPs as far as possible; *Type 2* carried over SLPs; and *Type 3* carried over MHPs. A VPN's traffic matrix will specify information on each of the three different types. The network will then try to accommodate the given requirements, maximizing the the network utilization. In this paper, we present a simulation based analysis of the system performance for different system configurations, e.g. different number of wavelengths on physical links, different number of VPNs that share one lightpath, etc.

Keywords: Optical WDM Networks, Virtual Private Networks, Wavelength Routed Networks, Quality of Service.

1 Introduction

A Virtual Private Network (VPN) may be defined as an overlay network that is built over a public network infrastructure, providing the VPN user with a private network using tunneling, encryption and authentication mechanisms [1]. VPNs are gaining an increased acceptance due to the economic benefits. VPNs may be

built above different types of public networks, such as Frame Relay, ATM or the Internet. The primary advantages of *VPNs over Internet* are their cost-effectiveness and flexibility. However, the disadvantages of VPNs over Internet are the lack of reliability and sufficient Quality of Service (QoS) mechanisms.

Optical wavelength division multiplexing (WDM) technology, that provides substantial bandwidth capacity, is becoming a practical reality with recent technological advances [2]. Such networks are expected to play an important role in the future wide area networks (WANs). There is a large number of research ideas on supporting "data directly over optics" on WDM networks. This has been fueled by the promise that the elimination of unnecessary network layers will lead to a vast reduction in the cost and complexity of the network [3].

In this paper, we explore how VPNs can be supported in optical WDM networks, in particular WDM mesh routed networks. The WDM routed network provides an "optical connection" layer which consists of several *lightpaths*. A lightpath is defined as an *all-optical* connection from the source node to the destination node, traversing several intermediate optical wavelength routing (or cross-connect) nodes. The optical core is composed of these wavelength router nodes which may possess a limited degree of wavelength conversion capability. *Access nodes* exist at the boundary of the backbone network and provide the interface between the electronic data equipment and the optical core. The access nodes perform E/O conversion when data enters the core, and O/E conversion when data leaves the core.

The network architecture considered is as follows: A network provider owns an optical WDM backbone network and provides capacity to users (e.g. large corporations) requiring VPN services. A VPN is specified by a set of nodes that need to be interconnected and *a priori* traffic demands for the VPN. This is similar to the virtual topology concept that has been studied earlier [4]. The difference in this work is that we consider a set of logical topology specifications, and the different types of lightpaths described later. The network provider's objective is to maximize the total amount of VPN traffic, meet the QoS specifications and optimally utilize the backbone capacity. The proposed architecture separates different VPNs in the optical

domain by providing lightpaths with different transmission qualities to meet the QoS requirements of the different VPNs.

The network provides three type of paths over which the VPN traffic is carried: (i) *Dedicated* lightpath (DLP) – an all-optical path spanning intermediate optical cross connects, which is used by exactly one VPN. The delay incurred by VPN traffic is due to propagation delay, wavelength conversion delay, and O/E and E/O conversions at the access nodes. (ii) *Shared* lightpath (SLP) – an all-optical path shared by multiple VPNs. Access nodes (where E/O conversion takes place) at the border of the optical backbone provide the necessary electronic buffering when contention arises for a lightpath. Here, once data enters the core, the only delays are due to propagation and wavelength conversion. However, there is the additional queuing delay at the access nodes; and (iii) *Multi-hop path* (MHP) – a hybrid path composed of a tandem of optical lightpaths with O/E and E/O conversion at the junction between two lightpaths. Here, additional delays can take place at the light-path junctions if queuing is necessary. This type of path is similar to the classic IP path where queuing is done at intermediate IP routers.

Depending on the QoS requirements of the VPN, one or more of these lightpath types are used to carry the VPN's traffic. *Type 1* traffic is carried over a DLP as far as possible, then *Type 2* traffic is carried over a SLP as far as possible; and finally *Type 3* traffic is carried over a MHP. Given a set of different VPNs and their specification, the network first attempts to meet the demands of Type 1 traffic using DLPs, Type 2 traffic using SLPs, and Type 3 using MHPs. When dedicated LPs are no longer available, it establishes shared LPs where the number of VPNs sharing a lightpath is limited based on certain performance specifications. When shared LPs are not feasible, multi-hop paths are set up.

We conduct simulations to investigate the performance of the network in terms of average packet delay. System parameters that are varied include the number of VPNs, number of wavelengths, and the traffic patterns.

The rest of the paper is organized as follows. In Section 2, a brief background on VPNs is provided.

In Section 3, the proposed framework for supporting VPN over WDM networks is presented. Preliminary results from simulation analysis are discussed in Section 4. A summary and description of ongoing research is provided in Section 5.

2 Background

A virtual private network uses a public network's infrastructure to make the connections among geographically dispersed nodes, instead of using cables owned or leased exclusively for one single network's use, as is typical for a wide area network (WAN). To the user, a VPN looks like a private network, even though it shares the network with other users. There are several uses for a VPN. It can be an extended intranet, connecting geographically distant facilities into a cohesive network. It can also be an extranet, linking customers and suppliers for increased efficiency. Although, there are several type of public networks that can be employed to create a VPN, the most popular and prominent VPNs are based on the Internet. The primary advantages of VPNs over Internet are cost-efficiency, flexibility and scalability [1].

