FORMULATION AND EVALUATION OF OPTIMIZATION

MODELS FOR MPLS TRAFFIC ENGINEERING

WITH QOS REQUIREMENTS

Approved by:

Dr. Richard Barr

Dr. Richard Helgason

Dr. Jeffery Kennington

Dr. Marion Sobol

Dr. Wayne Woodward

FORMULATION AND EVALUATION OF OPTIMIZATION MODELS FOR MPLS TRAFFIC ENGINEERING WITH QOS REQUIREMENTS

A Praxis Presented to the Graduate Faculty of the

School of Engineering

Southern Methodist University

in

Partial Fulfillment of the Requirements

for the degree of

Doctor of Engineering

with a

Major in Engineering Management

By

Saib K. Jarrar

(B.S., Technion, 1982)
(M.S, American University, 1987)

May 15, 2004

ACKNOWLEDGEMENTS

I would like to acknowledge all who directly or indirectly contributed to the preparation of this thesis. Thanks to Dr. Richard Barr for his guidance and critical input when it most mattered. Also thanks to him for enhancing the clarity and accuracy of this manuscript and peeking my interest in understanding the proper use of words. Thanks to Dr. Richard Helgason for leading the MPLS research team and establishing a framework that helped me stay on track and complete my research. Thanks to my wife Manal, son Fiross, and daughters Ranya and Rawan for enduring with me the long journey. Thank you Manal, for challenging me to finish the writing and offering you the first draft as my gift for your birthday. Finally, thanks to all those who knowingly or unknowingly inspired me to persevere.

Jarrar, Saib K.
B.S., Technion, 1982
M.S., American University, 1987

Formulation and Evaluation of Optimization Models for MPLS Traffic Engineering with QoS Requirements

Advisor: Professor Richard S. Barr

Doctor of Engineering Management conferred May 15, 2004

Praxis completed May 5, 2004

Multi-Protocol Label Switching (MPLS) is an evolving switching technology that is being integrated into Internet Protocol (IP) networks to overcome IP-routing deficiencies. MPLS facilitates traffic engineering (TE) by providing the mechanisms needed to control traffic flows in IP networks. Combined with differentiated services (Diffserv) capabilities, MPLS enables the implementation and support of multiple classes-of-service (CoS) types, each with specific quality-of-service (QoS) guarantees. Thus, MPLS facilitates network optimization to maximize resource utilization and enables the convergence of data, voice, and video applications over a common network infrastructure.

This praxis addresses multiple fundamental problems related to MBLS-based TE in IP networks including: the basic TE problem of constraint-based routing and admission control with single CoS type, the problem of load balancing, the problem of TE with multiple CoS types, and the problem of capacity planning. These problems are formulated as origin-destination integer multi-commodity network-flow models. The

models focus on revenue maximization, which is one of the primary goals of MPLS deployment by service providers.

To explore the effectiveness of these models in addressing practical problems from the telecommunications industry, a series of computational experiments are performed. A suite of network instances, with varying topologies that mirror realistic MPLS design problems, are constructed and tested on a wide range of parameter values. The results are evaluated to test a series of research hypotheses and develop insight into effectively engineered MPLS networks.

The models have practical applications and can be used by network administrators and managers in the TE design process. One of the main results of this research is the conclusion that partitioning a demand into multiple smaller demands for different classes, and routing them separately, indirectly realizes the benefits associated with load balancing and results in increasing the traffic delivery ratio. This result is of major significance since it suggests that service providers may exploit a revenue-generating service feature as a vehicle to increase network efficiency and promotes the adoption of Diffserv-aware traffic engineering.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

This dissertation is dedicated to the loving soul of my dad Khair A. Jarrar.
May your spirit continue to rest in peace.

Chapter 1

INTRODUCTION

Over the last decade the public Internet has evolved from a limited U.S. government-sponsored network serving the education and research communities to a gigantic global, robust, and ubiquitous commercial network. The Internet has evolved into a critical communications network at the heart of the new information-oriented economy serving both consumers and businesses. Internet growth—in terms of number of users and traffic volume—has been phenomenal and is expected to continue .

The growth and popularity of the public Internet has accelerated the adoption of the Internet Protocol (IP) [38] as a dominant communications technology. The Transmission Control Protocol/Internet Protocol (TCP/IP) [37] suite of protocols has been adopted as the protocol of choice by enterprise networks for both internetworking and applications. Carriers are now aggressively pursuing new Virtual Private Network (VPN) offerings that are based on IP technology [24, 25]. It is expected that these new services will replace current private-line and virtual data services such as Frame Relay (FR) and Asynchronous Transfer Mode (ATM).

Another important industry trend is the convergence of data communications and telecommunications. This convergence is driven by economic pressure to achieve cost savings and increase operational efficiencies. Enterprise customers are striving to embrace one common communications infrastructure to service their data, voice, and video applications. Technology advancements in packet voice and Voice over IP (VoIP), in particular, are accelerating and promoting that convergence over IP networks.

1.1 Motivation

Handling the explosive traffic growth and achieving convergence present serious challenges to the IP technology and service providers. Both Internet and IP-based VPN services are competitive businesses that require continual investment to keep pace with the increase in traffic. Due to IP deficiencies [6, 7, 8, 43] (notably its limitations in controlling and distributing traffic across the network), the conventional answer to dealing with traffic growth has been the over-provisioning of costly resources. In the current and future business environments this answer is not adequate. Carriers are under pressure to contain capital expenditures and are looking for other solutions to maximize the use of network resources. Promising new solutions have introduced traffic engineering (TE) capabilities to IP networks. TE is the process of controlling traffic flow through a network so that network performance and resource utilization are optimized. TE maximizes the benefits of an installed network infrastructure.

The convergence of voice, video, and data traffic imposes new requirements on IP networks. IP networks will need to support multiple traffic types with dissimilar characteristics and requirements. Voice traffic requires the same predictability and dependability as the existing Public Switched Telephone Network (PSTN). Voice and video traffic characteristics differ from those of data traffic. The current IP paradigm does not provide performance guarantees or service differentiation. All traffic receives best-effort service; i.e., all packets are treated equally with no regard to the needs of applications for some level of resource assurance or performance guarantees. IP networks are required to offer different grades or classes of service (CoS) with quality-of-service (QoS) guarantees. QoS traffic performance parameters include guaranteed bandwidth availability and upper bounds on packet delay, packet delay variation, and probability of packet loss.

Multi-Protocol Label Switching (MPLS) is a new switching technology that is intended to be integrated into IP networks and overcomes IP's deficiencies. MPLS facilitates traffic engineering by providing the mechanisms needed to control traffic flows in IP networks. It enables the implementation of QoS and enhances restoration in IP networks. This praxis addresses the modeling, analysis, and optimization of MPLS-based IP networks to maximize network resource utilization while meeting customer QoS requirements.

1.2 IP and MPLS Technology Overview

This section provides a high-level overview of the IP and MPLS technologies. It focuses on the areas of these technologies that provide relevant background information. First, an overview of IP technology is given, describing the current routing and forwarding paradigm and its shortcomings. Then, MPLS and its capabilities for implementing traffic engineering and QoS are described.

1.2.1 IP Overview

As mentioned above, the Internet utilizes the TCP/IP protocol suite for communications between devices [37, 38]. This suite defines protocols for the different layers in the protocol stack including the link, network (e.g., IP), transport (e.g., TCP), session, and application (e.g., e-mail, file transfer, or Web) layers. Typically, in an IP network the transport and applications protocols reside on hosts (end-user devices) such as personal computers, workstations, and servers. The network (or internetworking) function is primarily implemented by network nodes referred to as routers. Here, we focus on the IP internetworking paradigm and present the principles of IP routing and forwarding.

IP defines a standard packet format, addressing scheme, and routing and forwarding mechanisms. An IP packet consists of a fixed-size header and variable-size body (or payload). The header consists of several fixed-size fields including the source and destination addresses of the packet. The source address identifies the origin (sending) device. The destination address identifies the destination (receiving) device. IP addresses are assigned to interfaces on hosts and routers so that those devices are uniquely identified.

IP routing is the control-plane function of establishing reachability between network elements (i.e., routers) by exchanging topology and addressing information between these elements. IP defines different routing protocols for this purpose. Most commonly used is the Open Shortest Path First (OSPF) protocol [35]. IP routing is dynamic and distributed across the network; each router runs its own instance of OSPF. Through the exchanges of topology and addressing information with other routers, each router forms its own view of the network topology. Under normal circumstances, the network topology views of the different routers in a network will be synchronized and consistent. Each router considers itself to be the root and performs a shortest-path calculation on its own topology view to determine the shortest-path tree to all other destinations. The number of hops to a destination is the number of edges or links in a path determined by the router algorithm. Each edge or link in such a path is known as a hop. Based on its shortest-path calculations, a router establishes its own routing table. The table consists of one or more entries for each reachable destination. Each entry includes information identifying the next hop on the shortest path tree. The router then forms its own forwarding table, which is a subset of the routing table. When multiple routes exist to a given destination, the router considers predetermined criteria when determining the preferred route to that destination, and includes that route information in the forwarding table.

In addition, IP routing includes mechanisms to deal with network failure conditions such as node or link failure. A failure condition is detected automatically and routers exchange information to update their respective topology views. Subsequently, each router recalculates its shortest-path tree and updates its routing and forwarding tables accordingly.

IP forwarding is the data-plane function of transporting an IP packet from origin to destination. IP packet forwarding is connectionless, meaning there is no end-to-end network connection established between an origin-destination pair. Forwarding is done on a packet-by-packet basis using the destination address in the packet's header. The destination address in the packet remains intact as the packet traverses the network. IP forwarding of a packet is done hop-by-hop. Each router makes its own independent decision in forwarding a packet towards its destination, making use of its own forwarding table to determine the next hop. It is possible that two different packets from the same origin to the same destination may traverse different routes.

IP forwarding does not guarantee in-sequence delivery. Packets from an origin to a destination may not be received in the same order in which they were sent. Furthermore, IP forwarding does not guarantee packet delivery within the network. If a packet traverses a congested link, the packet may be dropped and the network is not responsible for its subsequent recovery. The IP paradigm relies on TCP running on the end-stations to ensure the integrity and in-sequence delivery of packets. A receiver end-station uses TCP to locally reorder miss-sequenced packets. The receiver end-station handles missing packets by requesting retransmission from the sender end-station.

Note also that the source address is not used for the purpose of forwarding the packet. IP forwarding is not necessarily symmetrical. A packet from origin to destination may traverse a particular route, while a return packet from the destination back to the origin may traverse a route, that is not simply that particular route in reverse order. The forward and reverse routes may have none, some, or all nodes and links in common. Finally, the traditional IP paradigm does not offer differentiation between different packets. All packets are treated equally and delivered on a best-effort basis.

The current IP routing and forwarding paradigm has several important deficiencies. IP routing may result in sub-optimal use of network resources and imbalance of traffic load on different links. IP routing considers topology information only in its shortest-path calculations. It does not consider traffic load and resource utilization information. Therefore, the shortest paths from different sources may overlap on some links, causing congestion on those links. This results in an increase in packet delay and a higher probability of packet loss. At the same time, links that are not on a shortest- path tree may remain underutilized even under heavy traffic loads. IP routing provides few controls to influence traffic flow across the network and exploit unused capacity.

As mentioned previously, the current IP forwarding paradigm offers best-effort service to all traffic received. All packets are treated equally regardless of the needs of applications for some levels of resource assurance and performance requirements. The ultimate result of these deficiencies is that performance is unpredictable and service differentiation and performance guarantees cannot be offered.

1.2.2 MPLS Overview

MPLS is an emerging switching technology increasingly being deployed in carriers' IP networks [5, 7, 30, 31, 39]. MPLS intends to support the transport of different protocols, but the focus to date has been on IP. MPLS integrates with, but does not replace, IP networks and helps to overcome much of IP's deficiencies [6, 7, 8, 43]. MPLS exploits IP routing and replaces the IP connectionless forwarding model with a connection-oriented forwarding model. MPLS facilitates traffic engineering by providing the tools and mechanisms to control traffic flow in IP networks [5, 6, 7, 8, 29, 30, 42, 43]. It enables the implementation of QoS and enhances restoration in IP networks.

MPLS defines a standard header format and label-assignment and signaling protocols [5, 39]. It enhances existing IP routing protocols and defines a new forwarding mechanism. MPLS employs a fixed-size header, pre-pending it to an IP packet. The header consists of a label field, a QoS field, and a label-stack-indicator field. The label field is used to switch a packet as it traverses the network. The QoS field is used to specify the required grade of service treatment for the packet. And the label-stack-indicator field allows stacking of labels and the building of some form of hierarchy.

MPLS introduces the following new concepts. A Label Switch Path (LSP) represents a concatenation of network links between two end nodes. A Forwarding Equivalence Class (FEC) is an aggregate of IP traffic between two nodes that requires the same treatment. The treatment can be characterized by QoS and performance parameters. A Label Edge Router (LER) is a router at the edge of the network that handles IP traffic and performs MPLS functions. And a Label Switch Router (LSR) is a transit router that switches MPLS packets. MPLS incorporates the following four key steps in handling an LSP to serve an FEC: path computation, path establishment, path selection, and packet forwarding across an LSP.

Path computation is the process of finding a path between two LERs for a given FEC that meets the treatment requirements for that aggregate. MPLS enhances IP routing protocols (such as OSPF) by including resource information (such as link capacity) and exchanging that information in routing updates. Accordingly, under MPLS a path is calculated based on both topological and resource information; hence, an LSP is no longer constrained to be a shortest path.

MPLS utilizes new signaling protocols to establish a path once it has been computed. The Resource ReSerVation Protocol (RSVP) has been enhanced for that purpose [5]. The source LER of an LSP, which is also referred to as the head-end of the LSP, initiates the setup of the LSP along the computed path. The signaling traverses the LSRs along the computed path towards the destination LER, which is also referred to as the tail-end of the LSP. Each LSR along the path assigns a label to form the LSP. There are different modes of assigning these labels. In a downstream mode, a downstream LSR along a computed path assigns a label and communicates that label to its upstream LSR.

Consequently, each LSR forms an entry in its label-switching table that consists of incoming interface, incoming label, outgoing interface, and outgoing label. At the end of this process, an LSP is formed from the source LER to the destination LER. Note that a label is locally significant and is unique only on a given interface. Also note that an LSP is a unidirectional virtual connection for traffic from the source LER to the destination LER.

Once an LSP is formed, the head-end LER associates a FEC with that LSP and binds the corresponding label. The binding is typically done by configuration and administrative means, but it can also be dynamic. For example, the network administrator may configure a static route to forward a group of packets across a particular LSP. The group of packets may be identified based on different criteria including destination IP address and class-of-service (CoS) setting in the IP header. Also, a FEC may be associated with one or more LSPs with some form of traffic load-sharing among the different LSPs.

Finally, traffic forwarding along an LSP proceeds as follows. The head-end LER processes incoming IP packets and identifies the FEC and associated LSP along which a packet is to be forwarded. The LER forms an MPLS packet by pre-pending an MPLS header to the entire IP packet. It specifies the proper label in the header and sends the MPLS packet across the link to the next LSR along the LSP. Each LSR processes the MPLS header (not the IP header), makes a table lookup based on incoming link and incoming label in the header, replaces that label with the outgoing label and sends the packet on the outgoing interface. The LSR may also process the QoS field in the MPLS header and provide the indicated treatment based on prioritization. As the packet reaches the tail-end, the LER strips off the MPLS header, processes the IP header, and forwards the packet based on the destination IP address.

MPLS technology offers several benefits. It preserves the connectionless property of IP and adds the benefit of a connection-oriented paradigm. MPLS provides the user with control mechanisms to perform traffic engineering and, as a result, facilitates optimized utilization of network resources. MPLS overcomes the limitations of shortest-path-only routing and allows the creation of traffic-engineered paths not necessarily the shortest. Thus, the user can exploit otherwise underutilized resources.

MPLS allows user-defined paths. For a given traffic demand from an origin to a destination node, a path can be established that considers traffic QoS requirements and attempts to meet performance goals. This approach results in more predictable and dependable service. Hence, MPLS facilitates the implementation of classes of service, the delivery of QoS guarantees, traffic engineering, and efficient use of network resources.

Finally, MPLS reduces the complexity of forwarding in IP networks. Processing of the IP header is much more resource intensive than the processing of the MPLS header. MPLS switching is faster than IP processing and forwarding and can be done at higher packet rates. This allows network equipment to support very-high-speed links without performance degradation.

1.3 Drivers for Traffic Engineering and QoS

As mentioned above, handling the explosive traffic growth and enabling convergence present serious challenges to service providers. Offering IP-based services is a competitive business that requires continual investment to keep pace with the increase in traffic demand and to support new applications [7, 8, 24, 25, 36]. Due to the IP deficiencies discussed above, the conventional approach to dealing with traffic growth has been the over-provisioning of costly resources. In the current and future business environment, this answer is not adequate. Carriers are under pressure to contain capital expenditures and are looking for other solutions to maximize their use of network resources.

Promising new solutions require traffic-engineering tools and exploiting the capabilities of MPLS. TE is the process of controlling traffic flow through a network so that network performance and resource utilization are optimized. It improves utilization of network resources by better distributing traffic across the network. TE maximizes the benefits of installed network infrastructure and provides cost savings by avoiding unnecessary expenditures.

The convergence of voice, video, and data traffic imposes new requirements on IP networks, which must support multiple traffic types with dissimilar characteristics and requirements. Voice traffic requires the same predictability and dependability as the existing PSTN. Voice and video traffic characteristics differ from those of data traffic [17]. Voice traffic is very sensitive to delay and delay variations, but can tolerate some degree of packet loss. Data traffic is typically elastic — can tolerate delay — and is oblivious to delay variations, but is sensitive to packet loss. The mix of multiple packet types with dissimilar performance characteristics requires IP networks to offer different grades of service with QoS guarantees. Service differentiation is a competitive requirement and can open new revenue opportunities for service providers.

1.4 Praxis Overview

This praxis addresses the modeling, analysis, and optimization of MPLS-based IP networks to maximize network resource utilization while meeting customer traffic QoS requirements. Specifically, it is concerned with the design of a virtual topology of MPLS LSPs for a given physical network topology. The physical network topology is defined in terms of a set of nodes, links, link capacities, and link costs. Also given are the traffic demand matrix, classes of service and corresponding QoS requirements, revenue per bandwidth and class of service, and penalties (in the form of back-credit to customers) for violating QoS requirements.

Developed herein is a series of network optimization models, each with different motivation, assumptions, and requirements for the designs produced. Chapter 2 formulates the basic traffic-engineering problem using MPLS as a revenue-maximization model. The problem deals with constraint-based routing and the design of paths to route traffic efficiently. The model assumes a single CoS type and that demand between an origin-destination (OD) pair is routed along a single path. The model maximizes revenue by admitting and routing the maximum demand possible while still meeting the resource and QoS constraints. A computational study is conducted to compare the performance of an offline strategy utilizing the optimization model with an online strategy, which implements a first-come-first-served (FCFS) algorithm. The impact of different factors on the performance of the two strategies is investigated. The basic model is enhanced to deal with load balancing using multiple paths per OD pair. The study also evaluates the benefits of load balancing and assesses the impact of different factors on performance.

Chapter 3 enhances the basic MPLS traffic engineering model to deal with multiple CoS types. Each CoS type is assigned a priority, defines its own QoS performance requirements, and is priced differently. A computational study for the case of two CoS types performs revenue analysis and assists in determining the relative increase in revenue per unit of demand of the higher-class traffic.

Chapter 4 deals with the problem of MPLS traffic engineering with over-subscription of link capacities. Over-subscription allows the model to admit demand that otherwise would have been rejected but a penalty may be assessed on traffic that exceeds link capacities. Over-subscription is intended to exploit fluctuations in demand and provide a statistical-multiplexing gain. The problem is formulated using a model that maximizes revenue and minimizes penalty. A computational study is performed to demonstrate the usability of the model for the purpose of capacity planning to accommodate traffic growth.

1.4.1 Illustration of the Problem

The following network example illustrates the basic problem addressed and highlights the advantage of MPLS LSPs over the conventional shortest-path IP routing. Figure 1-1 illustrates the network topology. The network consists of eight nodes. Nodes 1-4 are the "edge nodes" and are considered as source/sink nodes. Nodes 5-8 are considered as "backbone nodes" and act as transshipment nodes. All links are bi-directional with a capacity of 10 units of bandwidth. An administrative cost[1] or distance (which may represent delay) is associated with each link. The shortest-path algorithm uses these costs to calculate the shortest path from a given source.

[1] Typically, the network engineer operating the network determines and assigns the administrative costs based on the actual distance of the physical links.

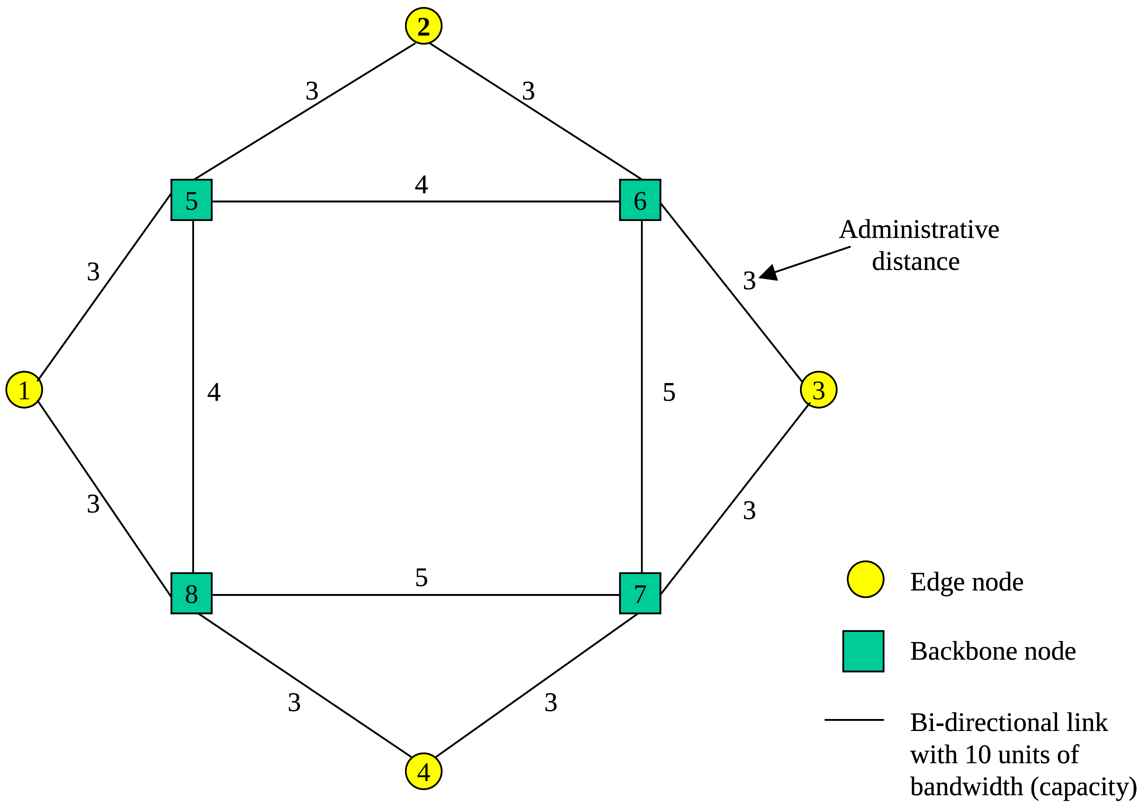

Figure 1-1 Topology of example network

Table 1-1 presents the traffic matrix for a single class of traffic. As mentioned above, for the purpose of this example only, nodes 1-4 are the origination/destination nodes. Furthermore, for illustration purposes and simplicity, the traffic is symmetric.

Table 1-1 Traffic matrix (one class of traffic)

To node \ From node	1	2	3	4
1			10	10
2			10	10
3	10	10		
4	10	10		

Figure 1-2 illustrates the routing of the demand from Table 1-1, based on shortest-path calculations. Note that on the links (4,8) and (3,6) the total load exceeds the link capacity. In that case the excess traffic is dropped (discarded) and not delivered across that link. Only 50% of the packets are delivered and some links are not utilized. However, this configuration provides the best delay performance; the weighted average delay is 8 units, and the minimum and maximum delays are 6 and 10 units, respectively. Furthermore, this configuration results in the fewest number of hops (links traversed) for the routed traffic; the weighted average of the number of hops is 2.5, and the minimum and maximum number of hops are 2 and 3, respectively.

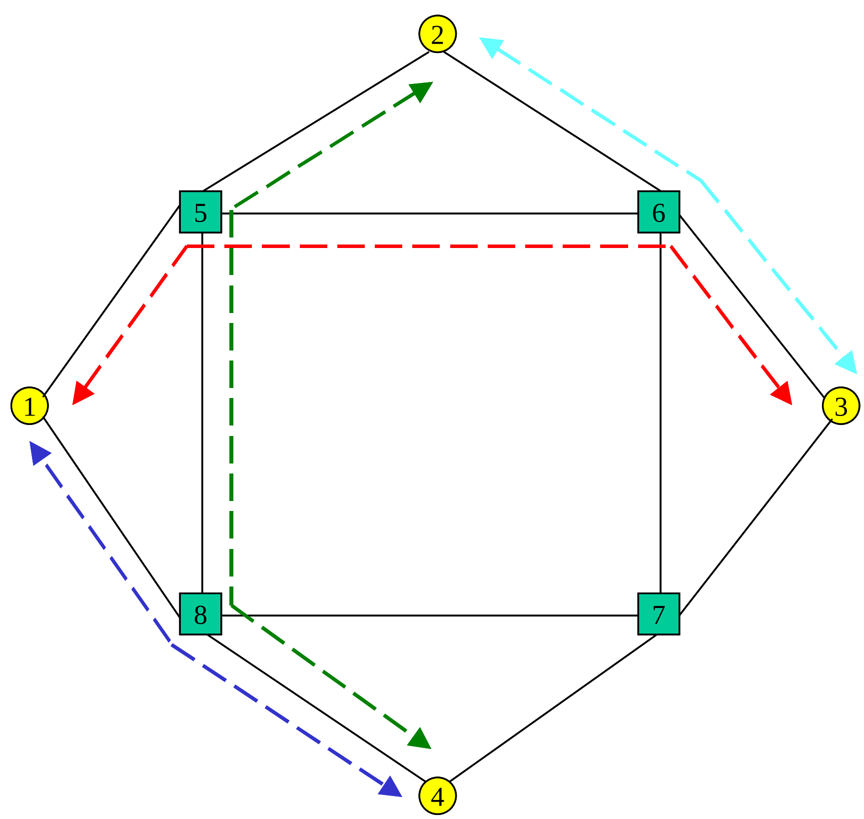

Figure 1-2 IP Shortest-Path Configuration

Figure 1-3 illustrates a possible LSP configuration that is intended to minimize packet loss for the same traffic matrix. A 100% packet delivery is achieved, but the delay performance is worse than that of the shortest-path configuration. The weighted average delay is 10.5 units, and the minimum and maximum delays are 10 and 11 units, respectively. The weighted average, and the minimum and maximum number of hops are all 3.

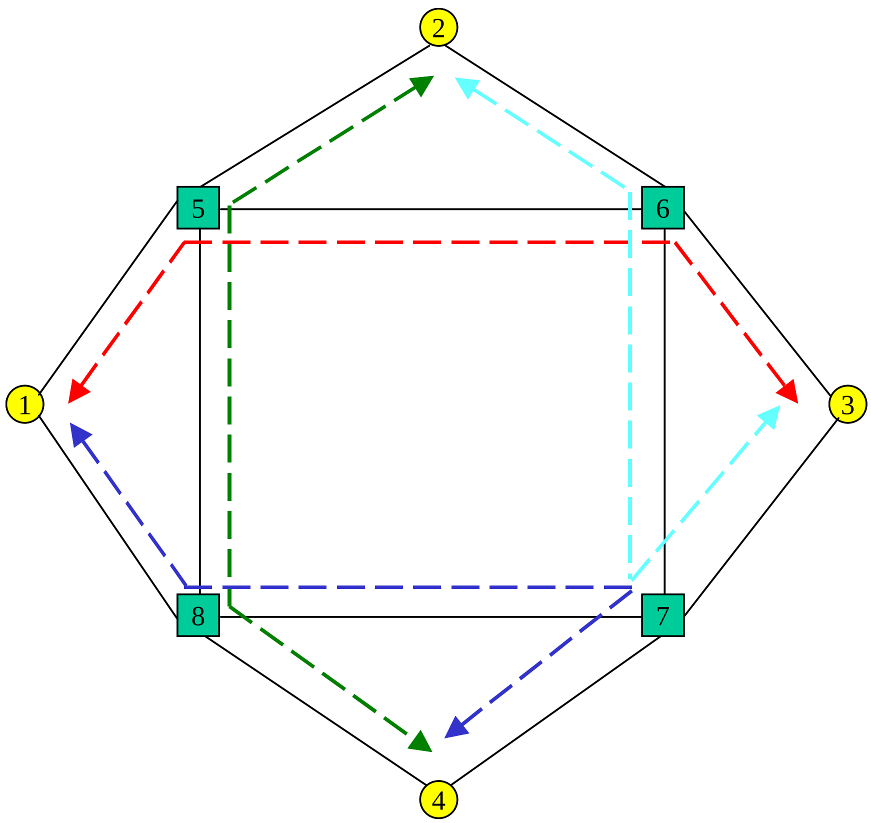

Figure 1-3 MPLS path configuration for one class of traffic

Table 1-2 presents a new traffic matrix with the total traffic load as before but with two different classes of service, A and B. The table shows that there are 5 units of A-class traffic and 5 units of B-class traffic between the respective OD pairs.

Table 1-2 Traffic matrix (class A/class B)

		From node			
		1	**2**	**3**	**4**
	1			5/5	5/5
To	**2**			5/5	5/5
node	**3**	5/5	5/5		
	4	5/5	5/5		

Figure 1-4 depicts an LSP configuration that is intended to minimize packet loss while providing preferential treatment to the A-class traffic. The figure illustrates two sets of LSPs, one to serve the A-class traffic and the other to serve the B-class traffic. Again, 100% packet delivery is achieved. The A-class delay performance is the same as that of the shortest-path delay performance, but the B-class delay performance is even worse than that of the single-class configuration. For the B-class traffic, the weighted average delay is 13 units, and the minimum and maximum delays are 11 and 15 units, respectively. The weighted average of the number of hops is 3.5, and the minimum and maximum number of hops are 3 and 4, respectively.

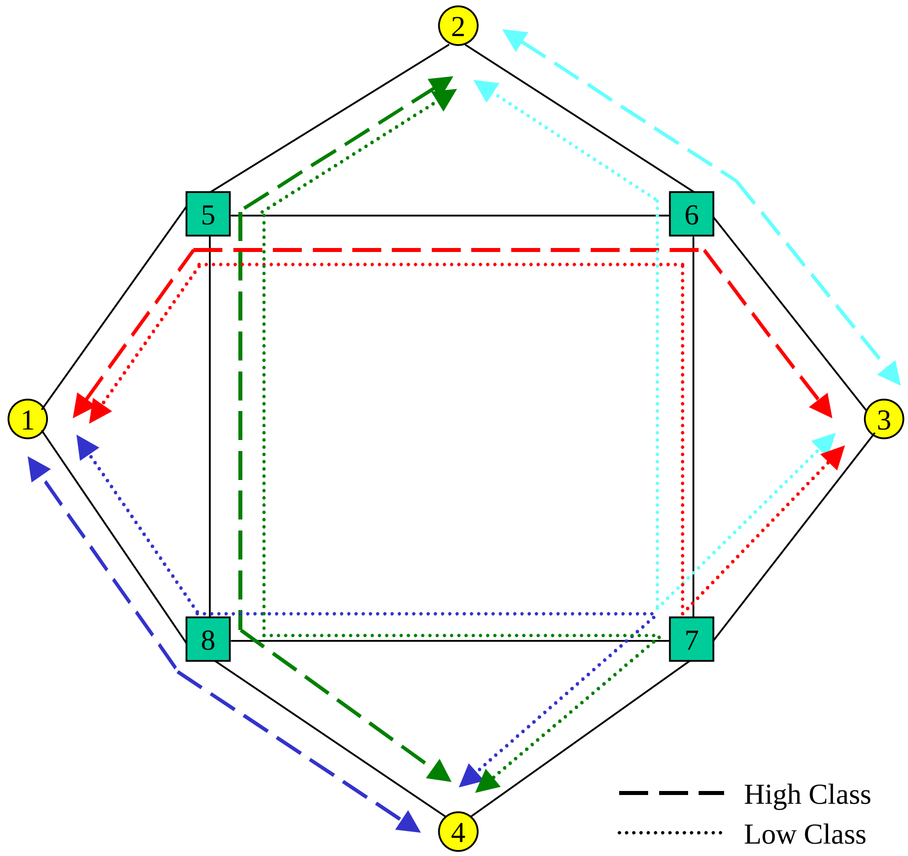

Figure 1-4 MPLS path configuration for two classes of traffic

Table 1-3 summarizes the results of the three configurations. The table shows that both (the single-class and the two-class) LSP configurations delivered 100% of the traffic while the shortest-path configuration delivered only 50% of the traffic. The improvement in packet delivery is achieved at the expense of higher delay. The average delay with the single-class LSP configuration is 10.5 units compared to 8 units with the shortest-path configuration. In the case of the two-class LSP configuration, the average delay for the A-class traffic is the same as that of the shortest-path. However, the B-class traffic suffers from the worst delay with an average of 13 units.

Table 1-3 Performance of the shortest-path and MPLS configurations

	Shortest-path	LSP, single class	LSP, two classes	
			Class A	Class B
Packet delivery	50%	100%	100%	100%
Min. delay	6	10	6	11
Avg. delay	8	10.5	8	13
Max. delay	10	11	10	15
Min. hops	2	3	2	3
Avg. hops	2.5	3	2.5	3.5
Max. hops	3	3	3	4

1.4.2 Approach and Methodology

Network flow models and algorithms provide the theoretical basis for the topic of this praxis [1, 12]. The main areas of relevance include shortest-path tree calculations, and formulation of capacitated multi-commodity flow networks with side constraints.

Several models using linear-programming (LP) and mixed-integer-programming (MIP) formulations are developed. The models are implemented using the Generalized Algebraic Modeling System (GAMS) and generated instances solved with the CPLEX optimization software package. The models are applied to a suite of network instances that mirror realistic MPLS design problems. The computational experiments are conducted using statistical experiment design and the test results are evaluated using the analysis of variance (ANOVA) technique.

1.4.3 Significance and Contributions

The praxis deals with a new technology that is expected to have a far-reaching impact on the data/telecommunications industry. The optimization of packet-network performance is becoming a critical component of competitiveness in the current telecommunications business environment. MPLS will play a major role in this effort. While MPLS is still evolving, it is expected to be widely deployed. MCI has been implementing MPLS in its Internet backbone and its private networks for next-generation IP-based data services.

Limited analytical work and research has been done in this area and few tools are available for planning and design activities. This research provides additional insight into the design aspects of MPLS-based networks, identifies factors that impact performance, evaluates benefits, and provides practical guidelines and recommendations for network administrators and managers engaged in the deployment of MPLS-based networks. Network designers and capacity planners can exploit the models and results of this research to optimize network resource utilization and maximize revenues for their companies.

Chapter 2

FORMULATION AND BENEFIT EVALUATION OF MPLS TRAFFIC ENGINEERING

This chapter presents new formulations and results of a computational study of two fundamental traffic-engineering problems using Multi-Protocol Label Switching (MPLS) in IP-based networks. The problems deal with constraint-based routing and the design of paths to route traffic efficiently and are referred to as TE1 and TE2. TE1 involves the routing of traffic along a single path per origin-destination (OD) pair. TE2 extends TE1 to load balancing using multiple paths per OD pair.

The problems are formulated as OD integer multi-commodity network flow problems with side constraints. This formulation is based on realistic assumptions and considers actual design practices in service provider networks.

Included are the results of computational studies that evaluate the benefits of traffic engineering. The first study compares the performance of an offline optimization model for problem TE1 with an online strategy. The second study evaluates the benefits of load balancing using an optimization model for problem TE2.

This chapter is organized as follows. Section 1 is an introduction and includes motivation for the research topic, a statement of the problem, and a survey of related literature. Section 2 includes mathematical formulations of the models for TE1 and TE2. Section 3 presents the methodology for the computational experiments, the tabulated results, and the results of hypothesis testing regarding the impact of different factors on the performance of the problem solutions. Section 4 summarizes the results and provides conclusions and recommendations.

2.1 Motivation

The growth and popularity of the public Internet has accelerated the adoption of the Internet Protocol (IP) as a dominant communications technology. Handling the explosive traffic growth on the Internet presents serious challenges to the IP technology and service providers. Both Internet and IP-based virtual private network (VPN) services are competitive businesses that require continual investment to keep pace with the increase in traffic.

The current IP routing-and-forwarding paradigm has several important deficiencies. IP routing may result in sub-optimal use of network resources and an imbalance of traffic load on different links because it considers topology information only in its shortest-path calculations. It does not consider traffic load and resource-utilization information. Moreover, IP routing provides few controls to influence traffic flow across the network and exploit unused capacity.

Due to IP deficiencies, the conventional answer to dealing with traffic growth has been network over-provisioning with costly resources. In the current and future business environments this answer is not adequate. Carriers are under pressure to contain capital expenditures and are seeking other solutions to maximize the use of existing network resources. The Internet Engineering Task Force (IETF) introduced and defined the architecture of MPLS [39] and defined requirements for Traffic Engineering (TE) [7] over MPLS. MPLS is a new switching technology that is integrated into IP networks and overcomes IP's deficiencies. TE is the process of controlling traffic flow through a network so that network performance and resource utilization are optimized. TE maximizes the benefits of an installed network infrastructure. MPLS facilitates traffic engineering by providing the mechanisms needed to control traffic flows in IP networks. MPLS overcomes the limitations of shortest-path-only routing and allows the creation of traffic-engineered paths that may not be simply shortest path. Thus, the user can exploit otherwise underutilized resources.

2.1.1 Problem Statements

TE1 can be stated as follows. The physical topology and link attributes of an MPLS network are given. The link attributes include capacity and an assigned administrative "cost" that reflects delay on the link. Also given is the traffic matrix, which represents aggregate traffic demand between each OD pair. The objective is to maximize network revenue by admitting and routing as much traffic as possible while observing the resource and traffic performance constraints. The resource constraints are the link capacities. The traffic performance constraints are typically expressed as the maximum number of hops and delay allowed for traffic between any OD pair. Hence, TE1 is a logical design problem that involves constructing a set of paths to route the traffic, with each OD pair's traffic routed along a single path. Also embedded in TE1 is a traffic admission control problem; in case not all traffic can be routed, the solution identifies the set of OD pairs whose traffic can be routed.

TE2 adds an additional dimension to TE1 by allowing each OD pair's traffic to be routed along multiple paths. Thus demand associated with an OD pair can be split according to predetermined ratio(s) and routed along multiple paths to achieve load balancing.

2.1.2 Survey of Related Literature

The Internet Engineering Task Force (IETF) has developed a series of requests for comments (RFCs) specifications that cover various aspects of MPLS and TE. A few important RFCs are briefly summarized here. RFC 3031 [39] specifies the architecture for multi-protocol label switching, which integrates the label-switching paradigm with the network layer routing. It defines the functions performed by an MPLS-capable router, which is referred to as a label switch router (LSR). The RFC defines the encoding of labels, operations on labels (pop, push, and swap) and procedures and modes of operation for label distribution protocols. The document specifies both the control and traffic-forwarding functions. The control functions consist of partitioning traffic to forwarding equivalence classes (FECs), assigning and binding labels to FECs, distributing labels among LSRs, and establishing label switched path (LSPs)[1], across which packets associated with FECs are transported. The forwarding function consists of a set of operations on labels that need to be performed by ingress, transit, or egress LSRs to switch the traffic along an LSP. The document also defines the data entries and mapping required by an LSR to forward labeled and unlabeled packets. Finally, it describes applications for MPLS and integration with IP routing including hop-by-hop routing, explicitly routed LSPs, and multi-path routing.

RFC 2702 presents a set of requirements for traffic engineering over MPLS. "It identifies the functional capabilities required to implement policies that facilitate efficient

[1] An LSP is interchangeably referred to as an LSP tunnel or just tunnel.

and reliable network operation in an MPLS domain" [7]. Specifically, these capabilities can be used to maximize bandwidth utilization and enhance traffic performance. The document utilizes the concept of a traffic trunk, which is an abstraction representing an aggregation of traffic flows of the same class mapped to an LSP. The set of capabilities consists of assigning attributes to traffic trunks (such as traffic rate, priority, and resiliency), assigning attributes to resources (such as link bandwidth and resource class), which are considered as topology attribute constraints, and utilize a constraint-based routing framework to select paths for traffic trunks considering both trunk and resource attributes.

RFC 3209 [5] specifies extensions to the Resource Reservation Protocol (RSVP) and describes its use as a label distribution and signaling protocol to establish explicitly routed LSP tunnels in MPLS networks. Explicitly routed LSPs are used for traffic-engineering purposes. The document specifies new objects such as the label and the explicit-route objects, defines message format, and describes procedures for processing signaling messages. The protocol defines the following features supprting the operation of LSP tunnels: performing down-stream-on-demand label allocation, distribution, and binding, establishing LSP tunnels with or without quality-of-service (QoS) requirements, dynamic rerouting of established tunnels, recording and observing the actual route traversed, diagnosing tunnels, and preempting established tunnels based on some policy.

RFC 3630 [29] describes extensions to the Open Shortest Path First (OSPF) routing protocol [35] to support intra-area traffic engineering. The extensions enhance the protocol to support resource-based routing in addition to topology-based routing. For this purpose, new link attributes (such as maximum bandwidth, maximum reservable bandwidth, and administrative group) are added and advertised using opaque link state advertisements. An extended link-state database (referred to as a traffic-engineering database) is built, which enables a router to perform constraint-based source routing.

Other related articles [6, 8, 43] discuss the applications and benefits of MPLS for traffic engineering in IP networks. The papers review the architectural components of MPLS, contrast MPLS with traditional IP routing, and describe the value of MPLS control mechanisms for establishing user-specified label-switched paths. The papers also describe how MPLS can increase network reliability by establishing backup paths and fast rerouting facilities.

Girish et al. [22] formulate multiple problems related to traffic engineering in MPLS-based networks using point-to-point (p-t-p) LSPs. They formulate the fundamental constraint-based routing as a mixed-integer program (MIP) that minimizes total cost. The formulation includes the number of hops as a constraint and can result in infeasible instances. They formulate separately the connection-admission-control problem, which determines whether or not to admit an LSP and, if admitted, determines the path for that LSP. They also formulate the rerouting problem, which deals with rerouting LSPs in case of link failure, and a capacity-planning model, which determines optimal link capacities.

Saito et al. [40] present a traffic-engineering scheme using multiple multipoint-to-point (m-t-p) LSPs in MPLS networks. The scheme consists of first designing the m-t-p LSPs and then assigning traffic demand (or flows) to those LSPs. Designing the m-t-p LSPs consists of two steps. First route selection is performed in which multiple diverse routes between an OD pair are selected. Then the m-t-p LSP design is formulated as an integer-program model that minimizes the number of m-t-p LSPs. The set of m-t-p LSPs that is selected includes all the selected routes and offers node diversity and load balancing for each OD pair. The m-t-p LSP set is determined based on topology information and does not consider link bandwidth. The flow assignment problem is formulated as an MIP problem that minimizes the maximum link load. Their computational results indicate that the number of LSPs and labels required for each link is considerably less than those required by a p-t-p LSP design. Furthermore, link load is reduced in comparison with shortest-path based routing.

Kar and Lakshman [28] present an online path-selection algorithm that establishes tunnels as demand requests arrive one-by-one and without a priori knowledge of future requests. Their minimum interference routing algorithm (MIRA) routes a new demand along the path that has minimum expected interference while satisfying anticipated future requests. It identifies “critical” links as those that, if heavily loaded, would result in rejecting some future demands. The algorithm exploits available ingress-egress pair information, even if the future demand is not known. Simulation results show that MIRA outperforms other algorithms based on minimum-hop routing in terms of LSP acceptance and rerouting around failed links. Aukia et al. [4] describe a software system called Routing and Traffic Engineering Server (RATES) that utilizes MIRA. The software combines a policy and flow database with a web-based interface for policy definition and demand requests. The server calculates paths and communicates with ingress nodes to initiate the establishment of LSPs.

Banerjee and Sidhu [10] present two online path-selection algorithms that consider both bandwidth and delay constraints. In addition, the algorithms consider metrics of max-flow reduction, path cost, and path load and exploit the concept of "critical" links. The algorithms attempt to achieve the multiple objectives of increasing network revenue, limiting network cost, and distributing network load. The algorithms are compared with other competitive approaches, including MIRA. Simulation results show that the authors' algorithms provide the most-favorable or close to the most-favorable performance with respect to multiple measured metrics.

Elwalid et al. [19] describe an online class of algorithms called Adaptive Multi-path Traffic Engineering (MATE). The objective of MATE is to reduce network congestion by adaptively balancing the load among multiple LSPs between an ingress-egress pair based on actual measurement of link load and traffic performance. The traffic engineering process is an iterative process of monitoring and performing load balancing. Network congestion is monitored by sending probe packets and collecting LSP statistics such as delay and packet loss. The optimality and stability of MATE are demonstrated and simulation results show that it removes traffic imbalances among LSPs and significantly reduces packet loss.

Fortz and Thorup [21] and Fortz et al. [20] argue that the traditional shortest-path routing protocols such as OSPF can be used effectively for traffic engineering by dynamically adjusting link weights based on a network-wide view of traffic and topology. The papers describe an approach for monitoring the traffic and topology, optimizing the setting of the static weights, and reconfiguring the routers with new weights when needed. The papers show that for a proposed AT&T IP backbone (with 90 nodes and 274 links) with projected traffic demands, adjusting the weights can result in maximum link utilization that is only three percent higher than that obtained by traffic engineering based on MPLS.

Barr and McLoud [33] present a new heuristic algorithm for solving the Bandwidth Packing Problem (BWP), which is applicable to the MPLS traffic-engineering problem formulated in 2.2.2 as an offline MIP optimization problem. The Invisible Hand Heuristic (IHH) uses an iterative process for routing commodities (traffic demands), where each demand acts in its own self-interest and repeatedly attempts to minimize the cost of its route. As commodities are routed, link costs are updated to reflect total flow and available capacity on affected links. Link cost is increased as the load on the link increases. The process continues until equilibrium is reached and all commodities are satisfied with their respective route costs. The algorithm has polynomial asymptotic bounds for both space and time. Extensive computational testing shows that the algorithm is very effective in solving large-scale problem instances providing near-optimal solution in extremely short algorithmic running times.

2.2 Mathematical Formulations

This section presents mathematical formulations of problems TE1 and TE2. Both are OD integer multi-commodity network flow problems with side constraints.