The chief mechanisms that enable VPN provisioning are tunneling and security. With tunneling, each packet is encapsulated by a new envelope or capsule that carries the addresses of the source and destination VPN servers. In this encapsulation process, the VPN software appends a new header, which contains a new source and destination address to the packet before sending it out on the Internet. Security is provided by encryption, authentication and other mechanisms.

Although providing VPN is cost-effective and flexible, there are a few problems. Quality of service (QoS) is difficult to guarantee when traffic is encrypted because the bits marking QoS cannot be read by the routers. Tunneling protocol cannot guarantee the minimum delay due to IP's best effort packet forwarding. The current VPNs over Internet are limited to handling low-priority enterprise traffic. With with rapid emergence of e-commerce and VPNs, reliability and security are becoming a great challenge for Internet-

based VPNs. With various enterprises turning to VPNs, providing diverse QoS becomes another important issue since different applications have different size and delay sensitivities. These requirements would accelerate the development of new technologies for the next generation VPNs. The next section examines how optical WDM networks can be used to support VPNs.

3 Proposed Architecture for VPN over WDM

In the context of future optical networks, providing QoS is one of the critical research issues. Traditional optical networks such as synchronous optical networks / synchronous digital hierarchy (SONET/SDH) have been perceived as high transmission rate networks without provision for any QoS to different traffic flows. Recently, some attention has been given to coarse-grain QoS using differentiated optical services [5]. By applying the virtual private network concept to WDM, we explore how QoS may be provided.

3.1 Basic Framework

In this section, we discuss the framework for employing virtual private networks over WDM as illustrated in Fig. 1. The wide area network connectivity is provided by a wavelength routed backbone network. The optical network consists of several switch nodes interconnected by multi-wavelength WDM links. The access nodes provide the electronic interface to end users, which may be regional networks that feed into the optical core network. The basic idea of this work involves segregating different VPN traffic types in the wavelength domain to provide support for tunneling and QoS.

The goal is to establish several VPNs on the physical topology, where each VPN is specified by a set of constituent nodes that comprise it, and the long term average traffic demands. In addition, different VPNs may have different QoS requirements such as bounded delay, guaranteed bandwidth, etc. Thus, each VPN can be viewed as a *logical* or *virtual topology* that is embedded on the physical topology [4].

For example, VPN 1 is specified by a topology that consists of the three access nodes, A, C and D. For

this VPN, lightpaths will be established in the optical network between the pairs $\{(A, C), (C, D), (A, D)\}$. We assume that the wavelength converter is not available in all the switch nodes in this paper. The lightpath for (A,C) is given by $\{1, 3, \lambda_1\}$; for (C,D) given by $\{3, 4, \lambda_2\}$; and for (A,D) given by $\{1, 3, 4, \lambda_3\}$. Thus, all of VPN 1's traffic will be carried by these lightpaths. For VPN2, only two nodes, (A, C) are specified and the lightpath assigned is $\{1, 3, \lambda_2\}$.

Given the above framework, we formulate the problem as follows. The inputs are the physical topology of wavelength routed networks, number of wavelengths on each physical link, number of transmitters and receivers on each switch node, a set of VPNs with their topologies, traffic demands and QoS requirements. The objective function is to maximize the amount of traffic carried by the VPNs, subject to the physical constraints and QoS requirements.

3.2 Different types of Lightpaths

The wavelength routed network provides the following three different lightpath types:

Dedicated Lightpath (DLP): is an all-optical path where the traffic is carried entirely in the optical domain within the backbone. The path is composed of links spanning intermediate optical routers with a transmission wavelength specified for each link. It is dedicated since it is allocated to carry exactly one VPN's traffic. This is the most expensive lightpath type, and its utilization will depend entirely upon the alloted VPN's traffic. The delay experienced by the VPN along this path will be the end-to-end propagation delay and wavelength conversion delay. Therefore, this lightpath can be alloted to VPN traffic that demands the highest level of QoS in terms of bandwidth or delay.

Shared Lightpath (SLP): is also an all-optical path, but it is shared among multiple VPNs. The maximum number of shared VPNs and the sharing mechanism is determined by the various QoS requirements. When a SLP is currently used by a VPN, traffic from another shared VPN that arrives at the access node is electronically buffered and transmitted after the considered VPN completes the transmission. When several

competing VPN traffic arrive during a busy period, an appropriate scheduling algorithm has to be used. This type of service is less expensive than the DLP, but the lightpath utilization is higher due to traffic contribution from several VPNs. The additional delay incurred in comparison to DLP will be the queuing delay at the access nodes.

Multi-hop (MHP): A multi-hop path is composed of a sequence of optical lightpaths in tandem. O/E and E/O conversion is done at the junction between two lightpaths. Since the component lightpaths may be shared among several MHPs, electronic buffering is required at the intermediate routers. This is the least expensive service among the three. The additional delay incurred in comparison to SLP will be due to O/E and E/O conversion at the junction between multiple lightpaths, and queuing at intermediate routers.

Using the lightpath setup, we implement the corresponding function of the tunneling mechanism directly at the optical layer. Thus, it is not necessary to apply a tunneling protocol which may append a new header to the original packet, thereby increasing the communication efficiency. Circuit-switched service is provided as far as possible. For mesh WAN networks, adopting a packet-switched routing scheme makes it harder to predict the overall delay. The lightpaths, both dedicated and shared, can provide guarantee on packet delay once they enter the optical core. Thus, the three types of paths can be used to design different types of quality of service, as described below.