2.2.1 Notation and Conventions

The following notation and conventions are used. Italicized upper-case letters, such as S, denote sets. The cardinality of a set S is denoted as $|S|$. The sets of real numbers and integers are denoted as R and Z, respectively. Bold lower-case letters, such as $\mathbf{b}$, represent vectors. An MPLS network topology is represented as a directed graph $G = (N, A)$, where N is the unordered set of nodes (vertices), and A is the set of directed arcs. A directed arc is an ordered pair (i, j) where $i, j \in N$. A node represents an MPLS LSR. The terms *node*, *LSR*, and *router* are used interchangeably. A physical transmission line or a trunk connecting two nodes is represented by two directed arcs, one in each direction. The terms *directed arc*, *arc*, and *link* are used interchangeably. The *capacity* of a link represents the bandwidth or the transmission speed of that link measured in units of bandwidth, such as mega-bits per second (Mbps). The "cost" of a link is an associated traffic-performance metric, such as delay, and not a monetary cost. A path from a node a to a node $b \neq a$ consists of a sequence (n_0, n_1, n_2, …, n_k) of distinct nodes such that $n_0 = a$, $n_k = b$, and $(n_{i-1}, n_i) \in A$ for i = 1, 2, …, k. The terms *path* and *LSP* are used interchangeably. A commodity represents distinct packet traffic to be routed from a source node to a destination node. The demand associated with a commodity is the data rate or bandwidth (measured in Mbps) consumed by that traffic. The terms *source* and *origin* are used interchangeably. The terms *demand* and *traffic demand* are used interchangeably

Using the above notation, the following symbols are used in the formulation of TE1:

N the set of node indices 1, 2, …, |N| in the network

A the set of directed arcs in the network; an arc is represented as an ordered pair (i, j) where $i, j \in N$.

b_{ij} the capacity, in units of bandwidth, of arc (i, j), $b_{ij} \in R$

c_{ij} the administrative cost associated with arc (i, j), $c_{ij} \in R$; typically the administrative cost represents a measure of delay or transmission time

K the set of commodities or OD demands to be routed

O_k the origin (source) node of commodity $k \in K$

D_k the destination node of commodity $k \in K$

d_k the demand of commodity $k \in K$ in units of bandwidth, $d_k \in R$

ℓ the maximum allowed delay (latency) that any commodity may incur while traversing the network from source to destination[2], $\ell \in R$

h the maximum allowed number of hops that any commodity may traverse from source to destination[3], $h \in Z$

μ a unit of revenue generated from delivering a unit of demand of any commodity, $\mu \in R$

ω a scaling factor or a weight used in the objective function, $\omega \in R$

2.2.2 Basic MPLS Traffic Engineering Single-Path Model (TE1)

Using the above notation, the basic problem of MPLS traffic engineering (TE1), which involves constraint-based routing and admission control, can be stated as follows. Given is a graph $G = (N, A, \mathbf{b}, \mathbf{c})$ that describes the physical topology and link attributes of an MPLS network. Also given is the node-to-node traffic represented by the set of commodities K. The objectives are to maximize revenue and minimize delay while meeting the resource and traffic performance constraints. The resource constraint is enforced by not allowing the combined traffic across all commodities on any link to exceed the capacity (b_{ij}) of that link. The performance constraints are expressed in terms

[2] This formulation can be generalized to associate separate maximum delay with each commodity.
[3] This formulation can be generalized to associate separate maximum number of hops with each commodity

of maximum delay and number of hops. Each commodity must be routed along a path that does not exceed the maximum delay and number of hops. The following decision variables are defined:

$$y^k = \begin{cases} 1, & \text{if commodity } k \text{ is routed,} \\ 0, & \text{otherwise,} \end{cases}$$

$$x_{ij}^k = \begin{cases} 1, & \text{if commodity } k \text{ is routed through a path that uses arc } (i, j), \\ 0, & \text{otherwise.} \end{cases}$$

The single-path problem can be stated in preemptive priority form as follows:

[TE0]

$$\text{Maximize } P_1\Big(\sum_{k\in K} \mu d_k y^k\Big) + P_2\Big(-\sum_{(i,j)\in A} c_{ij} \sum_{k\in K} d_k x_{ij}^k\Big) \tag{2.1}$$

subject to,

$$\sum_{j\in N|(i,j)\in A} x_{ij}^k - \sum_{j\in N|(i,j)\in A} x_{ji}^k = \begin{Bmatrix} y^k, & \mathit{if}\ i = O_k \\ -y^k, & \mathit{if}\ i = D_k \\ 0, & \mathit{otherwise} \end{Bmatrix} \quad \forall k\in K, \forall i \in N, \tag{2.2}$$

$$\sum_{k\in K} d_k x_{ij}^k \le b_{ij} \quad \forall (i, j) \in A, \tag{2.3}$$

$$\sum_{(i,j)\in A} c_{ij} x_{ij}^k \le \ell \quad \forall k \in K, \tag{2.4}$$

$$\sum_{(i,j)\in A} x_{ij}^k \le h \quad \forall k \in K, \tag{2.5}$$

$$x_{ij}^k, y^k \in \{0,1\} \quad \forall (i, j) \in A, \forall k \in K. \tag{2.6}$$

where P_i represents a priority of the corresponding objective function term, with 1 being the highest priority. Alternatively, the objective function can be represented using a linear combination of the two objectives and the model can be expressed as follows:

[TE1]

$$\text{Maximize} \sum_{k\in K} \mu d_k y^k - \omega \sum_{(i,j)\in A} c_{ij} \sum_{k\in K} d_k x_{ij}^k \tag{2.7}$$

Subject to (2.2) – (2.6).

In this formulation, the objective function consists of two terms with the first one being the primary objective and the dominant term. The first term represents the total revenue generated from routed commodities (i.e., delivered traffic). If a commodity is routed, its entire demand is delivered. The total demand delivered, and not necessarily the total number of commodities delivered, is maximized. The second term represents the total delay incurred by all the delivered commodities. The purpose of this term is to select the solution with the lowest delay among multiple alternate optimum solutions (yielding the same revenue) that may exist. The total delay is multiplied by the scaling factor ω, $0 \leq \omega \leq 1$. Typically, ω is set so that the second term will be small relative to the first (dominant) term.

Constraints (2.2) are the flow-conservation equations, which ensure a connected path for each routed commodity. The equations determine whether a commodity is routed and the set of arcs that comprise the path for that commodity. For each commodity, there are $|N|$ equations: one equation for the source node of that commodity, one for the destination node, and $|N|$ - 2 equations for the other nodes in the network, which may act as transit nodes. Flow for commodity $k \in K$ is routed only if binary variable $y^k = 1$; otherwise that particular demand is not satisfied. Unit supplies and demands, and binary flow variables cause a single LSP to be formed via these constraints.

Constraint set (2.3) enforces the arc capacity resource constraints. For each arc, the total traffic from all commodities whose paths include that arc cannot exceed the arc's bandwidth. Constraints sets (2.4) and (2.5) represent the performance requirements.

Constraint set (2.4) ensures that the delay along any path cannot exceed a predetermined upper delay limit ℓ and (2.5) ensures that the number of hops along any path cannot exceed a predetermined upper hop limit *h*. Constraint set (2.6) enforces the integrality condition on the binary decision variables.

The above formulation assumes that a single path is used for each commodity. If a commodity is routed, its entire demand will be delivered along that path. The solution determines which commodities will be routed and the path for each commodity.

2.2.3 MPLS Traffic Engineering with Load Balancing Model (TE2)

The formulation for TE1 is expanded for TE2 to allow for multiple paths to be used for each commodity. The demand of a commodity may be split based on assigned ratios among the different paths. All or part of each demand may be delivered. To formulate this problem we expand the above notation and define the following:

n the maximum number of paths for each commodity

p an index of a path for a commodity, where $p = 1, 2, \ldots, n$

K_n the set of new commodities that result from splitting each of the original commodities into n commodities, each of which can take a separate path; each commodity $k_p \in K_n$ is represented by the following three attributes

O_{k_p} the origin (source) node of commodity $k_p \in K_n$

D_{k_p} the destination node of commodity $k_p \in K_n$

d_{k_p} the demand of commodity $k_p \in K_n$ in units of bandwidth

r_p the ratio of the demand of any commodity $k \in K$ that may be routed on the p^{th} path for that commodity

Using the above notation the following relationships are observed:

$$\sum_{p=1}^{n} r_p = 1 \qquad \text{where } r_p \geq 0 \tag{2.8}$$

$$d_{k_p} = r_p d_k \qquad \forall k \in K, p = 1,2,\ldots,n \tag{2.9}$$

$$d_k = \sum_{p=1}^{n} d_{k_p} \qquad \forall k \in K \tag{2.10}$$

The following decision variables are defined for the multi-path model:

$$y^{k_p} = \begin{cases} 1, & \text{if commodity } k_p \text{ is routed,} \\ 0, & \text{otherwise,} \end{cases}$$

$$x_{ij}^{k_p} = \begin{cases} 1, & \text{if commodity } k_p \text{ is routed through a path that uses arc } (i,j), \\ 0, & \text{otherwise.} \end{cases}$$

The multi-path problem can be formulated as follows:

[TE2]

$$\text{Maximize} \sum_{k_p \in K_p} \mu d_{k_p} y^{k_p} - \omega \sum_{(i,j)\in A} c_{ij} \sum_{k_p \in K_p} d_{k\ p} x_{ij}^{k_p} \tag{2.11}$$

subject to,

$$\sum_{j\in N|(i,j)\in A} x_{ij}^{k} - \sum_{j\in N|(i,j)\in A} x_{ji}^{k_p} = \begin{Bmatrix} y^{k_p}, & if\ i = O_{k_p} \\ -y^{k_p}, & if\ i = D_{k_p} \\ 0\ , & otherwise \end{Bmatrix} \quad \forall k_p \in K_n, \forall i \in N, \tag{2.12}$$

$$\sum_{k_p \in K_p} d_{k_p} x_{ij}^{k_p} \leq b_{ij} \qquad \forall (i,j) \in A, \tag{2.13}$$

$$\sum_{(i,j)\in A} c_{ij} x_{ij}^{k_p} \leq \ell \qquad \forall k_p \in K_n, \tag{2.14}$$

$$\sum_{(i,j)\in A} x_{ij}^{k_p} \leq h \qquad \forall k_p \in K_n, \tag{2.15}$$

$$x_{ij}^{k_p}, y^{k_p} \in \{0,1\} \qquad \forall (i,j) \in A, \forall k_p \in K_n. \tag{2.16}$$

To explore the effectiveness of these models in addressing practical problems from the telecommunications industry, a series of computational experiments are performed. A suite of network instances, with varying topologies that mirror realistic MPLS design problems, are constructed and tested on a wide range of parameter values. Models TE1 and TE2 are then applied to create optimal routing solutions. The results are evaluated to test a series of research hypotheses and develop insight into effectively engineered MPLS networks. The following section describes this work.

2.3 Computational Experiments

Computational experiments are performed to conduct two separate comparative-benefit studies. The studies are concerned with solution-quality metrics and not model solution time. The first study compares the performance of two different strategies for designing LSPs in an MPLS network. The first strategy is an offline optimization strategy based on model TE1. The second strategy is an online algorithm with first-come-first-served (FCFS) strategy. The impact of different factors on the performance of the two strategies is investigated. The factors include: the number of OD pairs, the average demand per OD pair (denoted as $\bar{d}$), the demand range, the network topology, and the average node degree.

The second study evaluates the benefits of load balancing in MPLS networks. It compares the performance of single-LSP-per-OD-pair scheme based on model TE1 against the two-LSP-per-OD-pair scheme based on model TE2. The impact of different factors on the performance of the two designs is investigated. The factors include the number of OD pairs, the average demand per OD pair, and the network topology.

2.3.1 Organization of Tests

2.3.1.1 Test Networks Generation and Characteristics

The tests are performed on one realistic network and a number of randomly generated networks. The characteristics of the test networks are summarized in Table 2-1. The realistic network, referred to as N0, represents a typical topology of nationwide data communications network implemented by inter-exchange providers in the US. Network N0 is depicted in Figure 2-1. N0 consists of 20 nodes and 62 arcs with an average node degree of 3.1. The trunks connecting the nodes are bi-directional and full duplex, so that each trunk is represented as two directed arcs, each with the same capacity and cost. The thick trunks represent OC-48 transmission lines with 2488 Mbps of bandwidth capacity, and the thin trunks represent OC-12 transmission lines with 622 Mbps of bandwidth capacity. The arc cost reflects the actual circuit mileage of the corresponding transmission line. Since the shortest-path algorithm uses the arc cost as the metric to calculate the shortest path the OC-12 cost is increased by 1000 units so that an OC-48 trunk will be preferable to an OC-12 trunk.

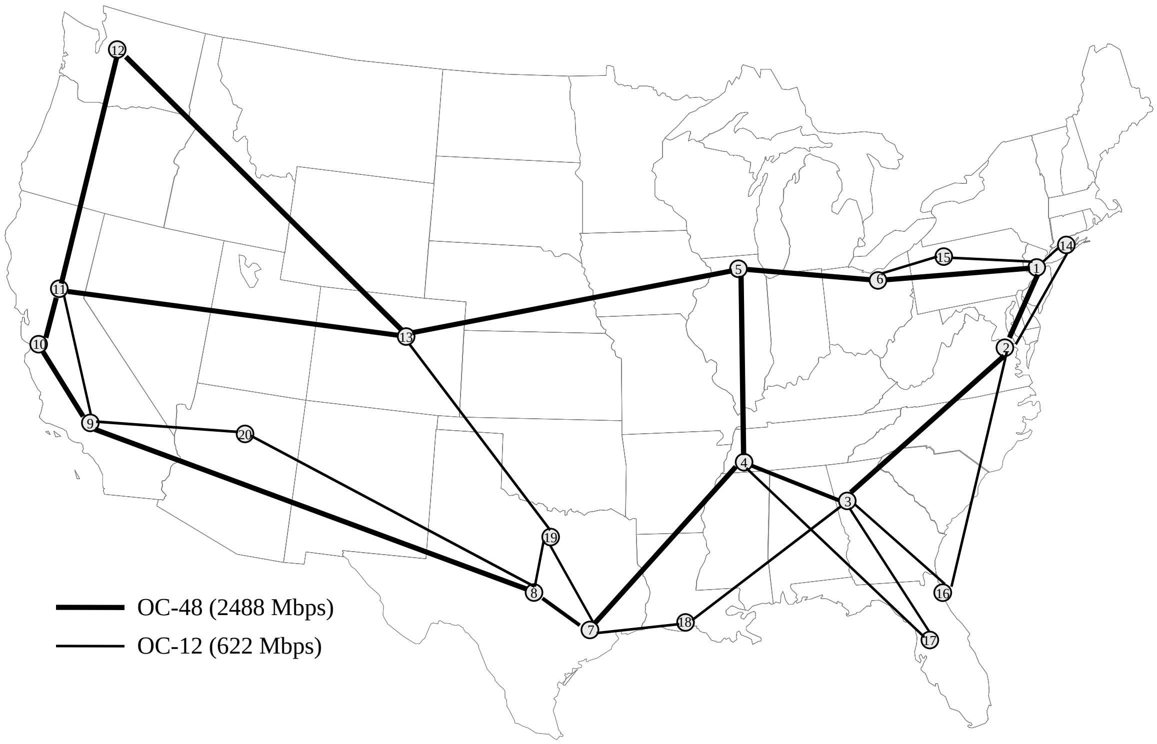

Figure 2-1 Network topology of realistic test network N0

Table 2-1 Experimental networks characteristics

Network(s)	Number of Nodes	Number of Arcs	Average Degree	Arc Capacity (Mbps)
N0	20	62	~3	2488 (30) 622 (32)
N1-N5	20	80	4	2488
N6-N10	20	160	8	1244

In addition to the realistic network, tests are performed on a number of randomly generated networks. The networks are generated using the RGEN problem generator developed by McLoud [33]. RGEN accepts as input parameters the number of nodes, number of arcs, arc cost range, and network topology type. (RGEN accepts other parameters but they are irrelevant for our purpose, which is merely generating network topologies.) All networks are generated with a mesh topology, as follows. RGEN starts by randomly selecting the position of nodes in a three-dimensional space before determining the set of arcs. In a mesh topology, the probability of an arc connecting a pair of nodes is inversely proportional to the distance between the two nodes raised to a certain power. Arc costs are set to the calculated distance and then scaled so that the minimum and 90^{th} percentile costs match the user-supplied input parameters. RGEN ensures that all networks are fully connected (a path exists between any two nodes in a network) and that there are no parallel arcs between any pair of nodes. Two sets of networks are generated, each consisting of five network instances with similar characteristics but differing topologies. The two sets allow examination of the impact of network topology and density on performance. Each network (N1-N5) in the first set consists of 20 nodes and 80 arcs, and with fixed arc capacity of 2488 Mbps. Each network (N6-N10) in the second set consists of 20 nodes and 160 arcs. For some of the experiments, an arc capacity of 1244 Mbps is assigned and for other experiments an arc capacity of 2488 Mbps is assigned. The choice of arc capacity depends on the purpose of the test and is explained below when the results are presented.

2.3.1.2 Traffic Generator

A traffic generator was developed to generate multiple sets of commodities for the different experiments. The generator accepts as input parameters the number of nodes in a network, the number of commodities, and the minimum and maximum demands. The generator selects OD pairs randomly and uniformly from the set of nodes and ensures that there are no duplicate OD pairs. The generator randomly selects the demands associated with OD pairs using a uniform distribution over the range specified by the minimum and maximum demands. The minimum and maximum demands are selected based on the test scenario. For example, when testing with N0 the maximum demand is always less than the 622 Mbps capacity of the lower bandwidth links. When testing with N1-N10, where all links have 2488 Mbps capacity, the maximum demand was less than half of that bandwidth. The minimum and maximum values are also adjusted depending on the number of commodities. Generally, when the number of commodities is doubled, the minimum and maximum demands are reduced by half, for aggregate demand compatibility.

When comparing the performance of model TE1 with the FCFS strategy, the minimum and maximum demand values are selected through an iterative process. The values are set so that when the FCFS strategy is applied to the set of commodities, at least 5% (and generally less than 30%) of the total demand is not satisfied. Setting to lower values would generate a set of commodities that would be fully satisfied by the FCFS strategy, and would not allow verification the benefit of model TE1. Setting to higher values would generate a set of commodities for which the majority of the demands would not be satisfied and make the test scenario unrealistic.

2.3.1.3 Offline Strategy Simulation Program

A program was developed to implement the online FCFS strategy. The program routes commodities from the set of commodities generated by the traffic generator one commodity at a time. The sequence in which commodities are routed is randomized. The program calculates the constrained shortest path for a given commodity by observing the link capacities. If a commodity is routed, the remaining bandwidth on the links along the shortest-path is adjusted based on the demand of that commodity, and the next commodity is handled. If a commodity cannot be routed due to lack of capacity, the commodity is rejected and the revenue associated with it is considered as missed revenue.

2.3.1.4 Computing Environment

All test cases are performed on a Compaq AlphaServer DS20E with dual 667 MHz processors and 4096 MB RAM. The machine is configured as a Network Queuing System and executes batch jobs. Each job on the system has access to approximately 20 MB RAM. The models are implemented using the GAMS [14] model description language, and integer-programming solutions are generated using CPLEX Linear Optimizer 7.0. Default settings for CPLEX are used with the exception that the MIP time limit is set to 3600 seconds and the relative optimality gap is set to 1%. The SAS software package is used to perform the statistical analysis.

2.3.2 Experimental Design for Study I

The first study compares the performance of an offline optimization strategy based on model TE1 with the performance of an online FCFS strategy. Since the main objective is to maximize revenue, the main performance metric used for comparison is the percentage of possible revenue that is missed. Since the model assumes that the revenue per unit of demand is constant, the percentage of revenue missed is defined as the ratio of total demand not delivered to the total demand. Hence, the lower the percentage, the better the performance of the corresponding strategy.

An ANOVA statistical analysis is performed on the percentage of revenue missed by each strategy and on the delta (difference) between the two percentages. The performance improvement by the optimization strategy is defined as the percentage of revenue missed by the FCFS strategy minus the percentage of revenue missed by the optimization strategy. The delta represents the incremental percentage of revenue gained by the optimization strategy. Therefore, the higher the delta, the more significant the improvement achieved by the optimization strategy.

2.3.2.1 Investigated Factors

The performance of each strategy and the performance improvement depend on multiple factors. The effect of the following factors and corresponding levels are explored:

1. The total number of OD pairs. All experimental networks consisted of 20 nodes, so the maximum number of distinct OD pairs is 380. To explore the effect of low, medium, and high levels of traffic meshing, the values of 80, 160, and 320 OD pairs are chosen, respectively. When studying the impact of the number of OD pairs on performance, the total demand is kept the same for the different levels by adjusting $\bar{d}$ accordingly. The average demand with 320 pairs is half of the average demand with 160 pairs, and the average demand with 160 pairs is half of that with 80 pairs.
2. The average demand per OD pair ($\bar{d}$). Multiple levels are explored, as described in the next section.
3. The range of demand per OD pair. Multiple levels are explored, as presented in the next section.
4. Network topology. The two sets of networks, N1-N5 and N6-N10, are used to explore the effect of network topology. The five networks in each set share common characteristics but differ only in topology.
5. Network density or average node degree. Two levels of 4 and 8 are explored. Each of the networks N1-N5 has a degree of 4, and each of the networks N6-N10 has a

degree of 8. To explore the impact of degree on performance using the two sets of networks, all networks are assigned the same bandwidth. Since the number of arcs in the second set is twice the number of arcs in the first set, the arc bandwidth in the second set is assigned half the bandwidth of arcs in the first set.

2.3.2.2 List of Hypotheses

Considering the above performance evaluation criteria and the list of factors that might impact performance, the following hypotheses are tested:

1. H_0 # 1: The average performance of the optimization model is not significantly different than that of the online FCFS strategy.
2. H_0 # 2: $\bar{d}$ does not affect the average performance of the optimization model nor the performance of the online FCFS strategy.
3. H_0 # 3: $\bar{d}$ does not affect the average performance improvement of the optimization model over the online FCFS strategy.
4. H_0 # 4: The number of OD pairs does not affect the average performance of the optimization model nor the performance of the online FCFS strategy.
5. H_0 # 5: The number of OD pairs does not affect the average performance improvement of the optimization model over the online FCFS strategy.
6. H_0 # 6: The range of traffic demand does not affect the average performance improvement of the optimization model over the online FCFS strategy.
7. H_0 # 7: Network topology does not affect the average performance of the optimization model nor the performance of the online FCFS strategy.
8. H_0 # 8: Network topology does not affect the average performance improvement of the optimization model over the online FCFS strategy.

9. H_0 # 9: The average node degree does not affect the average performance of the optimization model or the performance of the online FCFS strategy.
10. H_0 # 10: The average node degree does not affect the average performance improvement of the optimization model over the online FCFS strategy.

These hypotheses are tested on one or more of the network sets N0, N1-N5, and N6-N10. The tests are performed using the ANOVA statistical analysis on one-factor or multi-factor completely randomized-design experiments. The strategy (with two levels, optimization and FCFS) is one factor, and one or more of the factors listed above are used as the other factors. Each experiment is replicated 20 times with different sets of commodities. When comparing different networks, the same 20 sets of commodities are applied to all networks. For all tests the level of significance is set to $\alpha = 0.05$.

The tests are performed using the ANOVA statistical analysis on one, two, or three -factor completely randomized design experiments with fixed effects. When the performance of the two strategies is examined, the strategy (with two levels, optimization and FCFS) is one factor, and one or more of the factors listed above are used as the other factors. When the impact of the different factors on performance improvement is tested, the strategy is not used as a factor – the impact of one or more of the other factors on revenue Δ % is tested. Typically, a multi-factor experiment is conducted – if the interaction between the factors is found not to be significant, tests on the main effects are performed; otherwise, experiments with less factors are conducted. Each experiment is replicated 20 times with different sets of commodities. When comparing different networks, the same 20 sets of commodities are applied to all networks. For all tests the level of significance is set to $\alpha = 0.05$. the Duncan's multiple-range test [44] (hereafter referred to as the MR test) is performed when the number of OD pairs is used as a factor.

2.3.3 Experiment Test Results and Analysis for Study I

2.3.3.1 Experiment Test Results

The experimental results for the first study are presented in the Tables 2-2 through 2-7. The tables compare the performance of the offline optimization strategy based on model TE1 with the online FCFS strategy for different experiments. The tables show the following four performance metrics for each strategy: percentage of revenue missed, percentage of number of OD pairs missed, bandwidth utilization, and bandwidth efficiency. The *percentage of revenue missed* is defined above and is used as the performance metric for hypothesis testing. The other three metrics are not used for hypothesis testing but are analyzed to make general observations about the behavior of the two strategies. The *percentage of number of OD pairs missed* is defined as the ratio of number of OD pairs not delivered to the total number of OD pairs. *Bandwidth utilization* is defined as the ratio of total flow on all arcs to the total bandwidth of all arcs. *Bandwidth efficiency* is defined as the ratio of total demand delivered to the total bandwidth of all arcs. Finally, the *Revenue Δ* metric is the difference between the percentages of revenue missed by the FCFS and optimization strategies and represents the performance improvement by the optimization strategy. Each row in a table represents the averages for 20 replications of an experiment. Each table is described in turn and a preliminary analysis given. A statistical analysis and hypothesis test results are given in the next section.

Table 2-2 compares the performance of the two strategies when applied on network N0 as a function of the number of OD pairs and the average demand per OD pair. The table shows the results for 80, 160, and 320 OD pairs and varying demands per OD pair. $\bar{d}$ is adjusted based on the number of OD pairs so that the total demand remains the same. This allows comparing the results for different number of OD pairs and assessing its impact on performance. The results indicate that both strategies perform better for higher numbers of OD pairs and that optimization outperforms the FCFS strategy. The average performance improvement ranges from 7.78% to 13.77%.

Table 2-2 Average performance of optimization vs. FCFS for network N0 as a function of average demand and number of OD pairs

		Optimization				FCFS				
OD Pairs	Avg. Demand	Revenue Missed %	OD Pairs Missed %	Band. Util. %	Band. Eff. %	Revenue Missed %	OD Pairs Missed %	Band. Util. %	Band. Eff. %	Revenue Δ %
80	240	6.48	6.50	61.9	18.9	15.29	13.97	57.9	17.1	8.80
80	300	15.21	14.69	69.2	21.4	25.86	24.11	63.1	18.7	10.65
80	360	25.30	24.63	71.8	22.7	36.87	35.51	63.9	19.2	11.57
80	420	31.20	32.13	72.8	24.4	44.79	44.38	64.2	19.6	13.59
80	480	35.07	35.94	75.4	26.3	48.84	48.51	66.9	20.7	13.77
160	120	2.58	2.31	63.0	19.9	10.36	9.91	61.5	18.3	7.78
160	150	11.30	10.16	72.6	22.7	19.77	19.01	69.4	20.5	8.47
160	180	20.34	18.91	75.6	24.4	29.08	28.45	72.7	21.7	8.74
160	210	27.75	26.25	76.9	25.8	37.08	36.30	73.7	22.5	9.34
160	240	34.02	33.06	76.9	26.9	44.04	43.69	74.0	22.8	10.02
320	60	0.02	0.03	61.3	20.2	7.81	7.62	60.9	18.6	7.79
320	75	6.53	5.83	71.4	23.6	14.57	14.26	72.9	21.6	8.04
320	90	15.01	14.02	75.5	25.8	24.57	24.32	76.2	22.9	9.56
320	105	21.53	20.47	79.4	27.8	32.70	32.44	78.1	23.8	11.17
320	120	26.73	25.83	81.5	29.7	39.13	38.96	79.9	24.6	12.40

Table 2-3 compares the performance of the two strategies on network N0 as a function of the number of OD pairs and the demand range. The *demand range* is defined as the difference between the maximum and minimum demands per OD pair. $\overline{d}$ remains constant for each OD pairs value, specifically 240, 120, and 60 Mbps for 80, 160, and 320 OD pairs, respectively. The average demand and range are adjusted based on the number of OD pairs so that the total demand remains the same. As the previous table, this table also indicates that both strategies perform better for higher number of OD pairs and that the performance improvement by the optimization strategy does not depend on the demand range.

Table 2-3 Average performance of optimization vs. FCFS for network N0 as a function of demand range and number of OD pairs

		Optimization				FCFS				
OD Pairs	Demand Range	Revenue Missed %	OD Pairs Missed %	Band. Util. %	Band. Eff. %	Revenue Missed %	OD Pairs Missed %	Band. Util. %	Band. Eff. %	Revenue Δ %
80	320	7.30	6.69	60.5	18.7	15.97	13.49	57.3	16.9	8.67
80	160	6.82	7.31	60.4	18.9	15.71	15.24	57.6	17.1	8.89
80	80	8.46	9.12	58.2	18.6	16.76	16.73	56.7	16.9	8.30
80	40	8.89	9.31	57.3	18.5	16.61	16.61	56.7	16.9	7.72
80	20	9.10	9.31	57.1	18.5	16.59	16.59	56.7	16.9	7.49
160	160	3.11	2.63	63.4	19.9	10.93	9.98	61.2	18.3	7.81
160	80	2.38	2.22	63.0	19.9	10.28	10.07	61.6	18.3	7.91
160	40	1.95	1.94	62.3	20.0	9.91	9.88	61.4	18.3	7.96
160	20	1.87	1.91	62.2	20.0	9.62	9.62	61.5	18.4	7.75
160	10	1.88	1.91	62.1	19.9	9.56	9.57	61.5	18.4	7.68
320	80	0.06	0.05	61.3	20.1	7.76	7.39	60.7	18.6	7.70
320	40	0.00	0.00	61.4	20.2	7.87	7.78	61.1	18.6	7.87
320	20	0.00	0.00	61.5	20.3	7.98	7.97	61.4	18.7	7.98
320	10	0.00	0.00	61.6	20.3	7.99	8.00	61.6	18.7	7.99
320	5	0.00	0.00	62.3	20.5	8.03	8.04	62.2	18.8	8.03

Table 2-4 compares the performance of the two strategies as a function of network topology and number of OD pairs. The table shows the results for networks N1-N5, which share similar characteristics but differ in topology. All networks are tested with the same sets of commodities. The same $\bar{d}$ for each number of OD pairs is used for all networks. $\bar{d}$ is 600, 300, and 150 Mbps for 80, 160, and 320 OD pairs, respectively. The average demand is adjusted for the number of OD pairs so that the total demand remains the same. The results show that topology significantly affects the performance of both strategies and in a similar fashion; both strategies perform better for N1 and N2 and worse for N3-N5. However, the performance improvement seems to be less influenced by topology and no specific pattern is identified. The optimization strategy outperforms the FCFS strategy under all conditions. However the performance improvement is not as significant as in the case of N0. The performance improvement ranges from to 2.43% to 6.87%. This table also shows that both strategies perform better for higher number of OD pairs for all networks.

Table 2-4 Average performance of optimization vs. FCFS for networks N1-N5

		Optimization				FCFS				
Net. No.	OD Pairs	Revenue Missed %	OD Pairs Missed %	Band. Util. %	Band. Eff. %	Revenue Missed %	OD Pairs Missed %	Band. Util. %	Band. Eff. %	Revenue Δ %
N1	80	7.48	7.00	60.8	22.2	12.72	12.05	59.2	21.0	5.24
N2	80	10.00	9.19	62.5	21.6	16.87	16.01	61.0	20.0	6.87
N3	80	20.97	20.06	53.2	19.0	25.90	25.20	52.2	17.8	4.93
N4	80	20.72	19.81	58.2	19.1	25.09	24.38	56.9	18.0	4.37
N5	80	20.88	20.44	62.2	19.0	24.80	24.34	62.7	18.1	3.92
N1	160	4.95	4.53	61.8	23.0	8.56	8.26	61.0	22.2	3.61
N2	160	4.18	3.84	64.1	23.2	9.26	8.99	65.0	22.0	5.09
N3	160	19.35	18.69	54.6	19.6	22.66	22.31	55.4	18.7	3.30
N4	160	16.79	15.63	61.8	20.2	22.32	21.94	59.4	18.8	5.53
N5	160	19.33	18.34	66.0	19.6	23.08	22.67	67.0	18.7	3.75
N1	320	2.79	2.56	61.6	23.3	6.28	6.21	61.6	22.5	3.49
N2	320	2.96	2.66	63.0	23.3	6.60	6.53	67.3	22.4	3.64
N3	320	18.95	18.25	53.5	19.5	21.38	21.31	54.9	18.9	2.43
N4	320	13.81	12.86	61.6	20.7	19.61	19.38	60.6	19.3	5.79
N5	320	17.37	16.58	64.8	19.8	20.53	20.41	67.1	19.1	3.16

Table 2-5 is similar to Table 2-4 but shows the results when $\bar{d}$ is increased by 33%. $\bar{d}$ is 800, 400, and 200 Mbps for 80, 160, and 320 OD pairs, respectively. As expected, the increase in demand increases the percentage of demand not delivered, however, the total demand delivered is increased as shown by the bandwidth utilization and bandwidth efficiency. The increase in traffic affects both strategies in similar fashion, however the average performance improvement by the optimization strategy increases in comparison with the previous table. The performance improvement ranges from 4.6% to 9.06%.

Table 2-5 Average performance of optimization vs. FCFS for networks N1-N5 with 33% increase in demand

		Optimization				FCFS				
Net. No.	OD Pairs	Revenue Missed %	OD Pairs Missed %	Band. Util. %	Band. Eff. %	Revenue Missed %	OD Pairs Missed %	Band. Util. %	Band. Eff. %	Revenue Δ %
N1	80	22.73	19.75	65.9	24.7	28.93	26.67	64.6	22.7	6.20
N2	80	27.89	23.25	63.9	23.1	35.10	32.09	63.2	20.8	7.21
N3	80	33.18	30.00	56.9	21.4	38.48	36.43	55.8	19.7	5.31
N4	80	34.26	30.87	60.5	21.0	38.93	36.89	59.6	19.5	4.67
N5	80	33.03	30.25	65.3	21.4	38.19	36.44	66.0	19.8	5.17
N1	160	15.74	13.72	71.5	27.3	21.43	20.32	69.5	25.5	5.69
N2	160	16.93	14.66	71.2	26.9	25.43	23.97	70.4	24.2	8.50
N3	160	28.34	25.72	61.4	23.2	33.98	32.87	61.7	21.4	5.65
N4	160	30.36	26.88	64.8	22.6	35.79	34.69	63.2	20.8	5.43
N5	160	28.27	25.41	72.5	23.2	34.89	33.63	73.4	21.1	6.62
N1	320	11.46	10.00	72.0	28.3	17.39	16.86	72.0	26.4	5.94
N2	320	13.14	11.39	72.0	27.8	22.19	21.64	73.2	24.9	9.06
N3	320	25.40	23.61	61.8	23.8	30.55	30.20	62.9	22.2	5.15
N4	320	25.31	22.56	67.1	23.9	33.05	32.35	64.6	21.4	7.74
N5	320	25.04	23.05	73.6	24.0	30.24	29.76	76.8	22.3	5.21

Table 2-6 is similar to Table 2-5 and shows the results for networks N6-N10 with similar experiments, including the same sets of commodities. These statistics enable performance comparisons of the two strategies as a function of network topology and number of OD pairs. The results confirm previous observations: the network topology affects the performance of both strategies, and does so in similar fashion. The performance improvement by the optimization strategy is practically insignificant and the difference in improvement between different networks (for a given number of OD pairs) is generally less than 2%. The results also indicate that both strategies perform better as the number of OD pairs increases.

Table 2-6 also enables examination of the effect of node degree on the performance of both strategies. Comparing the values with the corresponding results in Table 2-5 indicates that both strategies perform significantly better with higher node degree for all networks and any number of OD pairs. For example, the overall average percentage of revenue missed by the optimization and FCFS strategies over N1-N5 and the three numbers of OD pairs are 24.74% and 30.97%, respectively. In contrast, the equivalent averages over N6-N10 are 4.63% and 7.8%, respectively.

Table 2-6 Average performance of optimization vs. FCFS for networks N6-N10

		Optimization				FCFS				
Net. No.	OD Pairs	Revenue Missed %	OD Pairs Missed %	Band. Util. %	Band. Eff. %	Revenue Missed %	OD Pairs Missed %	Band. Util. %	Band. Eff. %	Revenue Δ %
N6	80	7.89	8.00	62.7	29.4	12.91	11.29	60.4	27.8	5.02
N7	80	9.66	9.94	63.8	28.9	16.15	14.37	62.1	26.8	6.50
N8	80	12.16	12.06	60.8	28.1	17.20	15.25	58.7	26.5	5.04
N9	80	12.40	11.00	62.4	28.0	18.79	16.83	61.8	26.0	6.39
N10	80	13.04	12.25	61.4	27.8	18.09	16.38	60.3	26.2	5.05
N6	160	0.18	0.19	60.8	32.3	1.25	1.04	63.7	32.0	1.07
N7	160	0.37	0.31	63.2	32.3	3.20	2.81	66.0	31.4	2.83
N8	160	2.49	2.19	63.8	31.6	5.14	4.38	63.4	30.7	2.65
N9	160	4.47	3.78	63.8	31.0	7.91	7.15	66.2	29.8	3.45
N10	160	4.75	4.00	65.1	30.9	8.41	7.49	64.7	29.7	3.67
N6	320	0.00	0.00	57.0	32.0	0.00	0.00	60.1	32.0	0.00
N7	320	0.00	0.00	58.5	32.0	0.19	0.17	64.3	31.9	0.19
N8	320	0.18	0.17	59.9	31.9	0.50	0.43	62.8	31.8	0.32
N9	320	1.28	1.16	63.5	31.5	4.59	4.27	66.2	30.5	3.31
N10	320	0.63	0.58	63.8	31.8	2.63	2.39	65.9	31.1	2.01

Table 2-7 gives the experimental results for networks N6-N10 when the arc capacity is doubled from the Table 2-6 baseline values (from 1244 Mbps to 2488 Mbps) and the baseline demand is tripled ($\bar{d}$ is increased from 200 Mbps to 600 Mbps). When the demand is only doubled (to correspond to doubling the capacity), all demand is delivered (the results are not shown here). When the demand is tripled, more than twice the original demand is delivered. For example, for network N9 with 320 OD pairs and based on Table 2-6, the percentage of demand delivered is 98.72%. Based on Table 2-7, the percentage is 88.43%. The ratio of increase in demand delivered is then calculated as (88.43%*3)/98.72% which is approximately 2.69. Since revenues are proportional to traffic delivered, the conclusion in this case is that doubling the capacity of the network yields 2.69 relative increase in revenue. This result is also demonstrated by the significant increase in bandwidth efficiency. For network N0, it increases from 31.5% in Table 2-6 to 41.2% in Table 2-7.

Table 2-7 Average performance of optimization vs. FCFS for networks N6-N10 (arc capacity is doubled)

		Optimization				FCFS				
Net. No.	OD Pairs	Revenue Missed %	OD Pairs Missed %	Band. Util. %	Band. Eff. %	Revenue Missed %	OD Pairs Missed %	Band. Util. %	Band. Eff. %	Revenue Δ %
N6	320	5.73	4.89	83.0	45.3	11.77	11.13	86.7	42.4	6.04
N7	320	5.80	4.98	83.2	45.2	14.71	14.06	87.0	41.0	8.91
N8	320	14.17	12.67	79.1	41.2	17.75	17.18	79.2	39.5	3.58
N9	320	11.57	10.52	81.9	42.5	18.93	18.28	86.1	38.9	7.36
N10	320	14.81	13.39	79.4	40.9	19.66	19.18	81.8	38.6	4.84

2.3.3.2 Hypothesis Test Results

The data in on which Tables 2-2 to 2-7 are based is used to test the hypotheses listed in section 2.3.2.2. The results based on $\alpha = 0.05$ are as follows.

1. H_0 # 1: The average performance of the optimization model is not significantly different than that of the online FCFS strategy.
   a. H_0 is rejected on network N0, based on Table 2-2. Three separate two-factor experiments (one for each set of OD pairs) with strategy (two levels) and $\bar{d}$ (five levels) as the two factors were conducted. There was no significant interaction between the two factors. All tests show that the % of revenue missed by optimization is significantly less than the % of revenue missed by FCFS with a P-value less than 0.0001. Optimization outperforms FCFS with performance improvement ranging from 7.78% to 13.77%.
   b. H_0 is rejected on networks N1-N5, based on Tables 2-4 and 2-5. A three-factor experiment with strategy (two levels), network (5 levels), and OD pairs (three levels) as the factors were conducted. A significant interaction (P-value is 0.0001) was found between network and OD pairs. Subsequently, three separate two-factor experiments (one for each set of OD pairs) with strategy and network as the two factors were conducted. There was no significant interaction between the two factors (P-values are greater than 0.6). All tests show that the % of revenue missed by optimization is significantly less than the % of revenue missed by FCFS with a P-value less than 0.0001. Optimization outperforms FCFS with performance improvement ranging from 2.43% to 9.06%.

c. H_0 is rejected on networks N6-N10, based on Table 2-6. The test procedure and results are similar to those for networks N1-N5. Optimization outperforms FCFS with performance improvement ranging from 0.19% to 6.5%.

2. H_0 # 2: $\bar{d}$ does not affect the average performance of the optimization model or the performance of the online FCFS strategy.

   a. H_0 is rejected on network N0, based on Table 2-2. A two-factor experiment with strategy and $\bar{d}$ as the two factors was conducted for 160 OD pairs. There was no significant interaction between the two factors (P-value is 0.9075). $\bar{d}$ significantly affects the performance of both optimization and FCFS strategies. The percentage of revenue missed by either strategy increases as demand increases. For example, for 160 OD pairs, the average revenue missed by the optimization model increases from 2.58% to 34.02% from the lowest to the highest demand. Similarly, the missed revenue by the FCFS strategy increases from 10.36% to 44.04%.

   b. H_0 is rejected on networks N1-N5, based on Tables 2-4 and 2-5. Three separate two-factor experiments (one for each set of OD pairs) with strategy and $\bar{d}$ (two levels) as the two factors were conducted. There was no significant interaction between the two factors. All tests show that $\bar{d}$ significantly affects the performance of both optimization and FCFS strategies. The percentage of revenue missed by either strategy increases as demand increases. For example, for 80 OD pairs, the average revenue missed by the optimization model for the five networks increases from

16.01% to 30.22%. Similarly, the average by the FCFS strategy increases from 21.8% to 35.93%.

3. H_0 # 3: $\bar{d}$ does not affect the average performance improvement of the optimization model over the online FCFS strategy.
    a. H_0 is rejected on network N0, based on Table 2-2. Three separate one-factor experiments (one for each set of OD pairs) on the revenue Δ % with $\bar{d}$ as the factor were conducted. All tests show that $\bar{d}$ significantly affects the performance improvement (P-value is 0.0001). The improvement by the optimization model over the FCFS strategy increases as the demand increases, but the increase is somewhat limited. For example, for 80 OD pairs, the difference in revenue missed increases from 8.8% to 13.77% from the lowest to the highest demand. So, while the demand is doubled, the improvement increases by only approximately 5%.
    b. H_0 is rejected on networks N1-N5, based on Tables 2-4 and 2-5. The improvement by the optimization model over the FCFS strategy increases as the demand increases, but the increase is somewhat limited. For 80, 160 and 320 OD pairs, and for 33% traffic increase, the average improvement for the five networks increases only by 0.64%, 2.12%, and 2.92%, respectively.
4. H_0 # 4: The number of OD pairs does not affect the average performance of the optimization model or the performance of the online FCFS strategy.
    a. H_0 is rejected on network N0, based on Table 2-2. A two-factor experiment with strategy and number of OD pairs as the factors was conducted. There was no significant interaction between the two factors. The MR test is

conducted on the number of OD pair levels. The number of OD pairs significantly affects the performance of both optimization and FCFS strategies. The percentage of revenue missed by either strategy decreases as the number of OD pairs increases. For example, the average revenue missed (for the five levels of demand) by the optimization model is 22.65% and 13.96% for 80 and 320 OD pairs, respectively. Similarly, the revenue missed by the FCFS strategy is 34.33% and 23.76 % for 80 and 160 OD pairs, respectively.

b. H_0 is rejected on networks N1-N5, based on Tables 2-4 and 2-5. A three-factor experiment with strategy, network, and OD pairs as the factors were conducted. A significant interaction (P-value is 0.0001) was found between network and OD pairs. Subsequently, five separate two-factor experiments (one for each network) with strategy and OD pairs as the two factors were conducted. There was no significant interaction between the two factors. All tests show that the number of OD pairs significantly affects the performance of both optimization and FCFS strategies. The results of the MR tests on the number of OD pairs differ from one network to another. For example, the test for network 1 shows that the performance with 80 OD pairs is significantly lower that that with 160 or 320 OD pairs, but there is no significant difference between the performance with 160 and 320 OD pairs. In contrast, the test for network 2 shows that performance is significantly lower with 80 OD pairs than with 160 OD pairs, and in turn performance with 160 is lower than that with 320 OD pairs. Typically, the percentage of revenue missed by either strategy decreases as the number of OD pairs increases. For example (based on Table 2-5), the average percentages of revenue missed by optimization for the five networks are 30.22% and 20.07% for 80 and 320 OD pairs, respectively.

c. H_0 is rejected on networks N6-N10, based on Table 2-6. The test procedure and results are similar to those for networks N1-N5. The number of OD pairs significantly affects the performance of both optimization and FCFS strategies. Typically, the percentage of revenue missed by either strategy decreases as the number of OD pairs increases. The average percentages of revenue missed by optimization for the five networks are 11.03% and 0.42% for 80 and 320 OD pairs, respectively

5. H_0 # 5: The number of OD pairs does not affect the average performance improvement of the optimization model over the online FCFS strategy.

   a. H_0 is rejected on network N0, based on Table 2-2. The number of OD pairs has some impact on the average performance improvement. However, the impact is limited and with no particular pattern. The improvement for 80 OD pairs is higher than 320, which in turn is higher than 160. The improvement for 80 OD pairs is only 2.81% higher than the improvement for 160 pairs.

   b. H_0 is rejected on networks N1-N5, based on Tables 2-4 and 2-5. The number of OD pairs has some impact on the average performance improvement. However, the impact is limited and with no particular pattern. Based on Table 2-4, the improvement for 80 OD pairs is higher than 160, which in turn is higher than 320. The difference in improvement is less than 2%. Based on Table 2-5, the improvement for 320 and 160 OD pairs is higher than that of 80 OD pairs.

   c. H_0 is rejected on networks N6-N10, based on Table 2-6. The number of OD pairs has some impact on the average performance improvement.

6. H_0 # 6: The range of traffic demand does not affect the average performance improvement of the optimization model over the online FCFS strategy.

   a. H_0 is not rejected on network N0, based on Table 2-3. Three separate one-factor experiments (one for each set of OD pairs) on the revenue Δ % with the demand range as the factor were conducted. All tests show that the demand range per OD pair has no significant impact on the average performance improvement of the optimization model over the online FCFS strategy. The P-value for 80, 160, and 320 OD-pair tests is 0.4057, 0.9577, and 0.7006, respectively.

7. H_0 # 7: Network topology does not affect the average performance of the optimization model or the performance of the online FCFS strategy.

   a. H_0 is rejected on networks N1-N5, based on Tables 2-4 and 2-5. Two separate (one for 80 and another for 160 OD pairs) two-factor experiments with strategy and network as the factors were conducted. In both tests there was no significant interaction between the two factors (P-value is 0.6019 and 0.2682, respectively). Both tests show that topology significantly affects the performance of both strategies. For example, the percentage of revenue missed by optimization on N1 and N5 for 80 OD pairs is 7.48% and

20.88%, respectively. Similarly, with FCFS it is 12.72% and 24.8%, respectively.

b. H_0 is rejected on networks N6-N10, based on Table 2-6. The test procedure and results are similar to those for networks N1-N5. The topology significantly affects the performance of both strategies. For example, the percentage of revenue missed by optimization on N6 and N10 for 80 OD pairs is 7.89% and 13.04%, respectively. Similarly, with FCFS it is 12.91% and 18.09%, respectively.

8. H_0 # 8: Network topology does not affect the average performance improvement of the optimization model over the online FCFS strategy.

a. H_0 is rejected on networks N1-N5, based on Tables 2-4 and 2-3. A two-factor experiment on the revenue Δ % with network and OD pairs as the factors were conducted. A significant interaction (P-value is 0.0001) was found between network and OD pairs. Subsequently, three separate one-factor experiments (one for each set of OD pairs) with network as the factor were conducted. All tests show that topology affects the average performance improvement, but the impact is somewhat limited. For a given number of OD pairs, the difference between the network with the highest improvement and the network with the lowest improvement is less than 4%.

b. H_0 is rejected on networks N6-N10, based on Table 2-6. The test procedure and results are similar to those for networks N1-N5. The topology affects the average performance improvement, but the impact is somewhat limited. For a given number of OD pairs, the difference between the network with the highest improvement and the network with the lowest improvement is less than 3%.

9. H_0 # 9: The average node degree does not affect the average performance of the optimization model or the performance of the online FCFS strategy.
    a. H_0 is rejected when comparing networks N1-N5 with networks N6-N10, based on Tables 2-5 and 2-6. The degree significantly affects the average performance of both the optimization and FCFS strategies. The average percentage of revenue missed on N1-N5 is 24.47% compared with 4.63% on N6-N10. Similarly with FCFS, it is 30.97% on N1-N5 compared with 7.8% on N6-N10.
10. H_0 # 10: The average node degree does not affect the average performance improvement of the optimization model over the online FCFS strategy.
    a. H_0 is rejected when comparing networks N1-N5 with networks N6-N10, based on Tables 2-5 and 2-6. The degree affects the average performance improvement, but the impact is somewhat limited. The average improvement on N1-N5 is 6.24% compared with 3.17% on N6-N10.

In summary, these computational experiments involving the solution of MPLS traffic engineering problems provide empirical evidence that the performance and profitability of network designs is influenced by many factors. These results are further evaluated in section 2.4.

2.3.4 Experimental Design for Study II

The second study evaluates the benefits of load balancing in MPLS networks. It compares the performance of single-LSP per OD pair scheme based on model TE1 against two-LSP per OD pair (load-balancing) scheme based on model TE2. In the case of the load-balancing scheme the demand for an OD pair is split between the two LSPs based on predetermined split ratio. Model TE2 does not require both LSPs to be configured; so only portion of the demand of an OD pair may be delivered. The experiments evaluated split ratios of 0.5 and 0.9. The study attempts to determine whether load balancing improves performance and whether one ratio is more beneficial than the other.