Type 1 traffic: This type of traffic requires only dedicated lightpaths, and has stringent QoS requirements (for e.g., an upper bound on delay).

Type 2 traffic: For this kind of traffic, a shared lightpath is provided. The delay requirements are still high, requiring an all-optical path, but they are less stringent than that of the Type 1 traffic.

Type 3 traffic: For this kind of traffic with minimal or no QoS requirements, the multi-hop lightpath is provided.

Each VPN specifies the traffic demand for each traffic type. Given the set of VPN traffic and topology requirements and the physical topology, the task is to establish the lightpaths to meet these requirements. This requires the determination of the route and the wavelength assignment (RWA) for each lightpath. A survey of the different RWA is available in the literature [4, 6].

Let there be a total of V VPNs and N wavelength routed nodes in the network. Let $\mathrm{VPN}_i^p(s,d)$ represent the traffic demand carried for VPN i from source node s to destination node d for traffic type p. The following is the objective function:

$$max\left(\sum_{s,d\in N}\sum_{i\in V}\mathrm{VPN}_i^p(s,d)\right) \tag{1}$$

This is an NP-hard problem since it is a generalization of the routing and wavelength assignment (RWA) problem [7] that has been proven to be NP-hard. In the original problem, there is only single traffic type and a single VPN (i.e. logical topology), and there is no QoS considered when setting up the lightpaths.

3.3 Lightpaths Establishment Algorithm

In our proposed architecture, the following steps are taken to accomplish the LP establishment:

Step 1: We try to establish DLPs for all Type 1 traffic. The input to the algorithm is the set of entries in all the VPN_i^1 matrices. For each entry, a lightpath will be created. We adopt a heuristic algorithm presented in [8] to setup the lightpaths.

Traffic entries for which DLPs are not possible will be routed using SLPs. This will mean that their QoS requirements may be violated. The characterization of this problem will be a subject of future study.

Step 2: We try to establish the SLPs for Type 2 traffic. A single matrix $\mathcal{VPN}^2$ is created from the given VPN traffic matrix information as follows:

$$\mathcal{VPN}^2(s,d) = \sum_{i=1}^{V} \mathrm{VPN}_i^2(s,d)$$

This traffic matrix is then fed to the heuristic algorithm that establishes the lightpaths. Thus, for each entry in the matrix $\mathcal{VPN}^2$, a lightpath is established. Since each entry is the sum of the individual VPN matrix entries, the corresponding LP is shared among the different VPNs.

The heuristic algorithm used in Step 1 is used here too. In an effort to limit the number of VPNs sharing a SLP, we can establish multiple SLPs for one entry using thresholds based on the size of the entry or the number of shared VPNs.

Since it is possible that the entire traffic demand may not be met with SLPs (due to capacity limitations), some of the traffic may be carried over MHPs.

Step 3: Next, we establish the MHPs for Type 3 traffic. As before, we create a single matrix $\mathcal{VPN}^3$ from the given VPN traffic matrix information:

$$\mathcal{VPN}^3(s,d) = \sum_{i=1}^{V} \mathrm{VPN}_i^3(s,d)$$

For each entry in the Type 3 matrix, a multi-hop route composed of lightpaths is determined. Here again, due to the summation of different VPN requests, sharing is done implicitly. The routing algorithm we adopt here is the same as in [8] which is a simple shortest path algorithm.

3.4 Heuristic Algorithm

In this section, we describe the heuristic algorithm presented in [8]. The basic idea in this algorithm is to establish lightpaths in descending order of the traffic matrix entries. Therefore, the algorithm first assigns a

wavelength to the optical connection with the largest pairwise traffic demand. Then, it assigns a wavelength to the connection with the next largest pairwise traffic demand among the connections which do not use the links used by the first connection, and so on.

The algorithm first generates a connection-link indication matrix M of size $N^2 \times N^2$. The matrix is represented by $M = [m_{(ij),(lm)}]$, where the entry $m_{(ij),(lm)}$ is 1 if the path from $Node\ i$ to $Node\ j$ and that from $Node\ l$ to $Node\ m$ use a common link; otherwise it is 0. A simple shortest path algorithm is used to determine the route. In our simulation, we assume that each link has the same distance in the physical topology. Then this shortest path algorithm will give the same result as the least number of hops. The connection-link indication matrix m is generated based on the order of traffic demand. In the matrix M, each connection corresponds to one column. Once the matrix is obtained, the algorithm could be implemented as follows: Assign the wavelength to the first column. All the columns with elements equaling 0 in the first row are candidates for the next wavelength assignment and the first such column, say column i is chosen. Next, the wavelength is assigned to the first column with elements equaling zero in both $row\ 1$ and $row\ i$. The procedure is repeated until no such column can be found.

The complexity of this algorithm is $O(N^4)$. The complexity of this heuristic algorithm can be reduced to $O(N^3)$ as shown in [9]. However, for the sake of simplicity in illustrating the framework, we adopt the original heuristic algorithm in [8]. Future research is necessary to consider more efficient RWA algorithms.

3.5 Example

An example physical network is shown in Fig. 2. The graph is undirected and the distance of each link is shown. There are 2 wavelengths on each link.

Given the overall traffic matrix for Type 2 traffic ($\mathcal{VPN}^2$), in Table 1, we allocate the lightpaths for VPNs using the heuristic algorithm.

The first step is to determine the shortest paths between all source-destination pairs using a standard al-

Node ID	1	2	3	4	5
1	0	8	3	1	6
2	2	0	4	3	9
3	3	1	0	2	2
4	4	2	3	0	10
5	1	4	7	6	0

Table 1: Example $\mathcal{VPN}^2$ Traffic Matrix.

gorithm, such as Dijkstra's shortest-path algorithm [10]. The weight functions can vary, but in this example, we will use the distance metric.

Next, we establish the lightpaths. We use 5 units as the capacity threshold - that is, one lightpath is assigned for every 5 units of traffic. Allocating in the descending order of the traffic demand, we set up the following lightpaths: (4,5) will use $\{4, 5, w_1\}$ and $\{4, 5, w_2\}$, (2,5) will use $\{2, 3, 5, w_1\}$ and $\{2, 3, 5, w_2\}$, (1,2) will use $\{1, 2, w_1\}$ and $\{1, 2, w_2\}$, (5,3) will use $\{5, 3, w_1\}$ and $\{5, 3, w_2\}$, (1,5) will use $\{1, 5, w_1\}$ and $\{1, 5, w_2\}$, (5,4) will use $\{5, 4, w_1\}$ and $\{5, 4, w_2\}$, and (4,1) will use $\{4, 2, 1, w_1\}$. The (5,2) traffic pair could not be assigned a lightpath since wavelengths are not available on its shortest path. Thus, (5,2) is dropped to Type 3 traffic; similarly, (2,3) is dropped. Continuing with the example, (2,4) will use $\{2, 4, w_1\}$, (4,3) will use $\{4, 3, w_1\}$, (1,3) will be dropped, (3,1) will use $\{3, 2, 1, w_1\}$, (2,1) will use $\{2, 1, w_2\}$, (3,4) will use $\{3, 4, w_1\}$, (3,5) will be dropped, (4,2) will use $\{4, 2, w_2\}$, (1,4) will be dropped, (3,2) will use $\{3, 2, w_2\}$, and (5,1) will use $\{5, 1, w_1\}$.

The above example indicates that other routing and wavelength assignment algorithms may be used to improve efficiency. This is reserved for further research.

4 Simulation Results

In this section, we present performance results obtained using discrete event simulation. Since Type 1 traffic experiences only propagation delay, we are more interested in the performance of Type 2 and Type 3 traffic. Both of these two traffic types experience buffering delay at the access node. Contention delay arises due to the packets from other VPNs that share the same lightpath or the same routing paths. In the discussion below, V denotes the number of VPNs, N the number of nodes, and C the number of wavelengths per link.

4.1 Simulation Details

The performance metric studied is average packet delay which is defined as the time between packet generation and reception. The relevant details of the simulation are as follows:

- The physical topology considered is shown in Fig. 3. The graph is an undirected network with 24 optical nodes. Every node is associated with an access node.

- Initially, we use a single matrix per VPN, and there are a total of V traffic demand matrices. These matrices are added to get the overall traffic matrix T that is used for lightpath establishment.

  The algorithm tries to establish the maximum number of SLPs. Traffic that is not carried over these SLPs is carried over MHPs.

- During simulation, individual packets that make up a session are generated based on the overall traffic matrix, T and satisfying the Poisson distribution.

  Packets are of fixed length and the transmission time for one packet is one slot (where a slot is a fixed unit of time).

- Packets at Node i are generated independently of packets originating at other nodes. Packets are generated according to a Poisson process with rate $\lambda \sum_{j=1}^{N} T_{ij}$. A packet generated at Node i is destined

for Node j with probability $T_{ij} / \sum_{j=1}^{N} T_{ij}$. Thus, the matrix T depicts the predicted a priori traffic, while the simulation generates a variable number of packets using the parameter $\lambda \times$ corresponding T_{ij}.

- The O/E/O delay, incurred with multi-hop paths, is assumed to be 20 times that of the transmission time of a packet.

4.2 Discussion of Results

Varying traffic generation rate λ: The results presented in Fig. 4 are for a system with 3 VPNs and $C = 4$ wavelengths. The individual traffic demand matrices were randomly generated with each entry ranging from 0 to 10. As observed, the delay for Type 2 traffic is much smaller than that of Type 3 traffic which is mainly due to the multiple hops and the resulting O/E/O conversions. As expected, increasing λ results in increased delay.

Multiple LPs per traffic entry: For this experiment, the long term traffic pattern presented in [8] is used. The results shown in Fig. 5(a) are for $V = 4$ and $C = 16$ with one SLP per traffic matrix entry. We find that the delay for Type 2 traffic is high, and for some λ values, higher than that of Type 3 traffic. The large delay is due to the higher volume of Type 2 traffic and the fact that only one lightpath is provided for one pair of Type 2 traffic demand. The results shown in Fig. 5(b) indicate that using multiple lightpaths for large traffic demands reduces the delay. We provide up to 14 lightpaths for one pair of traffic demands from Type 2 VPNs. A simple threshold value based on the traffic demand entry is used to determine if multiple SLPs are needed.

Varying V and C: The results shown in Fig. 6 are based on randomly generated traffic patterns, that have about 80% non-zero entries. Furthermore, each entry in the traffic matrix has the same value. Fig. 6(a) shows that with the increased number of wavelengths, the delay for mixed traffic of Type 2 and Type 3 is

greatly decreased. Fig. 6(b) shows that for a network with 16 wavelengths, when the number of VPNs is increased, the delay for mixed traffic of Type 2 and Type 3 increases as expected.