As in the first study, the main performance metric used for comparison is the percentage of revenue missed. The lower the percentage, the better the performance of the corresponding scheme. The performance improvement by the load-balancing scheme is defined as the percentage of revenue missed by the single-LSP scheme minus the percentage of revenue missed by the load-balancing scheme at a given split ratio. The delta represents the incremental percentage of revenue gained by the load-balancing scheme. The higher the delta, the more significant the improvement achieved by that scheme.

2.3.4.1 Investigated Factors

The ANOVA statistical analysis is performed on the percentage of revenue missed by each scheme and on the delta between the two percentages. The performance of each scheme and the performance improvement depend on multiple factors.

The effect of the following factors and corresponding levels are explored:

1. The total number of OD pairs. All experimental networks consisted of 20 nodes, so the maximum number of distinct OD pairs is 380. To explore the

effect of low, medium, and high levels of traffic meshing, the values of 80, 160, and 320 OD pairs are chosen, respectively.

2. The demand per OD pair. Multiple levels are explored and presented in the next section.
3. The network topology. The two sets of networks N1-N5 and N6-N10 are used to explore the effect of network topology. The five networks in each set share common characteristics but differ only in topology.

2.3.4.2 List of Hypotheses

Considering the above performance evaluation criteria and the list of factors that might affect performance, the following hypotheses are tested:

1. H_0 # 11: The average performance of the one-LSP scheme is not significantly different than that of the two-LSP scheme.
2. H_0 # 12: The demand per OD pair does not affect the average performance improvement of the two-LSP scheme over the one-LSP scheme.
3. H_0 # 13: The split ratio does not affect the average performance of the two-LSP scheme.
4. H_0 # 14: The number of OD pairs does not affect the average performance improvement of the two-LSP scheme over the one-LSP scheme.
5. H_0 # 15: Network topology does not affect the average performance improvement of the two-LSP scheme over the one-LSP scheme.

These hypotheses are tested on one or more of the network sets N0, N1-N5, and N6-N10. The tests are performed using the ANOVA statistical analysis on one-factor or two-factor completely randomized design experiments. With a two-factor experiment the split ratio, which identifies the scheme (one-LSP or two-LSP scheme), is one factor with two levels (1 and 0.5) or three levels (1, 0.5, and 0.9), and the second factor is one of the factors listed above (OD demand, OD pairs, or network). First, a two-factor experiment is conducted – if the interaction between the two factors is found not to be significant, tests on the main effects are performed; otherwise, a one-factor experiment is conducted with the split ratio as the treatment at a one level of the second factor. Each experiment is replicated 20 times with different sets of commodities. When comparing different networks, the same 20 sets of commodities are applied to all networks. For all tests the level of significance is set to $\alpha = 0.05$. When more than two levels are involved in a test, the Duncan's multiple-range test [44] (hereafter referred to as the MR test) is performed.

2.3.5 Experiment Test Results and Analysis for Study II

2.3.5.1 Experiment Test Results

The experimental results are presented in the following tables. The tables compare the performance of the one-LSP scheme with the two-LSP scheme with split ratios of 0.5 and 0.9. On these tables, R1 is the average percentage of revenue missed using the one-LSP scheme. R2 and R3 (if present) are the average percentages of revenue missed using the two-LSP scheme with split ratios of 0.5 and 0.9, respectively. R1-R2 and R1-R3 represent the performance improvement by the two-LSP scheme with the corresponding split ratio. Each row in a table represents the averages for 20 replications of an experiment.

The "Statistical Comparison" column summarizes the result of running ANOVA on one or two-factor experiments. First, a two-factor experiment (for one level of OD pairs) with the split ratio as one factor and OD demand as the other factor is conducted. If interaction is found not to be significant, the significance test result on the main effect of the split ratio is summarized in this column. Otherwise, a separate one-factor test for each OD demand level is conducted and the significance test result is summarized for each row in a table. On this column, the notation "R1=R2" means that the corresponding null hypothesis is not rejected and that there is no significant difference between the means R1 and R2. Similarly, "R1>R2" means that the corresponding null hypothesis is rejected and that there is enough evidence to suggest that R1 is greater than R2.

When R1, R2, and R3 are involved, the MR test is performed and the following notation is used to present the results. "R1=R2=R3" means that the corresponding null hypothesis is not rejected and that there is no significant difference between the means R1, R2, and R3. "R1>R2>R3" means that the corresponding null hypothesis is rejected, and that R1 is significantly greater than R2 and R2 is significantly greater than R3. "R1>R2=R3" means that the corresponding null hypothesis is rejected, and that R1 is significantly greater than R2, but R2 and R3 are not significantly different. "R1≥R2=R3" means that the corresponding null hypothesis is rejected, and that R1 is significantly greater than R3, but R1 and R2 are not significantly different and R2 and R3 are not significantly different.

Table 2-8 compares the performance of the two schemes when applied on network N0 as a function of the number of OD pairs and the demand per OD pair in Mbps. The table shows the results for 80, 160, and 320 OD pairs and varying demands per OD pair. The demand per OD pair is adjusted based on the number of OD pairs so that the total demand remains the same. This allows comparing the results for different number of OD pairs and assessing its impact on performance. The table shows that both schemes perform better for higher numbers of OD pairs. Note that the lowest demands with 160 and 320 OD pairs shown in the table are 180 and 90 Mbps, respectively. With lower demands (that correspond to the demands shown for 80 OD pairs), all traffic is delivered and no revenue is missed. The table shows more significant performance improvement by the load-balancing scheme with 80 OD pairs than with 160 and 320 OD pairs. The highest improvement is 11.75% using a split ratio of 0.9 for 80 OD pairs and a demand of 320 Mbps per OD pair. However, in many cases the improvement by the load-balancing scheme is null or not significant.

Table 2-8 Average performance of one-LSP vs. two-LSP schemes for network N0 as a function of demand and number of OD pairs

OD Pairs	OD Demand	R1 %	R2 %	R3 %	R1-R2 %	R1-R3 %	Statistical Comparison
80	240	10.72	5.47	9.38	5.25	1.34	R1>R3>R2
80	260	10.69	9.90	9.88	0.78	0.81	R1=R2=R3
80	280	13.69	11.97	12.21	1.72	1.47	R1=R2=R3
80	300	13.69	13.69	13.29	0.00	0.40	R1=R2=R3
80	320	28.38	19.47	16.63	8.91	11.75	R1>R2=R3
80	340	28.38	21.31	17.33	7.06	11.05	R1>R2>R3
80	360	31.25	23.41	25.99	7.84	5.26	R1>R3>R2
80	380	31.25	23.41	26.97	7.84	4.28	R1>R3>R2
80	400	31.25	25.69	28.05	5.56	3.20	R1>R3>R2
80	420	35.75	33.41	28.83	2.34	6.92	R1>R2>R3
80	440	35.75	33.41	31.64	2.34	4.11	R1≥R2=R3
80	460	35.75	35.72	32.36	0.03	3.39	R1=R2>R3
160	180	0.56	0.24		0.33		R1=R2
160	200	2.06	2.06		0.00		R1=R2
160	220	4.28	4.28		0.00		R1=R2
160	260	11.59	9.63		1.97		R1>R2
160	300	15.72	15.72		0.00		R1=R2
160	340	20.59	20.59		0.00		R1=R2
160	380	25.98	23.22		2.77		R1>R2
320	90	0.00	0.00		0.00		R1=R2
320	100	0.36	0.13		0.23		R1>R2
320	110	3.05	2.18		0.87		R1>R2
320	130	8.69	8.69		0.00		R1=R2
320	150	14.38	13.42		0.95		R1>R2
320	170	18.34	17.30		1.05		R1>R2
320	190	21.11	20.68		0.43		R1=R2

Tables 2-9 to 2-13 compare the performance of the two schemes as a function of the demand per OD pair for 80 OD pairs when applied on networks N1 to N5. Tables 2-9 to 2-11 show the results also for 160 OD pairs. The impact of network topology on performance and performance improvement is analyzed by comparing the results among the different tables. The results show that topology significantly affects the performance of both schemes but has less significant impact on performance improvement. For the same demand size, the performance on different networks can be significantly different. For example, on network N1 with 80 OD pairs and 700 Mbps of demand per OD pair, R1, R2, and R3 are 15.13%, 9.97%, and 13.61%, respectively. In contrast, on network N4 the corresponding percentages are 28.31%, 22.59%, and 27.67. On the other hand, the performance improvement pattern and magnitude are very similar on the different networks. For example, on all networks and with 80 OD pairs and split-ratio of 0.5, the load-balancing scheme does not provide any improvement with demands of 460, 580, and 620 Mbps; the highest improvement on all networks is achieved with a demand of 660 Mbps. Similarly, on all networks with 80 OD pairs and split-ratio of 0.9, load-balancing scheme does not provide any improvement with a demand of 620 Mbps; the highest improvement on all networks is achieved with a demand of 540 Mbps.

Table 2-9 Average performance of one-LSP vs. two-LSP schemes for network N1 as a function of demand and number of OD pairs

OD Pairs	OD Demand	R1 %	R2 %	R3 %	R1-R2 %	R1-R3 %	Statistical Comparison
80	460	1.44	1.44	0.92	0.00	0.52	R1=R2=R3
80	500	5.75	3.22	3.91	2.53	1.84	R1=R3=R2
80	540	5.75	3.19	2.91	2.56	2.84	R1>R2=R3
80	580	5.75	5.75	5.28	0.00	0.47	R1=R2=R3
80	620	5.75	5.75	5.75	0.00	0.00	R1=R2=R3
80	660	15.13	9.97	14.18	5.16	0.94	R1=R3>R2
80	700	15.13	9.97	13.61	5.16	1.51	R1=R3>R2
160	230	0.35	0.22		0.13		R1>R2
160	250	1.66	0.83		0.83		
160	270	1.59	1.59		0.00		P-value
160	290	4.72	2.98		1.73		0.0293
160	310	4.75	4.72		0.03		
160	330	8.84	8.09		0.75		
160	350	8.84	8.84		0.00		

Table 2-10 Average performance of one-LSP vs. two-LSP schemes for network N2 as a function of demand and number of OD pairs

OD Pairs	OD Demand	R1 %	R2 %	R3 %	R1-R2 %	R1-R3 %	Statistical Comparison
80	460	1.38	1.38	0.66	0.00	0.71	R1=R2=R3
80	500	6.75	3.38	5.27	3.38	1.48	R1≥R3=R2
80	540	6.75	3.38	3.17	3.38	3.58	R1>R2=R3
80	580	6.75	6.75	6.25	0.00	0.50	R1=R2=R3
80	620	6.75	6.75	6.75	0.00	0.00	R1=R2=R3
80	660	17.88	11.59	17.72	6.28	0.16	R1=R3>R2
80	700	17.88	11.59	16.09	6.28	1.79	R1=R3>R2
160	230	0.22	0.00		0.22		R1=R2
160	250	1.50	0.67		0.83		
160	270	1.47	1.47		0.00		P-value
160	290	4.19	2.63		1.56		0.0651
160	310	4.19	4.19		0.00		
160	330	8.50	7.41		1.09		
160	350	8.50	8.50		0.00		
160	400	13.69	13.69		0.00		

Table 2-11 Average performance of one-LSP vs. two-LSP schemes for network N3 as a function of demand and number of OD pairs

OD Pairs	OD Demand	R1 %	R2 %	R3 %	R1-R2 %	R1-R3 %	Statistical Comparison
80	460	13.94	13.94	12.11	0.00	1.83	R1=R2=R3
80	500	19.50	16.60	16.50	2.91	3.01	R1>R2=R3
80	540	19.52	16.59	16.12	2.93	3.40	R1>R2=R3
80	580	19.50	19.50	18.91	0.00	0.59	R1=R2=R3
80	620	19.50	19.50	19.50	0.00	0.00	R1=R2=R3
80	660	27.63	23.25	24.36	4.38	3.27	R1>R3=R2
80	700	27.63	23.25	25.19	4.38	2.44	R1=R3≥R2
160	230	14.16	12.91		1.25		R1=R2
160	250	16.78	15.44		1.35		
160	270	16.78	16.78		0.00		P-value
160	290	19.47	18.13		1.34		0.1079
160	310	19.47	19.47		0.00		
160	330	22.34	20.88		1.47		
160	350	22.34	22.34		0.00		
160	400	26.34	26.34		0.00		

Table 2-12 Average performance of one-LSP vs. two-LSP schemes for network N4 as a function of demand

OD Pairs	OD Demand	R1 %	R2 %	R3 %	R1-R2 %	R1-R3 %	Statistical Comparison
80	460	9.75	9.75	7.39	0.00	2.36	R1=R2=R3
80	500	17.63	13.53	16.56	4.09	1.07	R1=R3≥R2
80	540	17.63	13.53	13.28	4.09	4.35	R1>R2=R3
80	580	17.63	17.63	17.15	0.00	0.47	R1=R2=R3
80	620	17.63	17.63	17.63	0.00	0.00	R1=R2=R3
80	660	28.31	22.59	27.67	5.72	0.64	R1=R3>R2
80	700	28.31	22.59	25.51	5.72	2.81	R1=R3≥R2

Table 2-13   Average performance of one-LSP vs. two-LSP schemes for network N5 as a function of demand

OD Pairs	OD Demand	R1 %	R2 %	R3 %	R1-R2 %	R1-R3 %	Statistical Comparison
80	460	13.44	13.44	11.74	0.00	1.69	R1=R2=R3
80	500	19.31	16.31	16.76	3.00	2.55	R1≥R3=R2
80	540	19.31	16.31	17.06	3.00	2.25	R1≥R3=R2
80	580	19.31	19.31	18.88	0.00	0.43	R1=R2=R3
80	620	19.31	19.31	19.31	0.00	0.00	R1=R2=R3
80	660	27.56	23.09	27.04	4.47	0.52	R1=R3>R2
80	700	27.56	23.13	25.36	4.44	2.20	R1≥R3=R2

Tables 2-14 to 2-18 compare the performance of the two schemes as a function of the demand per OD pair for 160 OD pairs when applied to networks N6 to N10. Again, these tables are mainly used to assess the impact of network topology on performance and performance improvement. The results reinforce the conclusion reached from analyzing the results for networks N1-N5; that topology significantly affects the performance of both schemes but has less significant impact on performance improvement. For the same demand size, the performance on different networks can be significantly different. For example, on network N7 with a demand of 1000 Mbps, R1 and R2 are both 3.86%. In contrast, on network N10 the corresponding percentages are 13.00% and 12.98%. On the other hand, the performance improvement pattern and magnitude are very similar on the different networks. For example, on all networks the load-balancing scheme provides no or insignificant improvement with demands of 800, 1000, 1100, and 1200 Mbps. Furthermore, the highest improvement on all networks is achieved with a demand of 1300 and 1400 Mbps. Although the improvement differs from one network to another, the major reason for the improvement is the demand size and not the topology differences. Only a single OD pair with a demand of 1300 Mbps can be carried on 2488 Mbps links. Splitting such a demand in half, allows packing the full demand of one OD pair and half the demand of another OD pair (recall that the model allows delivering only a portion of an OD pair demand). This explains the significant performance improvement for such large demands. The highest improvement achieved is 17.83% on network N9.

Table 2-14 Average performance of one-LSP vs. two-LSP schemes for network N6 as a function of demand

OD Pairs	OD Demand	R1 %	R2 %	R1-R2 %	Statistical Comparison
160	800	0.00	0.00	0.00	R1=R2
160	900	0.63	0.00	0.63	R1>R2
160	1000	4.74	4.74	0.00	R1=R2
160	1100	4.31	4.31	0.00	R1=R2
160	1200	3.65	3.54	0.10	R1=R2
160	1300	35.69	20.84	14.84	R1>R2
160	1400	34.94	19.53	15.41	R1>R2

Table 2-15 Average performance of one-LSP vs. two-LSP schemes for network N7 as a function of demand

OD Pairs	OD Demand	R1 %	R2 %	R1-R2 %	Statistical Comparison
160	800	0.00	0.00	0.00	R1=R2
160	900	4.38	1.21	3.16	R1>R2
160	1000	3.86	3.86	0.00	R1=R2
160	1100	3.88	3.88	0.00	R1=R2
160	1200	4.38	3.91	0.47	R1=R2
160	1300	35.38	25.98	9.39	R1>R2
160	1400	35.06	25.70	9.36	R1>R2

Table 2-16 Average performance of one-LSP vs. two-LSP schemes for network N8 as a function of demand

OD Pairs	OD Demand	R1 %	R2 %	R1-R2 %	Statistical Comparison
160	800	0.34	0.34	0.00	R1=R2
160	900	11.88	3.88	8.00	R1>R2
160	1000	11.88	11.86	0.02	R1=R2
160	1100	11.84	11.84	0.00	R1=R2
160	1200	11.25	11.25	0.00	R1=R2
160	1300	38.13	23.63	14.50	R1>R2
160	1400	38.13	23.60	14.52	R1>R2

Table 2-17 Average performance of one-LSP vs. two-LSP schemes for network N9 as a function of demand

OD Pairs	OD Demand	R1 %	R2 %	R1-R2 %	Statistical Comparison
160	800	2.34	2.34	0.00	R1=R2
160	900	10.74	5.98	4.77	R1>R2
160	1000	10.56	10.56	0.00	R1=R2
160	1100	10.56	10.56	0.00	R1=R2
160	1200	9.69	9.06	0.63	R1=R2
160	1300	37.09	19.27	17.83	R1>R2
160	1400	37.06	21.09	15.97	R1>R2

Table 2-18 Average performance of one-LSP vs. two-LSP schemes for network N10 as a function of demand

OD Pairs	OD Demand	R1 %	R2 %	R1-R2 %	Statistical Comparison
160	800	1.44	1.44	0.00	R1=R2
160	900	13.00	5.39	7.61	R1>R2
160	1000	13.00	12.98	0.02	R1=R2
160	1100	13.00	12.98	0.02	R1=R2
160	1200	17.97	13.44	4.53	R1=R2
160	1300	38.03	24.13	13.90	R1>R2
160	1400	37.97	23.41	14.56	R1>R2

2.3.5.2 Hypothesis Test Results

The data in on which Tables 2-9 to 2-18 are based is used to test the hypotheses listed in section 2.3.4.2. The results based on α = 0.05 are as follows.

1. H_0 # 11: The average performance of the one-LSP scheme is not significantly different than that of the two-LSP scheme.

    a. H_0 is rejected on network N0, based on Table 2-8. Load balancing improves performance especially for 80 OD pairs with the largest improvement of 8.91 and 11.75 for split-ratios of 0.5 and 0.9, respectively. The improvement is less significant and practically insignificant with 160 and 320 OD pairs.

    b. H_0 is rejected on networks N1-N5, based on Tables 2-9 to 2-13. The two-LSP scheme outperforms the one-LSP scheme and provides some improvement for certain demands and no improvement for other demands. For demands of 460, 580, and 620 Mbps there is no improvement. For demands of 500, 540, and 660 Mbps the improvement is in the range of 2.53% to 6.28% with a split ratio of 0.5.

    c. H_0 is rejected on networks N6-N10, based on Tables 2-14 to 2-18. The two-LSP scheme outperforms the one-LSP scheme and provides some improvement for certain demands and no improvement for other demands. The most significant improvement is achieved with large demands that exceed half the link's bandwidth. For example, the performance improvement on network N9 with demand of 1300 Mbps is 17.83%.

2. H_0 # 12: The demand per OD pair does not affect the average performance improvement of the two-LSP scheme over the one-LSP scheme.

    a. H_0 is rejected on network N0, based on Table 2-8. The demand affects the average improvement but not in a particular pattern; the improvement is not monotonic as the demand increases. However, for some demands the improvement can be explained by performing simple arithmetic of calculating the number of OD pairs that can be packed on a link and how load balancing affects that number. For example, two OD pairs with a demand 300 Mbps each can be packed on a 622 Mbps link. Splitting in this case does not provide any benefit. In contrast, only one OD pair with a demand of 320 Mbps can be accommodated on 622 Mbps link. Splitting such a demand with a ratio of 0.9, results in two demands of 288 and 32 Mbps. Now, two demands of 288 Mbps can be packed on the same link. The table shows that in this case the improvement is 11.75%.

    b. H_0 is rejected on networks N1-N5, based on Tables 2-9 to 2-13. The demand affects the average improvement, but not in a particular pattern; the improvement is not monotonic as the demand increases.

    c. H_0 is rejected on networks N6-N10, based on Tables 2-14 to 2-18. The demand affects the average improvement, but not in a particular pattern; the improvement is not monotonic as the demand increases. It is noticeable that for four out of the seven demands tested there is no improvement. It appears that the improvement behaves as a step function and depends on the ratio of link demand size to link bandwidth. For example, with a demand of 800

Mbps, which is approximately less than one third of the link bandwidth, there is no improvement. But with a demand of 900 Mbps, which exceeds one third of the link bandwidth, the improvement is significant. The same behavior occurs with 1200 Mbps, which is less than half the bandwidth, and a demand of 1300 Mbps, which is higher than half the bandwidth.

3. H_0 # 13: The split ratio does not affect the average performance of the two-LSP scheme.
   a. H_0 is rejected on network N0, based on Table 2-8 when comparing averages for a specific demand. However, H_0 is not rejected when comparing averages over all demands. The split ratio significantly affects performance but no one ratio provides higher performance for all demands. For some demands, 0.5 provides larger improvement, for other demands 0.9 provides the larger improvement, and yet for other demands there is no significant difference. The average improvement over all the demands with 80 OD pairs using split ratios of 0.5 and 0.9 are 4.14% and 4.15%, respectively.
   b. H_0 is rejected on networks N1-N5, based on Tables 2-9 to 2-13. The split ratio significantly affects performance but no one ratio provides higher performance for all demands. For some demands, 0.5 provides larger improvement, for other demands 0.9 provides the larger improvement, and yet for other demands there is no significant difference. When comparing averages over all demands and all networks, 0.5 provides less than 1% improvement, which is statistically significant but practically insignificant.

However, for all the five networks the highest performance improvement is achieved using split ratio of 0.5.

4. H_0 # 14: The number of OD pairs does not affect the average performance improvement of the two-LSP scheme over the one-LSP scheme.
   a. H_0 is rejected on network N0, based on Table 2-8. The performance improvement is significant with 80 OD pairs but less significant or practically insignificant with 160 and 320 OD pairs. With a split ratio of 0.5, the highest performance improvements achieved for 80, 160, and 320 OD pairs are 8.91%, 2.77%, and 1.05%, respectively.
   b. H_0 is rejected on networks N1-N3, based on Tables 2-9 to 2-11. The performance improvement is significant with 80 OD pairs but less significant or practically insignificant with 160 OD pairs. With a split ratio of 0.5, the highest performance improvements achieved for 80 and 160 OD pairs are 6.28% and 1.73%, respectively.
5. H_0 # 15: Network topology does not affect the average performance improvement of the two-LSP scheme over the one-LSP scheme.
   a. H_0 is not rejected on networks N1-N5, based on Tables 2-9 to 2-13. The topology does not significantly impact the performance improvement. The five networks exhibit comparable improvement and similar dependency on demand size and split ratio.
   b. H_0 is rejected on networks N6-N10, based on Tables 2-14 to 2-18. Topology does not seem to affect the general pattern of dependency on demand size, but it significantly affected the magnitude of improvement, but

especially for the two demands that exceeded half the link bandwidth. For example, with a demand of 1300 Mbps the improvement on network N9 is 17.83% while on network N7 it's only 9.39%. If the large demands were excluded, topology would not significantly impact performance improvement.

2.4 Summary and Conclusions

A new formulation of the basic traffic-engineering problem in MPLS-based packet networks is offered. The problem is formulated as a multi-commodity flow problem with side constraints. The optimization model maximizes revenue and determines which demands are admitted, hence also solving the admission control problem. Computational experiments are conducted to evaluate the benefits of the optimization model in comparison with an online FCFS strategy. The impact of multiple factors on the performance of both strategies and on the performance improvement of optimization over the FCFS strategy is examined. The factors include the number of OD pairs, average and range of demand per OD pair, network topology, and average node degree.

The results for instances of model TE1 show that optimization outperforms the FCFS strategy but offers only low to moderate gains. The improvement ranges from 7.78% to 13.77% for network N0 and from 0.19% to 9.06% for networks N1-N10; in some cases little improvement is observed. All factors significantly impact the performance of both strategies in similar fashion, and therefore have a less significant impact on performance improvement. No factor increased the performance improvement by more than 5%. The improvement increases as the traffic demand increases but only modestly: for example, doubling the demand for network N0 results in an increase of approximately 5%. Generally, optimization provided improvement as the load increased and the average link utilization approached or exceeded 60%. Both strategies perform better with higher numbers of OD pairs, when $\bar{d}$ is low. This is a favorable result for actual implementation in real networks, where demand is likely to exist between any pair of nodes and a full mesh of LSPs is configured. The two design factors, network topology and node degree, have a significant impact on performance. The results demonstrate that doubling the node degree while reducing the link bandwidth by half (so that the total network capacity is the same) results in a more efficient network.