Fig. 7(a) presents results obtained by varying C keeping $V = 4$ with traffic matrices having 60% non-zero entries. The graph shows that when the number of wavelengths increases, the delay for the total traffic decreases as expected. This is because we can support more Type 2 QoS VPN traffic demands when we have more channels on each physical link. The delay for the total traffic demand is reduced from 14 slots to about 5 slots. When the number of wavelengths is equal to 40, all of the traffic demand pairs have all-optical lightpath, therefore the delay is the same as the Type 2 VPN packet's delay.

In Fig. 7(b), results are presented for the case where number of Type 2 VPNs sharing one lightpath is set to 5 and 10 for $C = 16$. It is seen that when the number of Type 2 VPNs that share one lightpath is increased, the mean delay increases. The delay for the 5 VPN system is always smaller than 1.7 slots, and the delay for the 10 VPN system is larger than 2.2 slots.

5 Summary

In this paper, we present a framework for supporting VPNs with different QoS requirements over optical WDM networks. We formulate the off-line problem where a physical topology and a set of VPNs are provided, and the objective is to maximize the total traffic demand of VPNs that can be supported. With our simulation, we demonstrate that we can provide different QoS for different VPN traffic streams. In addition, we present a simulation analysis of the system performance with delay as the metric, varying different system parameters such as number of wavelengths, number of VPNs, and traffic generation rates.

References

[1] D. Fowler, *Virtual Private Networks.* Morgan Kaufmann Publishers, 1999.

[2] K. Sivalingam and S. Subramaniam, eds., *Optical WDM Networks: Principles and Practice.* Boston, MA: Kluwer Academic Publishers, 2000.

[3] P. Bonenfant and A. Rodriguez-Moral, "Optical Data Networking," *IEEE Communications Magazine*, vol. 38, pp. 63–70, Mar. 2000.

[4] G. Rouskas, "Design of Logical Topologies for Wavelength Routed Networks," in *Optical WDM Networks: Principles and Practice* (K. M. Sivalingam and S. Subramaniam, eds.), ch. 4, pp. 79–102, Boston, MA: Kluwer Academic Publishers, 2000.

[5] N. Golmei, T. Ndousse, and D. Su, "A differentiated optical services model for WDM networks," *IEEE Communications Magazine*, vol. 38, Feb. 2000.

[6] H. Zang, J. P. Jue, and B. Mukherjee, "A review of routing and wavelength assignment approaches for wavelength-routed optical WDM networks," *Optical Networks Magazine*, vol. 1, pp. 47–60, Jan. 2000.

[7] R. Ramaswami and K. N. Sivarajan, "Design of logical topologies for wavelength-routed optical networks," *IEEE Journal on Selected Areas in Communications*, vol. 14, pp. 840–851, June 1996.

[8] Z. Zhang and A. Acampora, "A heuristic wavelength assignment algorithm for multihop WDM networks with wavelength routing and wavelength re-use," *IEEE/ACM Transaction on Networking*, vol. 3, pp. 281–288, June 1995.

[9] Y. Qin, B. Li, and G. Italiano, "Low cost and effective heuristic wavelength assignment algorithm in a wide-area WDM based all optical network," in *The 14th International Conference of Information Networks (ICOIN'2000)*, Jan. 2000.

[10] N. M. Bhide, K. M. Sivalingam, and T. Fabry-Asztalos, "Routing Mechanisms Employing Adaptive Weight Functions for Shortest Path Routing in Multi-Wavelength Optical WDM Networks," *Journal of Photonic Network Communications*, Dec. 2000. (Accepted for Publication).

Yang Qin (IEEE Member) received the B.S degree in computer science from Southwest Jiaotong University, China, in 1989 and the M.Sc degree in Computer Science from Huazhong University of Science and Technology, China, in 1992. She got her PhD degree in Computer Science from Hong Kong University of Science and Technology in 1999. From December 1999 until November 2000 she has worked as a research associate in school of Electrical Engineering and Computer Science of Washington State University, Pullman, WA. She is presently an Assistant Professor in the School of Electrical and Electronic Engineering, Nanyang Technological University in Singapore. Her research interests include optical networks, photonic switching, performance evaluation, protocols supporting quality of service (QoS) and WDM routing networks. She is a member of IEEE. Email: eyqin@ntu.edu.sg

Krishna M. Sivalingam (ACM '93, IEEE SM '00 M '95) received his Ph.D. and M.S. degrees in Computer Science from State University of New York at Buffalo in 1994 and 1990 respectively. While at SUNY Buffalo, he was a Presidential Fellow from 1988 to 1991. Prior to that, he received the B.E. degree in Computer Science and Engineering in 1988 from Anna University, Madras, India. He is Boeing Associate Professor of Computer Science in the School of Electrical Engineering and Computer Science, at Washington State University, Pullman, where he was an Assistant Professor from 1997 to 2000. Earlier, he was an Assistant Professor at University of North Carolina Greensboro from 1994 until 1997. He has conducted research at Lucent Technologies' Bell Labs in Murray Hill, NJ, and at AT&T Labs in Whippany, NJ.