A load-balancing model, TE2, permits the demand of any OD pair to be split (using a predetermined ratio) among two LSPs. Computational experiments are conducted to compare TE2 (with different split-ratios) with TE1, the basic model with a single LSP per OD pair. The impact of multiple factors on the performance of both models and on the performance improvement by the load-balancing model is examined. The factors include the number of OD pairs, demand per OD pair, and network topology.

The results show that load balancing provides low to moderate improvement. In the case of network N0, an improvement of up to 11.75% is achieved depending on the demand per OD pair. The improvement increases with $\bar{d}$, but not monotonically. The improvement also depends on the ratio of link capacity to the demand per OD pair (i.e., the number of OD pairs that can be packed within a link's bandwidth). The most significant improvement of 17.83% (on network N9) is achieved when the demand per OD pair exceeds half the link's bandwidth. The larger the number of OD pairs (and assuming lower demand per OD pair) the smaller the improvement. The results also show that no single split-ratio dominates the others for all demand levels. Finally, network topology does not seem to impact the general behavior but can impact the magnitude of improvement itself.

In conclusion, MPLS enables network operators to perform traffic engineering and maximize resource utilization. It is recommended that the network operators implement a full mesh of LSPs, which will provide the immediate benefit of constructing a traffic matrix that can be used to study performance and influence the network design. The online FCFS strategy, which is currently implemented on network equipment, provides satisfactory performance in situations with light or moderate demand loads. As the network load increases and average link utilization exceeds 50%, it is recommended that network operators consider deploying an offline optimization strategy in order to further increase resource utilization. Initially, the basic optimization model can be deployed. As demand increases, the network operator should consider implementing the load-balancing model. In that case, operators should experiment with different split ratios to determine an optimum ratio for the traffic conditions. Instead of load balancing all demands, only large demands may need to be split.

Chapter 3

FORMULATION AND REVENUE ANALYSIS OF MPLS TRAFFIC ENGINEERING WITH MULTIPLE CLASSES OF SERVICE

This chapter presents a new formulation of the problem of Multi-Protocol Label Switching (MPLS) traffic engineering with multiple class-of-service (CoS) types. Each CoS type is assigned a priority, has its own quality-of-service (QoS) requirements, and is priced differently. The problem addresses admission control and constraint-based routing. The objective is to design multiple transmission paths that satisfy the QoS requirements of the different CoS types while maximizing revenue.

The problem is formulated as an origin-destination (OD) integer multi-commodity network flow problem with side constraints; the problem and the corresponding model are referred to as TE-MC. The results of computational testing with two service classes are presented. The objective of the computational testing is to perform revenue analysis and assist in assigning the relative increase in revenue per unit of the higher-class traffic (refer to Section 3.3.2).

This chapter is organized as follows. Section 1 is an introduction and includes motivation for the research topic, a statement of the problem, and a survey of related literature. The literature regarding traffic engineering with a single class of service is surveyed in section 2.1.2. Section 2 includes a mathematical formulation of the TE-MC model. Section 3 presents the methodology for the computational experiments, the tabulated results, and analysis of the results including breakeven price analysis. Section 4 summarizes the results and provides conclusions and recommendations.

3.1 Motivation

MPLS was originally developed to overcome the Internet Protocol (IP) deficiencies and facilitate Traffic Engineering (TE) in IP networks. MPLS provides the mechanisms needed to control traffic flows in IP networks, overcomes the limitations of shortest-path-only routing and allows the creation of traffic-engineered paths, leading to more efficient bandwidth utilization. MPLS has evolved to support other applications, in particular, virtual private networks (VPNs). Enterprise customers are increasingly adopting MPLS-based VPNs to replace existing wide area network (WAN) technologies such as Frame Relay. Such VPNs provide any-to-any connectivity among geographically dispersed customer sites, and support multiple CoS types with QoS performance guarantees. Enterprise customers are embracing MPLS-based VPNs to consolidate separate communications infrastructures, and converge multiple applications including data, voice, and video onto a common infrastructure. Network convergence is expected to provide enterprise customers with cost savings and increase operational efficiencies.

The convergence of voice, video, and data traffic requires the service provider MPLS/IP network to support multiple traffic types with dissimilar characteristics and requirements. Both voice and video traffic characteristics differ from those of data traffic. Such traffic is very sensitive to delay and delay variations, but can tolerate some degree of packet loss. Data traffic is typically elastic — can tolerate delay — and is oblivious to delay variations, but is sensitive to packet loss. So the current IP paradigm of best-effort service, where all packets are treated equally with no regard to the needs of applications for some level of resource assurance or performance guarantees, is no longer adequate. MPLS-based VPNs support service differentiation and provide multiple CoS types with QoS performance guarantees. Typically, QoS traffic performance parameters include guaranteed bandwidth availability and bounds on packet delay, packet delay variation, and probability of packet loss.

The support of multiple CoS types adds another dimension to the basic MPLS TE problem (TE1) discussed in Chapter 2. Instead of having a single LSP to carry aggregate traffic for an OD pair, traffic is now partitioned into multiple classes and the traffic of each class is carried through a distinct LSP. Each LSP has its own bandwidth and performance requirements and is routed separately. The routing also considers the priority of the CoS type that is mapped into that LSP. Also link bandwidth may be partitioned and pre-allocated for the different classes.

3.1.1 Problem Statement

TE-MC can be stated as follows. Given are the physical topology and link attributes of an MPLS network. The link attributes include capacity and an administrative cost that reflects delay on that link. Also given are multiple CoS types; each CoS is assigned a priority, and specifies its own QoS performance requirements. The priority of a CoS is preemptive; traffic with higher priority is admitted and routed first. Given is the traffic matrix, which represents aggregate traffic demand between any OD pair for each CoS. The objective is to maximize revenue by admitting and routing as much traffic as possible in CoS priority order while observing the resource and traffic performance constraints. The resource constraints are the link capacities. The traffic performance constraints are specific for each CoS and are typically expressed as the maximum allowable number of hops and maximum allowable delay for traffic between any OD pair. TE-MC is a logical design problem that involves admission control and the design of a set of paths to route traffic of different CoS types.

Another aspect of the TE-MC problem is assigning the revenue per unit of demand for each CoS type. Pricing of CoS types is driven by multiple factors including customer demand, market pressure, and the cost of providing the service. There are multiple components to the cost; one of these components reflects the cost of the network resources (bandwidth) allocated for traffic of each CoS type. The general premise is that traffic with a higher CoS has more stringent performance requirements, which further constrain network routing and may affect the total traffic that can be delivered. Therefore, a higher revenue per unit of demand should be assigned to a higher CoS type. This chapter explores this aspect of pricing and provides an approach for assigning the relative increase in revenue of a unit of demand for a higher CoS type.

3.1.2 Survey of Related Literature

The Internet Engineering Task Force (IETF) has developed a series of requests for comments (RFCs) specifications that cover various aspects of MPLS, TE, and Differentiated Services (Diffserv). A few important RFCs are briefly summarized here. RFC 3031 [39] specifies the architecture for multi-protocol label switching, and defines the functions performed by an MPLS-capable router, which is referred to as a label switch router (LSR). The RFC specifies both the control and traffic forwarding functions. RFC 2702 [7] presents a set of requirements for traffic engineering over MPLS. "It identifies the functional capabilities required to implement policies that facilitate efficient and reliable network operation in an MPLS domain." RFC 3209 [5] specifies extensions to the Resource Reservation Protocol (RSVP) and describes its use as a label distribution and signaling protocol to establish explicitly routed LSP tunnels in MPLS networks. Explicitly routed LSPs are used for traffic engineering purposes. RFC 3630 [29] describes extensions to the Open Shortest Path First (OSPF) routing protocol to support intra-area traffic engineering. The extensions enhance the protocol to support resource-based routing in addition to topology-based routing.

RFC 2475 [13] defines a scalable architecture for implementing differentiated services in the Internet to support multiple classes of service. The scalability is achieved by aggregating packet classification. At network boundaries packets are classified and marked accordingly using the Differentiated Services (DS) field in the IP header. The marking determines the per-hop forwarding behavior (PHB) a packet receives along its path. RFC 2475 defines the DS field in the IP header. RFC 2597 [26] defines an Assured Forwarding (AF) PHB group and the associated encoding of the DS field. RFC 3246 [18] defines an Expedited Forwarding (EF) PHB group and the associated encoding of the DS field. The AF classes are suited for data applications while the EF class is designed to support voice. RFC 3270 [30] defines the mechanisms for supporting Diffserv over MPLS networks. The RFC defines two types of Label Switched Path (LSP), E-LSP and L-LSP. E-LSP can transport traffic of multiple classes, and the EXP field in the MPLS header identifies the scheduling behavior and drop precedence associated with the corresponding class. L-LSP transports traffic of a specific class; the label itself indicates the scheduling behavior associated with the class and the EXP field determines the drop precedence of a packet. The RFC specifies the operation of both types of LSPs in supporting Diffserv classes. RFC 3564 [31] specifies requirements for supporting Diffserv-aware traffic engineering (DS-TE). Instead of performing traffic engineering at an aggregate level (all traffic per OD pair), DS-TE performs traffic engineering at a per-class level by mapping each Diffserv class of service to a separate LSP (i.e., L-LSP). DS-TE allows finer-grain optimization of network resources and better meets the performance requirements of each class.

Srikitja and Tipper [41] formulate the problem of designing VPNs in MPLS networks using multiple multipoint-to-point (m-t-p) trees or LSPs. The design is modeled as an MIP problem that minimizes the bandwidth cost of laying out VPNs supporting multiple CoS types while meeting the QoS requirements. The model utilizes a path formulation (trees are pre-computed) and also incorporates multi-hour support to accommodate traffic variations during a day. The use of m-t-p LSPs reduces the number of LSPs as compared to point-to-point (p-t-p) LSPs and can provide bandwidth savings. Computational results show that without bandwidth aggregation, the m-t-p design yields results similar to p-t-p full-mesh design. However, when bandwidth of different VPNs is aggregated and traffic is not necessarily routed along shortest-path trees, significant reduction in bandwidth can be achieved.

Mitra and Ramakrishnan [34] formulate the problem of MPLS traffic engineering with multiple CoS types using the multi-commodity flow problem maximizing revenue. The paper distinguishes between CoS types that have QoS requirements and a best-effort class with no QoS requirements, formulates the problem for each type, and then combines them to incorporate priorities. The admission control/routing for the QoS traffic utilizes a route-based (path-based) formulation while the best-effort traffic utilizes a link-based formulation. They argue that the route-formulation is more suitable for QoS traffic since it is easier to incorporate policy constraints for different classes and for different OD pairs. A set of admissible routes for each OD pair and service class is pre-computed. The formulation allows for load balancing and partial delivery of demand. The formulation incorporates priority by using a multi-critera objective function where the first objective is to maximize the revenue of QoS traffic and the second objective is to maximize revenue from best-effort traffic using residual capacity. The paper describes techniques for solving the combined problem using multi-layer decomposition and presents computational results on a sample network. The formulation is also extended to deal with stochastic traffic patterns.

3.2 Mathematical Formulation

This section presents the mathematical formulation of the model for MPLS traffic engineering with multiple classes of service. The model is formulated as an OD integer multi-commodity network flow problem with side constraints.

3.2.1 Notation and Conventions

The notation and conventions of Section 2.2.1 are used. Using that notation, the following symbols are used in the formulation of TE-MC:

N the set of node indices 1, 2, …, |N| in the network

A the set of directed arcs in the network; an arc is represented as an ordered pair (i, j) where $i, j \in N$.

b_{ij} the capacity, in units of bandwidth, of arc (i, j), $b_{ij} \in R$

c_{ij} the administrative cost associated with arc (i, j), $c_{ij} \in R$; typically the administrative cost represents a measure of delay or transmission time

n the number of CoS types, each with its own QoS requirements

q a CoS type, where $q = 1, 2, \ldots, n$

K_q the set of commodities with CoS type q; each commodity $k_q \in K_q$ is represented by the attributes O_{k_q}, D_{k_q}, and d_{k_p}

O_{k_q} the origin node of commodity $k_p \in K_n$ where $O_{k_q} \in N$

D_{k_q} the destination node of commodity $k_p \in K_n$ where $D_{k_q} \in N$

d_{k_p} the demand of commodity $k_q \in K_q$ in units of bandwidth

ℓ_q the maximum allowed delay (latency) that any commodity of CoS type q may incur while traversing the network from source to destination[1], $\ell_q \in R$

h_q the maximum allowed number of hops that any commodity may traverse from source to destination[2], $h_q \in Z$

μ_q a unit of revenue generated from delivering a unit of demand of any commodity of CoS type q, $\mu_q \in R$

ω a scaling factor or a weight used in the objective function, $\omega \in R$

3.2.2 MPLS Traffic Engineering with Multiple Classes of Service Model

Using the above notation the problem of MPLS traffic engineering with multiple classes of service can be stated as follows. Given is a graph $G = (N, A, \mathbf{b}, \mathbf{c})$ that describes the physical topology and link attributes of an MPLS network and n CoS types,

[1] This formulation can be generalized to associate separate maximum delay with each commodity.
[2] This formulation can be generalized to associate separate maximum number of hops with each commodity

each with its own priority and QoS requirements. The priority of CoS type q is represented by the value of q; the lower the value the higher the priority. The traffic performance requirements of a CoS type q are represented by the maximum delay ℓ_q and the maximum number of hops h_q. Also given is the node-to-node traffic for each CoS type q represented by the set of commodities K_q. The objective is to maximize revenue by admitting and routing the maximum traffic possible while observing the priority of traffic (higher priority traffic is admitted and routed first), and meeting the resource and traffic- performance constraints. The resource constraint is enforced by not allowing the combined traffic across all commodities on any link (i, j) to exceed the capacity b_{ij} of that link. The performance constraints are expressed in terms of allowable maximum delay and allowable maximum number of hops. A commodity of CoS q is to be routed along a path whose delay and number of hops do not exceed ℓ_q and h_q, respectively. The following decision variables are defined:

$$y^{k_q} = \begin{cases} 1, & \text{if commodity } k_q \text{ is routed,} \\ 0, & \text{otherwise,} \end{cases}$$

$$x_{ij}^{k_q} = \begin{cases} 1, & \text{if commodity } k_q \text{ is routed through a path that uses arc } (i, j), \\ 0, & \text{otherwise.} \end{cases}$$

The multi-class problem can be formulated as follows:

[TE-MC]

$$\text{Maximize} \sum_{q} \mathrm{P}_q \left(\sum_{k_q \in K_q} \mu_q d_{k_q} y^{k_q} - \omega \sum_{(i,j)\in A} c_{ij} \sum_{k_q \in K_q} d_{k_q} x_{ij}^{k_q} \right) \tag{3.1}$$

subject to,

$$\sum_{j \in N|(i,j) \in A} x_{ij}^{k_q} - \sum_{j \in N|(i,j) \in A} x_{ji}^{k_q} = \left\{ \begin{array}{cl} y^{k_q}, & if\ i = O_{k_q} \\ -y^{k_q}, & if\ i = D_{k_q} \\ 0\ , & otherwise \end{array} \right\} \quad \forall q, \forall k_q \in K_q, \forall i \in N, \tag{3.2}$$

$$\sum_{q} \sum_{k_q \in K_q} d_{k_q} x_{ij}^{k_q} \le b_{ij} \quad \forall (i, j) \in A, \tag{3.3}$$

$$\sum_{(i,j) \in A} c_{ij} x_{ij}^{k_q} \le \ell_q \quad \forall q, \forall k_q \in K_q, \tag{3.4}$$

$$\sum_{(i,j) \in A} x_{ij}^{k_q} \le h_q \quad \forall q, \forall k_p \in K_n, \tag{3.5}$$

$$x_{ij}^{k_q}, y^{k_q} \in \{0,1\} \quad \forall (i, j) \in A, \forall q, \forall k_q \in K_q. \tag{3.6}$$

where P_q indicates the priority level, with P_1 being the first (highest) priority.

The above model is represented as preemptive priority model where the CoS type signifies the priority. The solution determines which commodities will be routed and the path for each commodity. The objective function is the sum over all classes of the expression in parenthesis, which consists of two terms with the first one being the primary objective and the dominant term. The first term represents the total revenue generated from routed commodities (i.e., delivered traffic) of a CoS type *q*. If a commodity is routed, its entire demand is delivered. The total demand delivered, and not necessarily the total number of commodities delivered, is maximized. The second term represents the total delay incurred by all the delivered commodities of a CoS type *q*. The purpose of this term is to select the solution with the lowest delay among multiple alternate optimum solutions (yielding the same revenue) that may exist. The total delay is multiplied by the scaling factor ω, $0 \leq \omega \leq 1$. Typically, ω is set so that the second term evaluates to small values and the first term remains dominant.

Constraints (3.2) are the flow-conservation equations, which ensure a connected path for each routed commodity. For each commodity, there are $|N|$ equations: one equation for the source node of that commodity, one for the destination node, and $|N|$ - 2 equations for the other nodes in the network, which act as transit nodes. Flow for commodity $k_q \in K_q$ is routed only if binary variable $y^{k_q} = 1$; otherwise, that particular demand is not satisfied. Unit supplies and demands, and binary flow variables cause a single LSP to be formed via these constraints.

Constraint set (3.3) enforces the arc-capacity constraints. For each arc, the total traffic from all commodities (of all CoS types) whose paths include that arc cannot exceed the arc's bandwidth. Constraints sets (3.4) and (3.5) represent the performance requirements. Constraint set (3.4) ensures that the delay along any path cannot exceed a predetermined upper delay limit ℓ_q for a commodity of CoS type q. Constraint set (3.5) ensures that the number of hops along any path cannot exceed a predetermined upper hop limit h_q for a commodity of CoS type q. This constraint approximates setting a bound on delay variations. The major cause for delay variations is queuing at different nodes. Including queuing delays would complicate the model and introduce non-linearties [2, 3]. Limiting the number of hops is expected to reduce delay variations. Constraint set (3.6) constrains the decision variables to be binary.

3.3 Computational Experiments

Computational experiments are performed with the TE-MC model and two classes of service (q = 2), referred to as classes A and B. Class A has the higher priority and more stringent traffic performance requirements (lower delay and number of hops). The objective of the experiments is to determine the percentage of incremental revenue per unit of demand that class A traffic needs to generate. For that purpose, the performance of the TE-MC model is compared with that of the single-class model TE1 (see section 2.2.2). The total revenue generated by the demand delivered by TE1 is compared to the sum of revenues generated by class A and B traffic in model TE-MC. Then, a breakeven analysis is performed to calculate the percentage of incremental revenue per unit of demand for class-A traffic.

3.3.1 Organization of Tests

3.3.1.1 Test Network Characteristics

The tests are performed on the realistic network N0, depicted in Figure 2-1, which represents a typical topology of nationwide data communications networks implemented by inter-exchange providers in the U.S. N0 consists of 20 nodes and 62 arcs with an average node degree of 3.1. The trunks connecting the nodes are bi-directional and full duplex; hence each trunk is represented as two directed arcs, each with the same capacity and cost. The thick trunks represent OC-48 transmission lines with 2488 Mbps of bandwidth capacity, and the thin trunks represent OC-12 transmission lines with 622 Mbps of bandwidth capacity. The arc cost reflects the actual circuit mileage of the corresponding transmission line. Since the shortest-path algorithm uses the arc cost as the metric to calculate the shortest path the OC-12 cost is increased by 1000 units so that an OC-48 trunk will be preferable to an OC-12 trunk.

3.3.1.2 Traffic Generator

A traffic generator was developed to generate multiple sets of commodities for the different experiments. The generator accepts as input parameters the number of nodes in a network, the number of OD pairs, and the minimum and maximum demands. The generator selects OD pairs randomly and uniformly from the set of nodes and ensures that there are no duplicate OD pairs. The generator randomly selects the demands associated with OD pairs using a uniform distribution over the range specified by the minimum and maximum demands. The minimum and maximum demands are selected based on the test scenario. For example, when testing with N0 the maximum demand is always less than the 622 Mbps capacity of the lower bandwidth links.

3.3.1.3 Computing Environment

All test cases are performed on a Compaq AlphaServer DS20E with dual 667 MHz processors and 4096 MB RAM. The machine is configured as a Network Queuing System and executes batch jobs. Each job on the system has access to approximately 2 GB RAM. The models are implemented using the GAMS [14] model description language, and integer-programming solutions are generated using CPLEX Linear Optimizer 7.0. Default settings for CPLEX are used with the exception that the MIP time limit is set to 3600 seconds and the relative optimality gap is set to 1%.

3.3.2 Experiment Design

The experiments are designed for the purpose of calculating the relative increase in revenue per unit of demand that class-A traffic needs to generate. This is motivated by a common industry situation in which a service provider currently offers a single class of service and intends to change to two classes of service. Class A will offer higher QoS performance and class B will offer the same performance as the current class. Since class-A traffic has more stringent performance requirements, its routing is more constrained. This may reduce the total traffic that can be routed over the existing network without any capacity augmentation, depending on the total load on the network. With a lightly loaded network, the stricter routing of the class-A traffic is unlikely to affect the delivery of the class-B traffic. However, as the overall load increases, the total demand of both classes of traffic that can be delivered is expected to be less than the total load that can be delivered with a single class with less stringent performance requirements. To compensate for the potential loss in revenue, the service provider is likely to offer the class-A service at a higher price.

The following breakeven analysis is performed to establish a base for determining the relative increase in revenue per unit of demand of class A. The total revenue generated by class A and B traffic in model TE-MC is required to equal the revenue generated by the single-class model TE1 when the total traffic load is equal. We denote the revenue generated by single-class, class-A, and class-B traffic as r, r_1, and r_2, respectively. Then, $r = r_1 + r_2$, where,

$$r = \mu \sum_{k \in K} d_k y^k, \quad r_1 = \mu_1 \sum_{k_1 \in K_1} d_{k_1} y^{k_1}, \text{ and } r_2 = \mu_2 \sum_{k_2 \in K_2} d_{k_2} y^{k_2}$$

Assuming that $\mu = \mu_2$, and substituting the above terms in the equation, the relative percentage increase in revenue per unit of demand of class A, $\mu_{1\%}$, is calculated as follows:

$$\mu_{1\%} = \left(1 - \frac{\mu_1}{\mu}\right) \cdot 100\% = \left(1 - \frac{\sum_{k \in K} d_k y^k - \sum_{k_2 \in K_2} d_{k_2} y^{k_2}}{\sum_{k_1 \in K_1} d_{k_1} y^{k_1}}\right) \cdot 100\% \tag{3.5}$$

The computational experiments calculate $\mu_{1\%}$ for different network loads. Using the traffic generator, a set of commodities for a given number of OD pairs and demand per OD pair is generated. The model TE1 is run with $\mu = 1$ and a solution for the single-class is obtained. Then, for the same set of OD pairs, the demand per OD pairs is split (using a predetermined ratio) into two demands, one for each class. So, the cardinality of the new set of commodities is twice the original one, but the total traffic demand remains the same. The model TE-MC is run with the new set and with $\mu_1 = \mu_2 = 1$ and a solution for the two-class problem is obtained. Using the two solutions, $\mu_{1\%}$ is determined from (3.5). This procedure is repeated for a series of demands.

3.3.3 Experiment Results

The experimental results are presented in Tables 3.1 and 3.2. The tables compare the performance of the single-class model (TE1) with the two-class model (TE-MC with $q = 2$) and show the calculated $\mu_{1\%}$ values for different demands. The total traffic demand is equal for both models. The OD Demand column shows the demand per OD pair for the single-class model. In these experiments we assumed that 50% of the traffic is designated as class A and 50% as class B. So a single-class demand per OD pair is split into two equal demands for the same OD pair. The performance of the two models is compared in terms of percent traffic delivered and bandwidth utilization. For the single-class model, the demand delivered ratio is calculated as the ratio of total demand delivered to total demand. For the two-class model, the demand delivered ratio is calculated for each class separately. The bandwidth utilization is defined as the ratio of total flow on all arcs to the total bandwidth of all arcs. Each row in a table represents the averages for 20 replications of an experiment.

Table 3.1 shows the results for a half mesh of 190 single-class, 190 class-A, and 190 class-B commodities for the same 190 OD pairs. Table 3.2 shows the results for a full mesh of 380 single-class, 380 class A, and 380 class B commodities. The demand per OD pair in Table 3.2 is half the corresponding demand in Table 3.1. By maintaining the same total demand for the two different numbers of OD pairs, the results in the two tables can be compared and the impact of the number of OD pairs on performance and on the value of $\mu_{1\%}$ can be examined.

Table 3-1 Average performance of TE1 and TE-MC and $\mu_{1\%}$ as function of demand for 190 OD pairs

OD Demand	One Class		Two Classes			% Class A Revenue Increase ($\mu_{1\%}$)
	Demand Delivered %	Band. Util. %	Class A Delivered %	Class B Delivered %	Band. Util. %	
100	99.4	60.4	100	98.7	60.5	0.0
120	94.1	65.5	100	89.4	66.8	-1.2
140	86.4	66.7	100	72.6	69.0	0.2
160	79.5	68.6	99.8	59.6	71.1	-0.4
180	75.6	70.2	99.3	44.1	71.5	7.8
200	73.5	73.7	99.4	37.1	75.4	10.6
220	65.5	71.2	97.3	25.2	75.5	8.8
240	63.4	72.6	94.1	24.5	77.3	8.6
260	61.3	73.9	89.9	19.9	77.3	14.2
280	59.1	74.5	86.4	19.1	77.0	14.6
300	59.1	80.8	84.4	18.9	79.1	17.5
320	46.8	65.5	79.5	13.7	76.5	0.7
340	46.8	69.7	77.6	12.9	77.4	4.1
360	44.4	66.3	75.6	13.5	78.1	-0.3
380	44.2	70.0	75.6	12.5	82.3	0.3
400	44.4	73.7	73.5	12.0	81.9	4.6

Table 3-2 Average performance of TE1 and TE-MC and $\mu_{1\%}$ as function of demand for 380 OD pairs

OD Demand	One Class		Two Classes			% Class A
	Demand Delivered %	Band. Util. %	Revenue Increase ($\mu_{1\%}$)	Class B Delivered %	Band. Util. %	Revenue Increase ($\mu_{1\%}$)
50	100	60.4	100	100	60.5	0
60	95.8	67.1	100	91.6	67.7	0
70	87.9	68.0	100	77.9	70.6	-2.11
80	82.6	69.9	100	66.3	73.8	-1.05
90	76.8	70.8	100	54.2	76.5	-0.53
100	74.2	72.7	100	42.6	78.3	5.79
110	71.1	76.3	100	30.5	79.7	11.58
120	68.9	77.9	95.8	30.0	81.5	12.64
130	65.5	81.7	92.1	26.3	82.0	13.71
140	63.4	82.0	87.9	24.2	81.7	16.77
150	62.4	86.3	85.8	24.2	84.2	17.18
160	57.1	82.6	82.6	19.5	81.9	14.65
170	55.8	83.8	80.5	21.6	84.9	11.76
180	54.2	83.8	77.4	17.9	83.6	17.01
190	54.2	88.6	76.3	17.9	86.4	18.62
200	52.9	88.8	74.2	18.4	88.1	17.73

3.3.4 Analysis of Results

The results show that both TE1 and TE-MC perform better (in terms of demand delivered) with larger numbers of OD pairs and correspondingly smaller demand per OD pair. With TE1, the average (over all demand sizes) ratio of demand delivered for 380 and 190 OD pairs is 70.2% and 65.2%, respectively. With TE-MC, the average ratio of class-A demand delivered for 380 and 190 OD pairs is 90.8% and 89.5%, respectively. The effect of larger numbers of OD pairs with smaller demand is less significant since class-A traffic is routed first while adequate bandwidth is available. However, the impact on the average delivery ratio of class-B traffic is significant. Since class-B traffic is routed after class-A traffic has been routed, there is less bandwidth available and the network is capable of routing more of the smaller demands. The average ratio of class-B demand delivered for 380 and 190 OD pairs is 41.4% and 35.9%, respectively.

The tables also show that for most demand sizes TE1 delivers more traffic than TE-MC, yet the system-wide bandwidth utilization by TE-MC is higher. Based on Table 3.1, the average ratio of demand delivered with TE1 and TE-MC is 65.2% and 62.7%, respectively; the average bandwidth utilization with TE1 and TE-MC is 70.2% and 74.8% respectively. So while the ratio of demand delivered decreased by 2.5% the average utilization increased by 4.6%. Similarly, based on Table 3.2, the average ratio of demand delivered with TE1 and TE-MC is 70.2% and 66.1%, respectively; the average bandwidth utilization with TE1 and TE-MC is 77.5% and 78.8%, respectively. So while the ratio of demand delivered decreased by 4.1%, the average utilization increased by 1.3%. These results demonstrate the penalty (from the service provider perspective) associated with delivering the higher-class traffic: the total traffic that can be delivered decreases while the bandwidth consumption increases.

To compensate for this increase in resource utilization and protect against potential loss in revenue, the revenue per unit of demand of class-A traffic needs to be increased. The calculated $\mu_{1\%}$ values for the different demands are presented in the last column of both tables. A value of 0 indicates that the same amount of demand is delivered with the TE-MC model as with TE1 model; therefore there is no resource-based justification for increasing revenue per unit of demand of class-A traffic. A negative value indicates that TE-MC delivered more total demand than TE1; this implies that the revenue per unit of demand of class-A traffic may be discounted but in reality this is counter-intuitive and will not be done. Alternatively, this can be viewed as a gain in network efficiency. A positive value indicates that TE-MC delivered less total demand than TE1; therefore the revenue per unit of demand of class A needs to be increased to offset the decrease in revenue due to lower delivery of class-B traffic.

Examining the results in both tables shows that $\mu_{1\%}$ is positive for most of the demand sizes. However, the presence of the unexpected negative values of $\mu_{1\%}$ and the fact that $\mu_{1\%}$ does not monotonically increase as the aggregate demand increases suggests that other factors influence this value. In the case of 190 OD pairs, the value of $\mu_{1\%}$ falls into three ranges. For demands from 100 to 160 Mbps, $\mu_{1\%}$ is either 0 or negative. For demands from 180 to 300 Mbps, $\mu_{1\%}$ is positive; it significantly increases at 180 Mbps and then gradually and non-monotonically increases as demand increases up to a demand size of 300 Mbps. For demands from 320 to 400 Mbps, $\mu_{1\%}$ is positive; it significantly increases at 180 Mbps and then gradually and non-monotonically increases as demand increases up to a demand size of 300 Mbps. The highest $\mu_{1\%}$ value in this range is 17.5%.

The values of $\mu_{1\%}$ decrease sharply in the range from 320 to 400 Mbps. This decrease is notable and provides the hint for the other factor involved. The demand of 320 Mbps exceeds half the bandwidth of an OC-12 link, which is 622 Mbps. The study of load balancing in Chapter 2 shows that splitting such a demand into two smaller demands help increase the total demand delivered. This is also consistent with the above result that the demand delivery ratio is higher with a larger number of demands and smaller demand size. So, splitting the 320 Mbps (or above) demand in half into two demands for the two classes and routing them separately achieves the benefits of load balancing.

In the case of 380 OD pairs, the general trend is that $\mu_{1\%}$ increases as the demand increases, but not monotonically. The highest $\mu_{1\%}$ value is 18.62%. The load-balancing effect is less significant in this case. Since demand sizes are already small, splitting them into even smaller demands has less impact on increasing the total traffic delivery ratio. This is again consistent with the results obtained in Chapter 2.

3.4 Summary and Conclusions

A new formulation of the problem of MPLS traffic engineering with multiple classes of service is presented. The problem is formulated as an OD integer multi-commodity network flow problem with side constraints. The optimization model is implemented as a preemptive priority model to reflect CoS priorities. The model maximizes revenue and determines which demands are admitted, hence also solving the admission-control problem. Computational experiments are conducted to compare the performance of the two-class model with the single-class model for the realistic network N0. Subsequently, a breakeven revenue analysis is conducted to help determine the value of the relative increase in revenue per unit of demand of class-A traffic, $\mu_{1\%}$.

The results show that the traffic delivery ratio is higher with a larger number of OD pairs and correspondingly smaller demand sizes. Partitioning a demand into multiple demands for different classes and routing them separately achieves the benefits associated with load balancing and increases the traffic-delivery ratio. The results demonstrate that delivering the higher-class traffic may reduce the total traffic delivered and increase bandwidth utilization, depending on the size of demand per OD pair. To compensate for this increase in resource utilization and protect against potential loss in revenue, increasing the revenue per unit of demand of class-A traffic is justified. The calculated $\mu_{1\%}$ values depend on the number of OD pairs, the demand per OD pair, and the effect of load balancing. The results show that $\mu_{1\%}$ can be as high as 18.62%.

In conclusion, supporting multiple classes and utilizing separate LSPs indirectly realizes the benefits of load balancing and increases network efficiency. This result can be exploited by service providers and highlights the value of L-LSPs. Determining $\mu_{1\%}$ is not straightforward, since it depends on traffic load and the number of OD pairs. A network administrator is advised to perform analysis similar to that performed here by taking the existing demand levels, extrapolating them to some realistic future levels, calculating $\mu_{1\%}$ for each level, and then combining statistically (e.g., averaging) to set an overall $\mu_{1\%}$. The resulting value provides a powerful resource-based rational for management's CoS pricing decisions.

Chapter 4

FORMULATION AND EVALUATION OF MPLS TRAFFIC ENGINEERING WITH OVER-SUBSCRIPTION

This chapter presents a new formulation of the problem of Multi-Protocol Label Switching (MPLS) traffic engineering with over-subscription of link capacities. As with airlines' over-booking policies, capacity over-subscription is justified by the stochastic nature of traffic and the likelihood that OD demand variations will accommodate additional network loading. Hence, over-subscription allows demand to be routed that otherwise would have been rejected, but assesses a penalty (reflecting risk) on traffic that exceeds link capacities. The model involves constrained-based routing with the objectives of maximizing revenue and minimizing penalty.

The problem is formulated as an origin-destination (OD) integer multi-commodity network flow problem with side constraints; the problem and the corresponding model are referred to as TE-OS. The results of the computational testing are presented and are used to demonstrate the suitability of the model for capacity planning to accommodate traffic growth.

This chapter is organized as follows. Section 1 is an introduction and includes motivation for the research topic, a statement of the problem, and a survey of related literature. Section 2 includes a mathematical formulation of the TE-OS model. Section 3 presents the methodology for the computational experiments, the application of the TE-

OS model for capacity planning, the tabulated results, and analysis of the results. Section 4 summarizes the results and provides conclusions and recommendations.

4.1 Motivation

The models presented in previous chapters strictly enforce the link-capacity constraints — the total traffic on any link is not allowed to exceed the capacity of that link. Since those models are intended to be used by a service provider for admission control, the capacity constraints ensure that admitted traffic is guaranteed to be delivered. In practice, many service providers do not reject customer demand but attempt to accommodate the "excess" demand (demand exceeding a network's capacity) by increasing network capacity. Since the previous models reject any excess demand, the actual load on links when all excess demand is admitted was not calculated. Therefore, capacity planning requires a different model: one that admits all demand and increases (over-subscribes) some link capacities to accommodate demand. Using over-subscription and admitting all demand helps identify where additional capacity is needed. Links with traffic overflow are the first candidates for capacity augmentation and the higher the load on a link, the greater the need for augmenting that link's capacity.

Over-subscription also acknowledges the benefits of the statistical multiplexing of demand between different OD pairs. In modeling MPLS TE problems as deterministic, multi-commodity network flow problems, it is assumed that the demand matrix represents the constant rate of peak traffic between OD pairs and that all demands are simultaneous. In reality, OD traffic in packet-based networks fluctuates and does not flow at a constant peak rate; moreover, the peaks of different demands do not necessarily coincide. As a result of these traffic fluctuations, the actual total demand at any given time is typically lower than the total peak demand. Traffic fluctuations can be exploited to realize a statistical multiplexing gain; instead of having adequate network capacity to accommodate the total peak demand, lower capacity can accommodate the actual demand. Since the traffic matrix is still given in terms of constant peak rates, to realize the statistical multiplexing gain, link capacities are over-subscribed. In that case, when traffic overflows on a link it does not imply that the excess traffic is dropped with certainty. Instead, the implication is that the excess traffic may be dropped with some probability. Typically, service providers do not guarantee 100% packet delivery and allow for some packet drops. If the packet drops exceed the level specified in the service level agreement (SLA), the customer is entitled to a billing adjustment known as back-credit. So, over-subscription of link capacity is used by service providers to realize statistical multiplexing advantages and maximize the use of existing infrastructure while considering the risk of losing some revenue in the form of back-credit to customers.

4.1.1 Problem Statement

The physical topology and link attributes of an MPLS network are given. The link attributes include capacity and an administrative cost that reflects delay on that link. Also given is the traffic matrix, which represents aggregate traffic demand between any OD pair. It is assumed that the demand is expressed as a peak traffic rate, but the actual demand fluctuates. To handle the fluctuations in demand it is assumed that link capacity is over-subscribed so that total flow on a link is allowed to exceed link capacity. It is also assumed that traffic exceeding link capacity may be dropped and an assessed penalty proportional to the amount of dropped traffic. The objective is to maximize revenue by admitting and routing as much traffic as possible and minimizing dropped traffic with associated penalties while observing the resource and traffic performance constraints. The resource constraint is the over-subscribed link capacities. The traffic performance constraints are typically expressed as the maximum allowable number of hops and maximum allowable delay between any OD pair.

4.1.2 Survey of Related Literature

The literature regarding MPLS traffic engineering with a single class of service is presented in Section 2.1.2. RFC 2702 [7] requires MPLS traffic engineering to support over-subscription of resources. It defines the maximum allocation multiplier (MAM) of a resource as an administratively configurable attribute that determines the proportion of the resource that can be allocated for traffic demand. A resource can be under-allocated or over-allocated.

In the literature surveyed, MPLS traffic engineering with over-subscription has not been addressed as a separate problem. The MAM is incorporated in the basic MPLS TE problem (as discussed in Chapter 2) as a multiplier of link bandwidth for which the link capacity constraint is enforced. Over-subscription has not been previously investigated as a means for capacity planning or for the purpose of estimating traffic loss and determining financial penalties.

4.2 Mathematical Formulation

This section presents a mathematical formulation of the TE-OS problem. The model is formulated as an OD integer multi-commodity network flow problem with side constraints.

4.2.1 Notation and Conventions

The notation and conventions of Section 2.2.1 are used. Using that notation, the following symbols are used in the formulation of TE-OS:

N the set of node indices 1, 2, …, |N| in the network

A the set of directed arcs in the network; an arc is represented as an ordered pair (i, j) where $i, j \in N$.

b_{ij} the capacity, in units of bandwidth, of arc (i, j), $b_{ij} \in R$

c_{ij} the administrative cost associated with arc (i, j), $c_{ij} \in R$; typically the administrative cost represents a measure of delay or transmission time

n the number of over-subscription factors

β_f the value of the f^{th} over-subscription factor, which is a multiplier of link capacities, where $f = 1, 2, \ldots, n;$ $\beta_f \in R$, $\beta_f > 1$, and $\beta_j > \beta_i$ if $j > i$

φ_f the probability of dropping packets when traffic on a link exceeds link capacity up to the f^{th} over-subscription factor times the link capacity where $f = 1, 2, \ldots, n;$ $\varphi_f \in R$, and $\varphi_j > \varphi_i$ if $j > i$

K the set of commodities or OD demands to be routed

O_k the origin (source) node of commodity $k \in K$

D_k the destination node of commodity $k \in K$

d_k the demand of commodity $k \in K$ in units of bandwidth, $d_k \in R$

ℓ the maximum allowed delay (latency) that any commodity may incur while traversing the network from source to destination, $\ell \in R$

h the maximum allowed number of hops that any commodity may traverse from source to destination, $h \in Z$

μ a unit of revenue generated from delivering a unit of demand of any commodity, $\mu \in R$

ω_1 a scaling factor used in the objective function to scale the penalty, $\omega_1 \in R$

ω_2 a scaling factor used in the objective function to scale the total delay, $\omega_2 \in R$

4.2.2 MPLS Traffic Engineering with Over-subscription Model

Using the above notation, the problem of MPLS traffic engineering with over-subscription of link capacities can be stated as follows. A graph $G = (N, A, \mathbf{b}, \mathbf{c})$ describes the physical topology and link attributes of an MPLS network. Also given is the node-to-node traffic represented by the set of commodities K. It is assumed that the demand d_k of a commodity $k \in K$ is expressed as a peak traffic rate (but the actual demand fluctuates). Considering the statistical characteristics of the traffic, in order to take advantage of statistical multiplexing, link capacities are over-subscribed so that total allowed flow on a link may exceed link capacity. In this formulation, over-subscription is generalized to include multiple factors instead of a single one. The multiple factors define multiple ranges of excess flow on a link. Excess traffic within an over-subscription range is subject to being dropped with some probability. The over-subscription factors β_f and associated drop probabilities φ_f are given. It is also assumed that a penalty proportional to the amount of dropped traffic is assessed when traffic exceeds link capacity. The objectives are to maximize revenue by admitting and routing as much traffic as possible, to minimize total dropped traffic (with the associated penalty), and to minimize total delay while observing the resource and traffic-performance constraints. The three objectives are represented as a linear combination with scale factors ω_1 and ω_2 (as detailed below). The resource constraint pertains to link capacity and is not strictly enforced. The combined traffic across all commodities on any link is allowed to exceed the capacity (b_{ij}) of that link. For that purpose, multiple slack variables are defined for each link that represent excess traffic in different ranges. The performance constraints are expressed in terms of maximum allowable delay and maximum allowable number of hops. Each commodity is to be routed along a path that does not exceed the maximum delay and number of hops. The following decision variables are defined:

$$y^k = \begin{cases} 1, & \text{if commodity } k \text{ is routed,} \\ 0, & \text{otherwise,} \end{cases}$$

$$x_{ij}^k = \begin{cases} 1, & \text{if commodity } k \text{ is routed through a path that uses arc } (i,\ j), \\ 0, & \text{otherwise,} \end{cases}$$

and z_{ij}^f are the slack variables representing the total traffic flow on arc (i, j) that exceeds link capacity within the range defined by the f^{th} and $(f\text{-}1)^{\text{st}}$ over-subscription factors as specified by constraint (4.4) below. The over-subscription problem can be formulated as follows:

[TE-OS]

$$\text{Maximize} \sum_{k \in K} \mu d_k y^k - \omega_1 \sum_{f=1}^{n} \varphi_f \sum_{(i,j) \in A} z_{ij}^f - \omega_2 \sum_{(i,j) \in A} c_{ij} \sum_{k \in K} d_k x_{ij}^k \tag{4.1}$$

subject to,

$$\sum_{j \in N \mid (i,j) \in A} x_{ij}^k - \sum_{j \in N \mid (i,j) \in A} x_{ji}^k = \begin{Bmatrix} y^k, & if\ i = O_k \\ -y^k, & if\ i = D_k \\ 0\ , & otherwise \end{Bmatrix} \quad \forall k \in K, \forall i \in N, \tag{4.2}$$

$$\sum_{k \in K} d_k x_{ij}^k \le b_{ij} + \sum_{f=1}^{n} z_{ij}^f \qquad \forall (i,j) \in A, \tag{4.3}$$

$$z_{ij}^f \le \left(\beta_f - \beta_{f-1}\right) b_{ij} \quad \forall (i,j) \in A,\ f = 1,\dots,n \text{ and } \beta_0 = 1 \tag{4.4}$$

$$\sum_{(i,j) \in A} c_{ij} x_{ij}^k \le \ell \qquad \forall k \in K, \tag{4.5}$$

$$\sum_{(i,j) \in A} x_{ij}^k \le h \qquad \forall k \in K, \tag{4.6}$$

$$x_{ij}^k, y^k \in \{0,1\} \qquad \forall (i,j) \in A, \forall k \in K, \tag{4.7}$$

$$z_{ij}^f \in R \quad \forall (i,j) \in A,\ f = 1,\dots,n. \tag{4.8}$$

In this formulation, the objective function consists of three terms with the first one being the primary objective and the dominant term. The first term represents the total revenue generated from routed commodities (i.e., delivered traffic). If a commodity is routed, its entire demand is delivered. The total demand delivered, and not necessarily the total number of commodities delivered, is maximized. The second term represents the penalty due to over-subscription and possible violation of packet-delivery performance guarantees. The penalty is assessed by estimating the amount of dropped traffic and scaling it by factor ω_1 to express it as a financial penalty. In the computational experiments, ω_1 is set to 1. However, it can be set to a greater value if the amount of dropped traffic and the associated penalty are to be reduced. The third term represents the total delay incurred by all the delivered commodities. The purpose of this term is to select the solution with the lowest delay among multiple alternate optimum solutions (yielding the same revenue) that may exist. The total delay is multiplied by the scaling factor ω_2. Typically, $0 < \omega_2 << 1$ so that the value of the third term is less than one and the first and second terms remain dominant.

Constraints (4.2) are the flow-conservation equations, which ensure a connected path for each routed commodity. For each commodity, there are $|N|$ equations: one equation for the source node of that commodity, one for the destination node, and $|N|$ - 2 equations for the other nodes in the network, which act as transit nodes. Flow for commodity $k \in K$ is routed only if binary variable $y^k = 1$; otherwise that particular demand is not satisfied. The unit supplies and demands and binary flow variables cause a single LSP to be formed via these constraints.

Constraints (4.3) enforce the arc capacity constraints. For each arc, the total traffic from all commodities whose paths include that arc is allowed to exceed the arc's bandwidth. Multiple slack variables are used to represent excess traffic on arcs. Constraint set (4.4) sets the upper bound for each slack variable for each arc, which defines multiple ranges of traffic overflow on an arc. Constraints sets (4.5) and (4.6) represent the performance requirements. Constraint set (4.5) ensures that the delay along any path cannot exceed a predetermined upper delay limit ℓ and (4.6) ensures that the number of hops along any path cannot exceed a predetermined upper hop limit h. Constraint set (4.7) constrains the decision variable x_{ij}^k and y^k to be binary, and constraint set (4.8) defines the variables z_{ij}^f as non-negative real numbers.

4.3 Computational Experiments

Computational experiments are performed to demonstrate the effectiveness of the TE-OS model for capacity planning and for estimating the amount of traffic dropped as demand grows. The experiments illustrate the robustness of the model across different loads and changes in over-subscription parameters.

4.3.1 Organization of Tests

4.3.1.1 Test Network Characteristics

The tests are performed on the realistic network N0 depicted in Figure 2-1. N0 represents a typical topology of nationwide data communications network implemented by inter-exchange providers in the US. N0 consists of 20 nodes and 62 arcs with an average node degree of approximately 3. The trunks connecting the nodes are bi-directional and full duplex, so each trunk is represented as two directed arcs with the same capacity and cost. The thick trunks represent OC-48 transmission lines with 2488 Mbps of bandwidth capacity, and the thin trunks represent OC-12 transmission lines with 622 Mbps of bandwidth capacity. The arc cost represents the actual circuit mileage of the corresponding transmission line. Since the shortest-path algorithm uses the arc cost as the metric to calculate the shortest path the OC-12 cost is increased by 1000 units so that an OC-48 trunk will be preferable to an OC-12 trunk.

4.3.1.2 Traffic Generator

A traffic generating computer program was developed to create multiple sets of commodities for the different experiments. The generator accepts as input parameters: the number of nodes in a network, the number of OD pairs, and the minimum and maximum OD demands. The generator randomly selects: OD pairs (and without replacement) from the set of nodes and the demands associated with those pairs (using a uniform distribution over the range specified by the minimum and maximum demands). The minimum and maximum demands are selected based on the test scenario. For example, when testing with N0 the maximum demand is always less than the 622 Mbps capacity of the lower bandwidth links.

4.3.1.3 Computing Environment

All test cases are performed on a Compaq AlphaServer DS20E with dual 667 MHz processors and 4096 MB RAM. The machine is configured as a Network Queuing System and executes batch jobs. Each job on the system has access to approximately 2 GB RAM. The models are implemented using the GAMS [14] model description language, and integer-programming solutions are generated using CPLEX Linear Optimizer 7.0. Default settings for CPLEX are used with the exception that the MIP time limit is set to 3600 seconds and the relative optimality gap is set to 1%.

4.3.2 Experiment Design

The experiments are designed for capacity planning purposes and for estimating the amount of traffic dropped as demand grows. The experiments are intended to show the relationships between revenue, penalties, and the number of links within each traffic overflow range and increases in demand. These relationships can be used to determine the percent of traffic growth that the network can accommodate with existing capacity, and identify the links with the highest levels of traffic overflow. Those links would be the first candidates for capacity augmentation. The model and the experiments do not determine the amount of increase in bandwidth on the candidate links.