His research interests include wireless networks, optical wavelength division multiplexed networks, and performance evaluation. He has served as a Guest Co-Editor for a special issue of the IEEE Journal on Selected Areas in Communications on optical WDM networks. He is co-recipient of the Best Paper Award at the IEEE International Conference on Networks 2000 held in Singapore. He has published an edited book on optical WDM networks in 2000. His work is supported by AFOSR, Laboratory for Telecommunication Sciences, NSF, Cisco, Bellcore, Alcatel, Intel, and Washington Technology Center. He holds three patents in wireless networks and has published several papers including 18 journal publications. He has served on several conference committees including ACM Mobicom 2001, Opticom 2001, Opticom 2000, ACM Mobicom 1999, MASCOTS 1999, and IEEE INFOCOM 1997. He is a Senior Member of IEEE and a member of ACM. Email: krishna@eecs.wsu.edu

Bo Li (IEEE SM) received B.S (cum laude) and M.S degrees in computer science from Tsinghua University, Beijing, China, in 1987 and 1989, respectively, and a Ph.D degree in computer engineering from the University Mas-

sachusetts at Amherst in 1993. Between 1994 and 1996 he worked on high-performance routers and ATM switches in IBM's Networking Systems Division, Research Triangle Park, North Carolina. He joined the faculty of the Computer Science Department of Hong Kong University of Science and Technology in January 1996. He has been on the editorial board of ACM Mobile Computing and Communications Review (MC2R), ACM/Baltzer Journal of Wireless Networks (WINET), Journal of Communications and Networks (JCN), and IEEE Journal of Selected Areas in Communications (IEEE JSAC Wireless Series). He has been co-guest of special issues for IEEE Communications Magazine, IEEE Journal of Selected Areas in Communications and the upcoming SPIE/Baltzer Optical Networks. He has been involved in organizing numerous conferences, particularly IEEE INFOCOM, since 1986, and is Vice-Chair of INFOCOM 2001.

His current research interests are: wireless mobile networking supporting multimedia, voice and video transmission over the Internet, and all-optical networks using WDM. He is a member of ACM. Email: bli@cs.ust.hk.

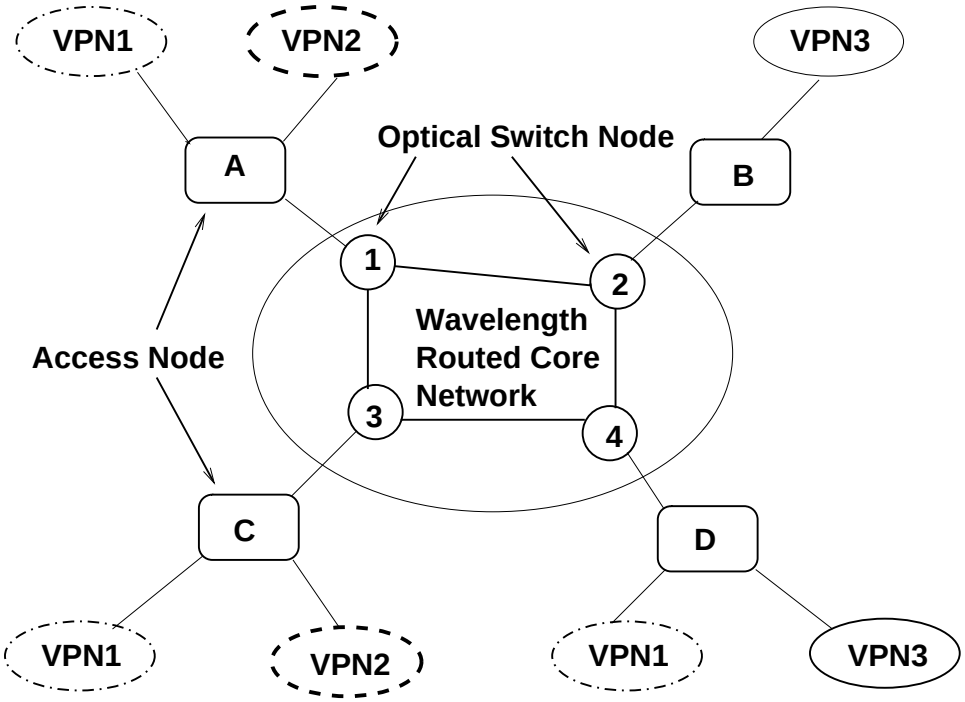

Figure 1: Framework for VPN over optical WDM networks.

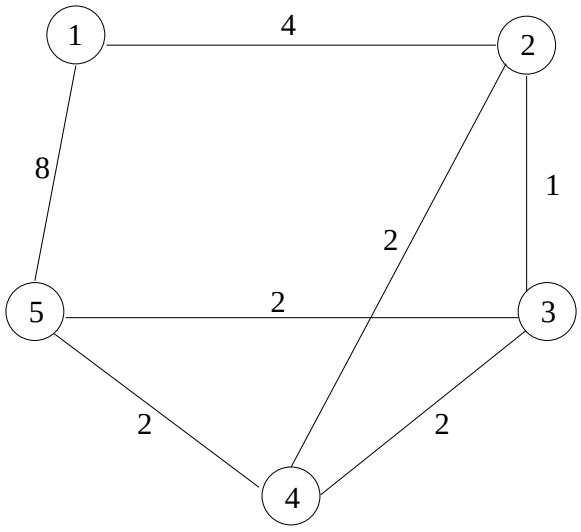

Figure 2: Physical topology used in the example to explain the SLP generation algorithm.

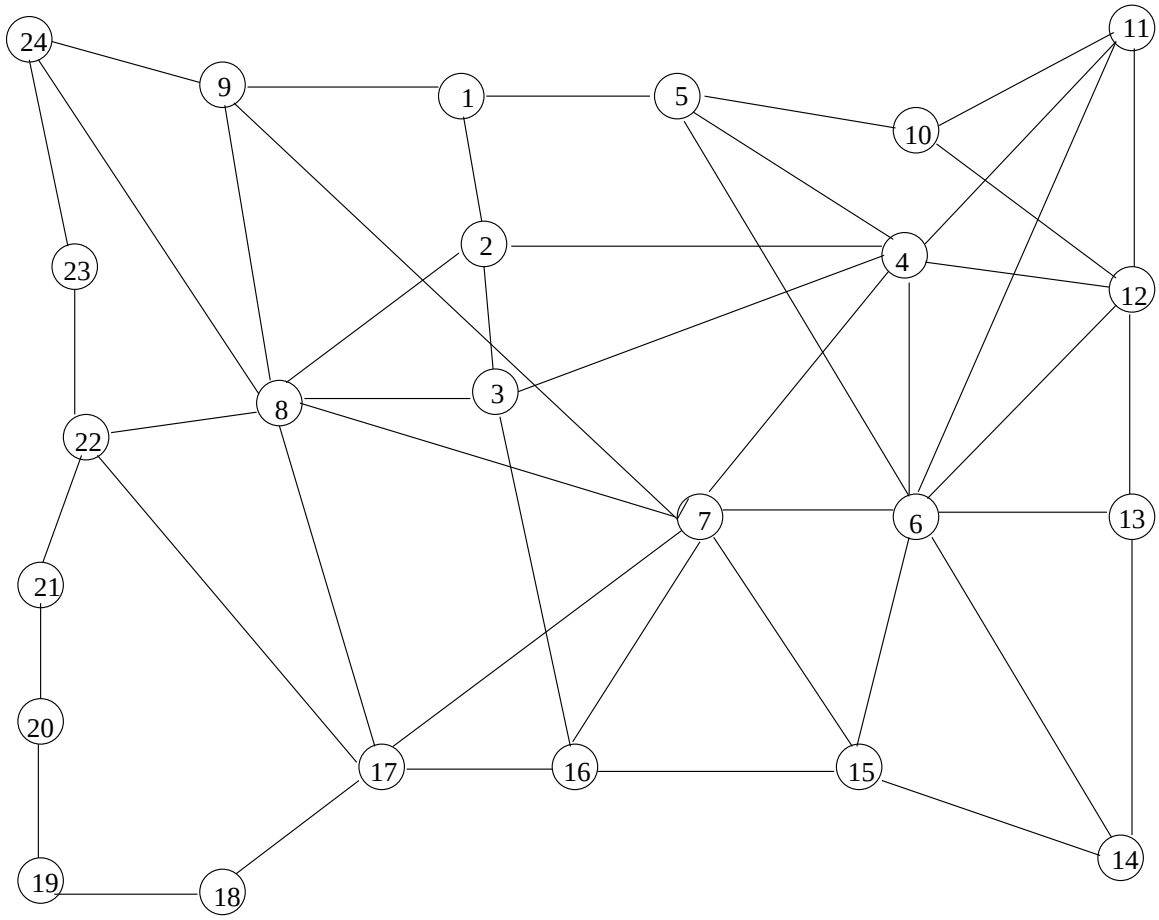

Figure 3: The physical topology used in the performance analysis.

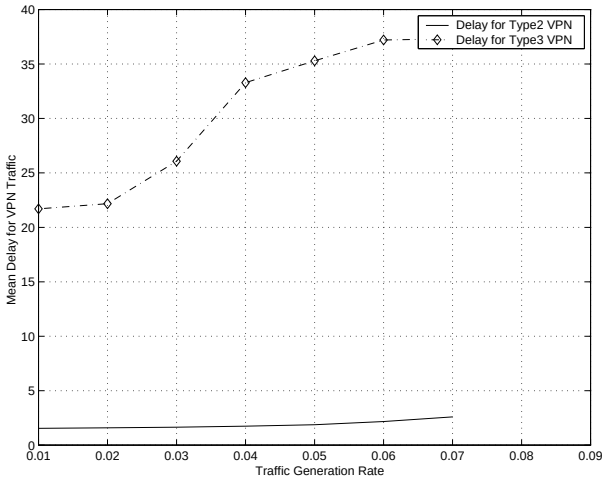

Figure 4: Delay for Type 2 and Type 3 traffic types, for $V = 3$ and $C = 4$.

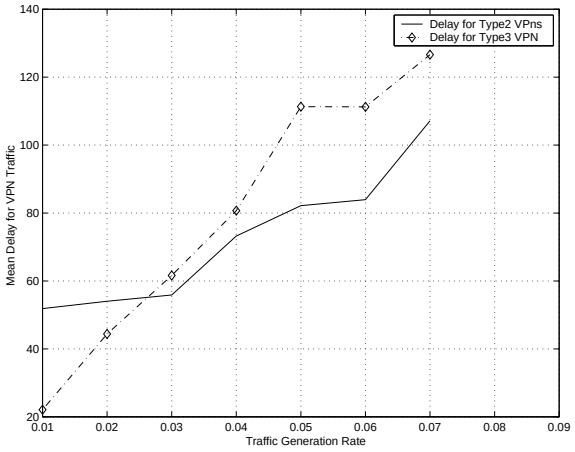

(a) Single lightpath per traffic entry

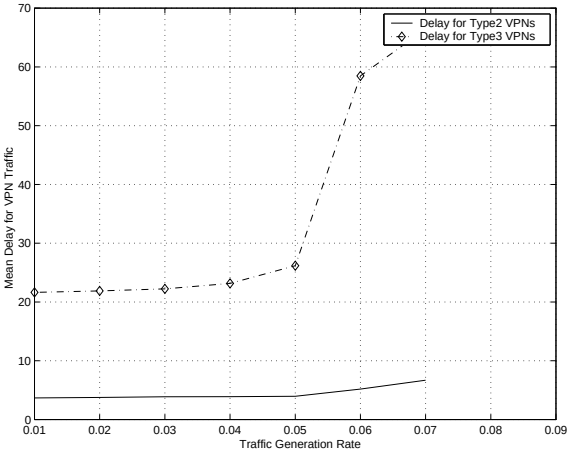

(b) Multiple lightpaths per traffic entry

Figure 5: Delay analysis using single and multiple LPs per traffic entry.

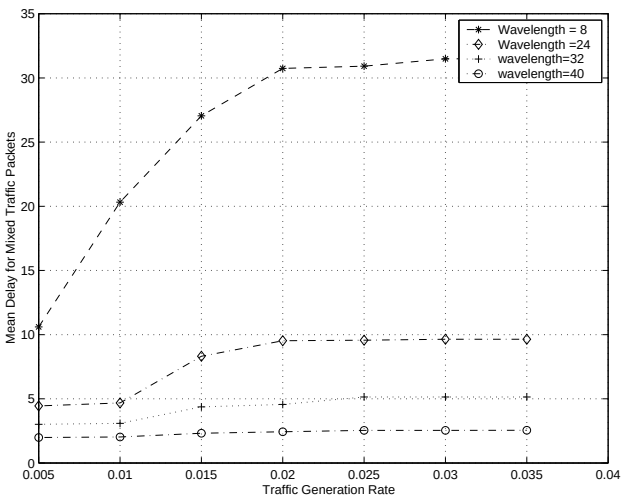

(a) Varying number of wavelengths, C

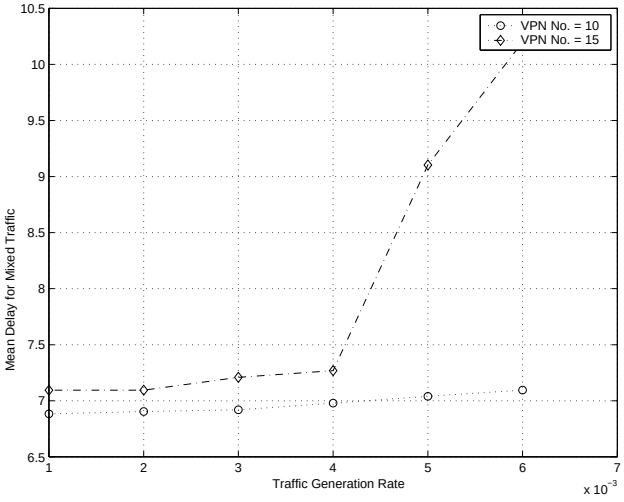

(b) Varying number of VPNs, V

Figure 6: Delay analysis, varying number of wavelengths (C) and number of VPNs (V) with 80% non-zero entries.

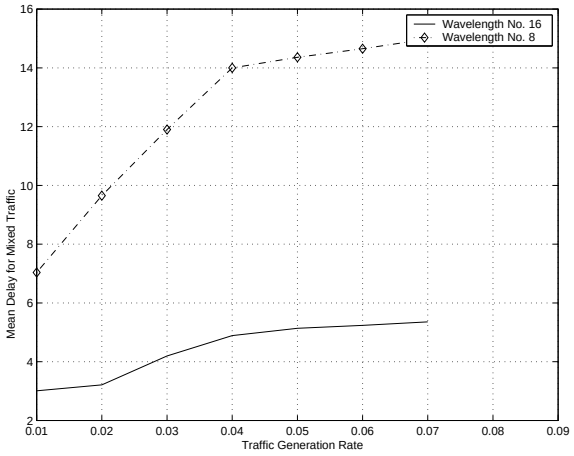

(a) Varying number of wavelengths, C

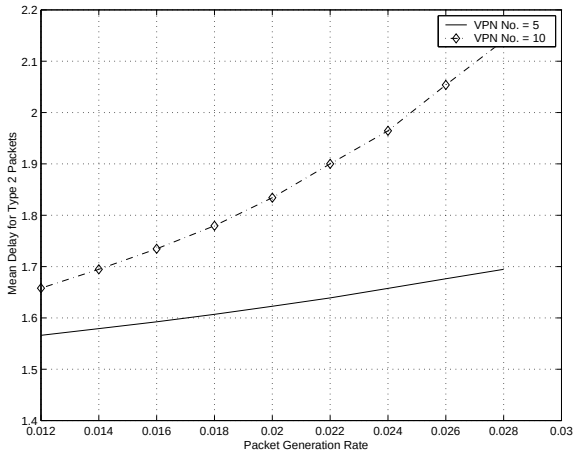

(b) Varying number of VPNs, V

Figure 7: Delay analysis, varying number of wavelengths (C) and number of VPNs (V) with 60% non-zero entries.