A series of experiments is conducted to establish these relationships. First, an initial (base) traffic matrix that consists of a full mesh of commodities (i.e., 380 OD pairs) is generated. The demands are uniformly distributed with minimum, average, and maximum demands per OD pair of 10, 50, and 90 Mbps, respectively. The traffic load is chosen so that network utilization is high with minimal or no traffic overflow[1]. The initial matrix is used as the baseline for generating subsequent traffic matrices that represent demand growth. Ten additional traffic matrices are generated by increasing the traffic in increments of 10% up to 100% of the initial demands. The model TE-OS is solved for each traffic matrix and the results are tabulated. To investigate the impact of traffic variations on link overflow trends, the above series of experiments is repeated with different sets of demand parameters for the initial traffic matrix. Table 4-1 shows experiment sets A-G and associated parameter values.

When solving the model, the values of the over-subscription parameters are pre-specified. The following values are selected for the first experiment set: $n = 4$; $\beta_1 = 1.1$, $\beta_2 = 1.2$, $\beta_3 = 1.4$, $\beta_4 = 100$; and $\varphi_1 = 0.1$, $\varphi_2 = 0.2$, $\varphi_3 = 0.4$, $\varphi_4 = 0.9$. Although somewhat arbitrary, setting n to 4 provides adequate partitioning of traffic overflow into ranges and grouping of links into these ranges. If additional granularity is needed, a larger n may be chosen. It is assumed that for the first three ranges of traffic overflow the drop probability is proportional to the size of the range, and for the fourth range the drop probability is significantly higher to simulate a finite buffer. This is consistent with buffering and drop strategies implemented on actual equipment. Also, β_4 is set to a high value to allow unlimited over-subscription so that any demand will be accepted (with its

[1] Choosing lower load levels just shifts the trends and does not contribute to the study (in practice, the initial traffic matrix would be the current traffic matrix).

associated penalty). The fourth range will be referred to as the *critical range* and links that overflow into that range will be considered to be severely congested and in need of bandwidth augmentation.

The over-subscription parameters approximate the statistical characteristics of packet traffic and the behavior of network queuing[2]. The impact of the setting of these values on the link-overflow trends is explored to some extent. The above values are changed in other experiment sets listed in Table 4-1 and the model is solved for the same sets of traffic matrices and the results are compared.

Table 4-1 Experiment sets parameter values

Experiment Set	OD Pairs	OD Demand (Mbps)			Over-subscription Parameters (n = 4)							
		Min	Avg	Max	β_1	β_2	β_3	β_4	φ_1	φ_2	φ_3	φ_4
A	380	10	50	90	1.1	1.2	1.4	100	0.1	0.2	0.4	0.9
B	380	20	50	80	1.1	1.2	1.4	100	0.1	0.2	0.4	0.9
C	380	40	50	60	1.1	1.2	1.4	100	0.1	0.2	0.4	0.9
D	190	20	100	180	1.1	1.2	1.4	100	0.1	0.2	0.4	0.9
E	380	10	50	90	1.1	1.2	1.3	100	0.1	0.2	0.3	0.9
F	380	10	50	90	1.1	1.2	1.5	100	0.1	0.2	0.5	0.9
G	380	10	50	90	1.1	1.2	1.4	100	0.4	0.6	0.8	0.9

4.3.3 Experiment Results

The experimental results are presented in the following tables and figures. Each table contains the results for a set of experiments with eleven traffic matrices. The base traffic matrix and the values of the over-subscription parameters define experiment sets A-G (Table 4-1). Tables 4-2 through 4-7 are organized as follows. The first column indicates the percentage increase in traffic from the base traffic matrix. The value in the first row is set to zero to indicate the base traffic matrix. The second column contains the

[2] A more exact algorithm for determining these values is left for future study.

percent of OD pairs that are rejected. The third column contains the total net revenue, which is the value of the objective function in equation (4.1) with μ =1, ω_1 = 1, and ω_2 = 10^{-8}. The fourth column calculates the percent increase in net revenue from the revenue in the first row. The fifth column contains the penalty as defined by the second term in equation (4.1). The sixth column calculates the ratio of penalty to gross revenue; the ratio is calculated by dividing the second term by the first term in equation (4.1). The seventh column contains the number of links with no overflow (i.e., links with utilization less than 100%). The last four columns show the number of links with overflow traffic in the four ranges specified by the β_f parameter values.

Table 4-2 shows the results of solving the TE-OS model for experiment set A. Figure 4-1 shows a histogram of the percent of links in the different overflow ranges. The results show that:

1. all OD pairs are accommodated when demand increase are 50% or less.
2. net revenue consistently increases with increased demand.
3. over-subscription occurs on a small portion of the 62 links with demand growth of 20% or less, but occurs on the majority of arcs when growth is 50% or more.
4. there are no links in the critical range with traffic increments up to 40%. Link (5, 13) is the first link to overflow into the critical range when the demand increases by 50%. Link (3, 18) is the next link to overflow into the critical region when the demand increases by 60%.

Table 4-2 Solution results of the TE-OS model for experiment set A

Demand Growth %	OD Pairs Rejected %	Net Revenue $	Net Rev. Inc. %	Penalty $	Penalty to Rev. %	Links with no Overflow	Links with overflow in the Range (100%-110%]	(110% - 120%]	(120% - 140%]	> 140%
0	0	18,227	0	11	0.1	57	5	0	0	0
10	0	20,050	10.0	13	0.1	56	6	0	0	0
20	0	21,857	19.9	30	0.1	53	9	0	0	0
30	0	23,528	29.1	183	0.8	40	18	4	0	0
40	0	24,974	37.0	561	2.2	33	17	8	4	0
50	0	26,117	43.3	1,241	4.5	26	14	11	10	1
60	2.4	27,105	48.7	1,381	4.8	25	15	11	9	2
70	4.2	27,969	53.4	1,697	5.7	19	19	11	11	2
80	7.6	28,667	57.3	2,057	6.7	15	17	10	17	3
90	9.5	29,699	62.9	1,721	5.5	23	15	12	9	3
100	11.6	30,419	66.9	1,784	5.5	19	19	9	12	3

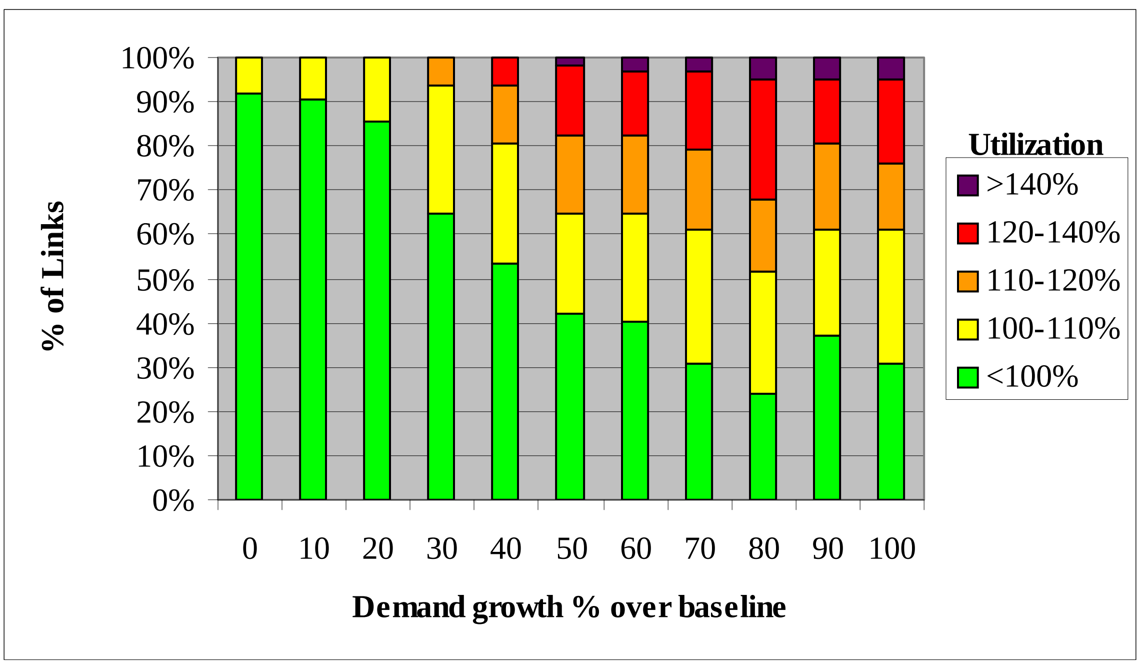

Figure 4-1 Link overflow histogram for experiment set A

Table 4-3 Solution results of the TE-OS model for experiment set B

Demand Growth %	OD Pairs Rejected %	Net Revenue $	Net Rev. Inc. %	Penalty $	Penalty to Rev. %	Links with no Overflow	Links with overflow in the Range (100%-110%]			
							(100%-110%]	(110% - 120%]	(120% - 140%]	> 140%
0	0	18,419	0	9	0	57	5	0	0	0
10	0	20,261	10.0	10	0	54	8	0	0	0
20	0	22,072	19.8	42	0.2	50	12	0	0	0
30	0	23,735	28.9	221	0.9	42	17	3	0	0
40	0	25,173	36.7	625	2.4	29	19	11	3	0
50	0.3	26,309	42.8	1,244	4.5	27	12	12	11	0
60	2.9	27,298	48.2	1,393	4.9	27	13	12	9	1
70	5.0	28,201	53.1	1,529	5.1	22	15	15	8	2
80	8.2	29,043	57.7	1,627	5.3	22	15	13	9	3
90	9.5	29,786	61.7	2,078	6.5	23	11	11	12	5
100	11.6	30,627	66.3	2,017	6.2	24	11	10	14	3

Table 4-4 Solution results of the TE-OS model for experiment set C

Demand Growth %	OD Pairs Rejected %	Net Revenue $	Net Rev. Inc. %	Penalty $	Penalty to Rev. %	Links with no Overflow	Links with overflow in the Range			
							(100%-110%]	(110% - 120%]	(120% - 140%]	> 140%
0	0	18,795	0	9	0	55	7	0	0	0
10	0	20,675	10.0	9	0	54	8	0	0	0
20	0	22,496	19.7	68	0.3	53	9	0	0	0
30	0	24,162	28.6	282	1.2	42	15	4	1	0
40	0	25,533	35.9	792	3.0	23	24	11	4	0
50	1.1	26,655	41.8	1,256	4.5	27	12	14	7	2
60	2.9	27,569	46.7	1,652	5.7	23	13	14	11	1
70	5.5	28,472	51.5	1,744	5.8	23	12	14	11	2
80	7.6	29,287	55.8	1,973	6.3	18	21	9	10	4
90	9.5	30,039	59.8	2,280	7.1	18	17	9	13	5
100	12.9	30,979	64.8	1,758	5.4	19	19	13	7	4

Tables 4-3 and 4-4 include the results of solving the TE-OS model for experiment sets B and C, respectively. The results of experiment sets A, B, and C are compared to examine the impact of demand variance on the different trends. These sets share the same over-subscription parameters and the same average demand value but differ in the minimum and maximum demand values. Figure 4-2 compares the penalty-to-revenue ratios for the three experiment sets. The graphs are plotted for traffic increases of up to 50%. (Operating at higher traffic levels results in rejecting some of the demands and the ratios are no longer valid for comparison.) The graphs for the three experiment sets are similar and indicate that the impact of traffic variations on the ratios is minimal.

Experiment set B identifies link (13, 5) as the first link to overflow into the critical range when traffic increases by 60%. Experiment set C identifies the same link when traffic increases by 50%. Both sets identify links (5, 13) and (3, 18) as the subsequent links to overflow into the critical range.

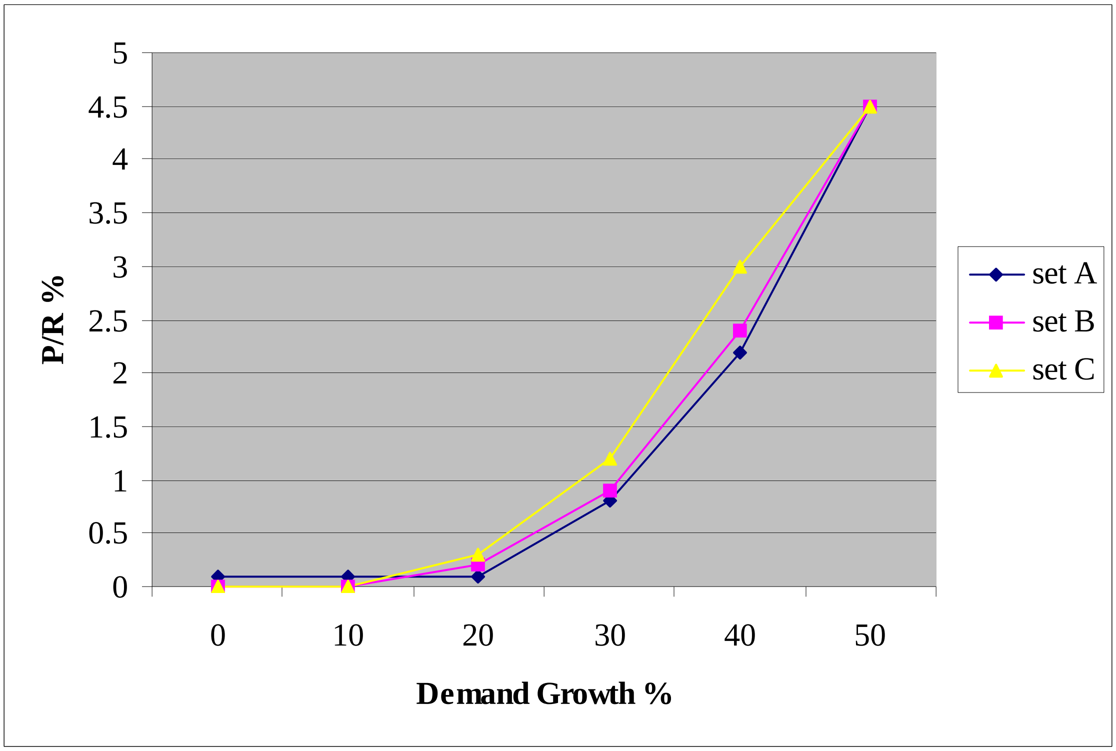

Figure 4-2 Comparison of penalty to revenue ratios for experiment sets A, B, and C

Table 4-3 shows the results of solving the TE-OS model for experiment set D. The results show that some demand is rejected when the traffic increases by 40% or more, and that there are no links in the critical range with traffic increments up to 50%. The links (3, 18), (7, 18), and (9, 20) are the first links to overflow into the critical range when the demand increases by 50%. The links (5, 13) and (13, 5) are the most-utilized backbone links but are operating in the third range. Figure 4-3 shows a histogram of the percent of links in the different overflow ranges.

Table 4-5 Solution results of the TE-OS model for experiment set D

Demand Growth %	OD Pairs Rejected %	Net Revenue $	Net Rev. Inc. %	Penalty $	Penalty to Rev. %	Links with no Overflow	Links with overflow in the Range (100%-110%]	(110% - 120%]	(120% - 140%]	> 140%
0	0	18,554	0	27	0.1	54	8	0	0	0
10	0	20,378	9.8	61	0.3	50	9	3	0	0
20	0	22,118	19.2	179	0.8	44	13	3	2	0
30	0	23,604	27.2	551	2.3	34	12	11	5	0
40	2.1	24,792	33.6	769	3.0	27	16	12	7	0
50	2.6	25,744	38.8	1,587	5.8	25	10	12	12	3
60	5.3	26,703	43.9	1,535	5.4	21	12	11	17	1
70	7.4	27,700	49.3	1,642	5.6	23	9	16	13	1
80	8.4	28,385	53.0	2,123	7.0	20	10	13	15	4
90	13.2	29,115	56.9	1,824	5.9	20	11	9	21	1
100	13.7	29,772	60.5	2,141	6.7	19	8	14	19	2

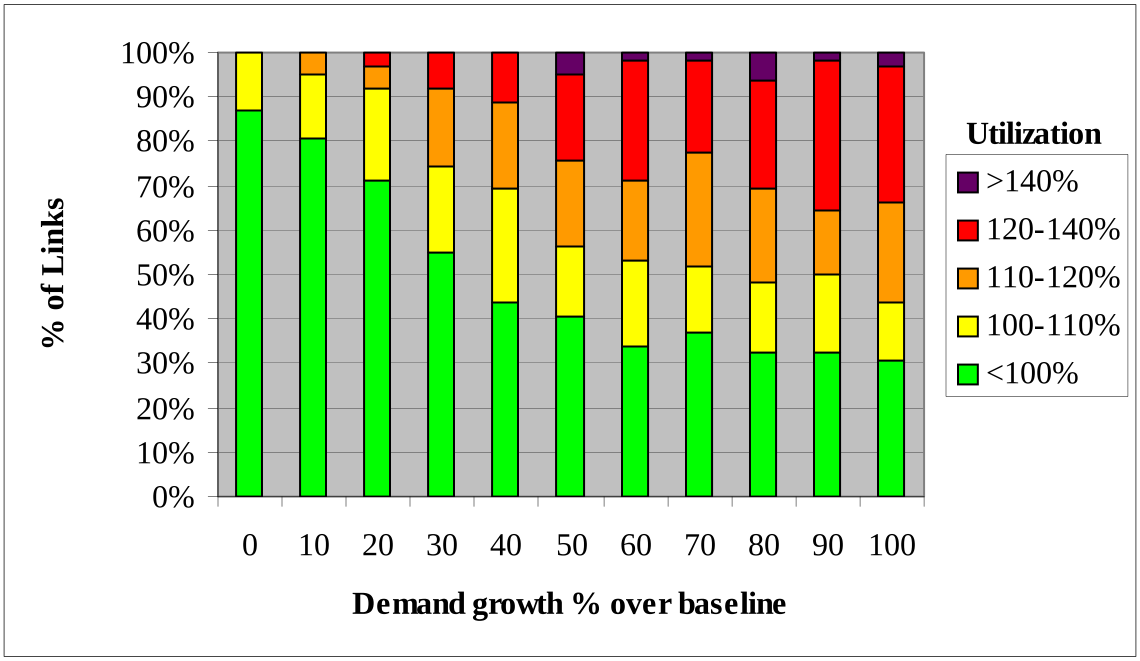

Figure 4-3 Link overflow histogram for experiment set D

Tables 4-6, 4-7, and 4-8 include the results of solving the TE-OS model for experiment sets E, F and G, respectively. These results are compared with the result of experiment set A to examine the impact of the over-subscription parameter values on the different trends. As documented in Table 4-1, these sets share the same base traffic matrix but differ in over-subscription parameter values. In experiment sets E and F the critical range is changed but the other parameters remain intact. In experiment set E, the critical range is expanded by setting β_3 to 1.3. In experiment set F, the critical range is reduced by setting β_3 to 1.5. Figure 4-4 compares the penalty-to-revenue ratios for experiment sets A, E, and F, with values plotted for traffic increments of up to 50%. (Operating at higher traffic increments results in rejecting some of the demands and the ratios are no longer valid for comparison.) Figure 4-4 shows that the values are roughly equivalent with demand increments of up to 40%. However, when demand increases by 50%, experiment set F shows a lower ratio. Reducing the critical range results in having the first link overflow into that range only when the traffic increases by 70%. Experiment set F identifies link (5, 13) as the first link to overflow into the critical range. In both experiment sets A and E there are no over-subscribed links in the critical range when traffic increases by 50% or less. However, experiment set A identifies link (5, 13) as the first critical-range link and experiment set E identifies links (4, 7) and (3, 18) as critical-range links. With experiment set E, link (5,13) overflows into the critical range when traffic increases by 60%.

Table 4-6 Solution results of the TE-OS model for experiment set E

Demand Growth %	OD Pairs Rejected %	Net Revenue $	Net Rev. Inc. %	Penalty $	Penalty to Rev. %	Links with no Overflow	Links with overflow in the Range			
							(100%-110%]	(110% - 120%]	(120% - 140%]	> 140%
0	0	18,227	0	11	0.1	57	5	0	0	0
10	0	20,050	10.0	13	0.1	56	6	0	0	0
20	0	21,857	19.9	30	0.1	53	9	0	0	0
30	0	23,528	29.1	183	0.8	40	18	4	0	0
40	0	24,985	37.1	549	2.2	31	19	10	2	0
50	0.5	25,957	42.4	1,236	4.5	29	11	10	10	2
60	2.6	27,110	48.7	1,263	4.5	28	11	11	7	5
70	5.8	27,905	53.1	1,405	4.8	29	10	10	8	5
80	7.4	28,743	57.7	1,565	5.2	27	12	7	10	6
90	9.7	29,623	62.5	1,399	4.5	25	14	12	7	4
100	11.3	30,377	66.7	1,770	5.5	24	15	4	13	6

Table 4-7 Solution results of the TE-OS model for experiment set F

Demand Growth %	OD Pairs Rejected %	Net Revenue $	Net Rev. Inc. %	Penalty $	Penalty to Rev. %	Links with no Overflow	Links with overflow in the Range (100%-110%]	(110% - 120%]	(120% - 140%]	> 140%
0	0	18,227	0	11	0.1	57	5	0	0	0
10	0	20,050	10.0	13	0.1	56	6	0	0	0
20	0	21,857	19.9	30	0.1	53	9	0	0	0
30	0	23,528	29.1	183	0.8	40	18	4	0	0
40	0	24,958	36.9	576	2.3	31	18	10	3	0
50	0.8	26,070	43.0	1,057	3.9	27	15	17	3	0
60	2.4	27,032	48.3	1,504	5.3	22	18	12	10	0
70	3.9	27,725	52.1	2,093	7	15	20	14	12	1
80	6.1	28,750	57.7	1,907	6.2	23	13	13	13	0
90	8.2	29,390	61.2	2,473	7.8	18	11	14	17	2
100	11.1	30,214	65.8	2,235	6.9	14	19	12	15	2

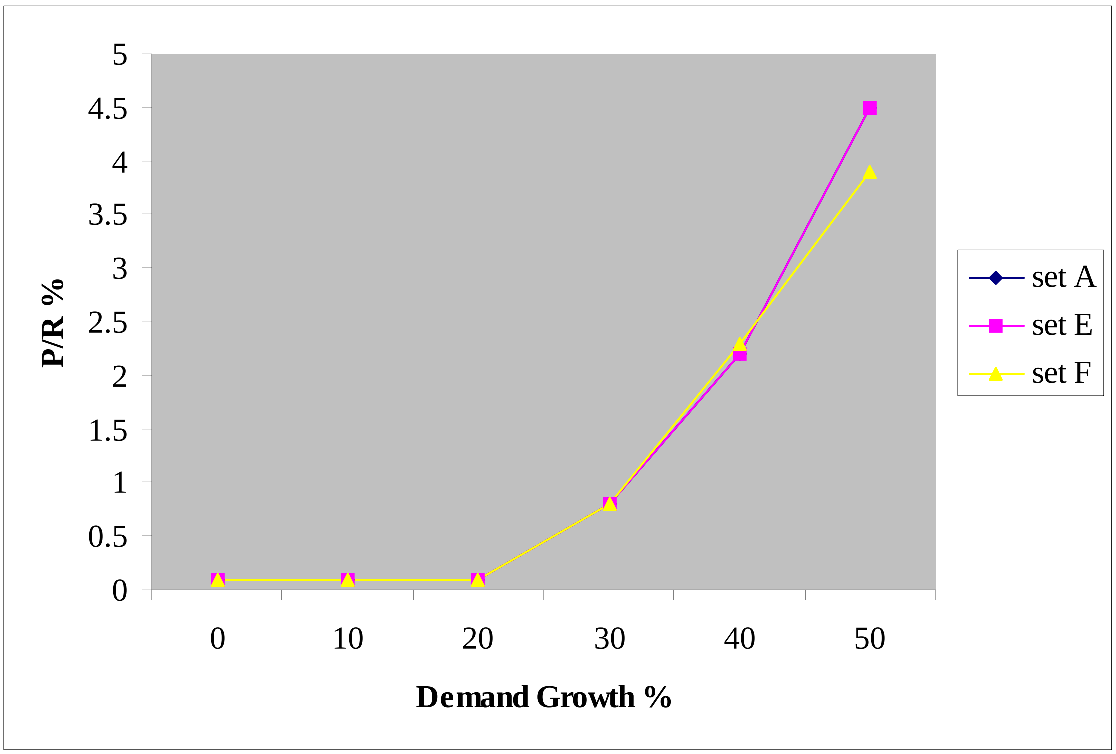

Figure 4-4 Comparison of penalty to revenue ratios for experiment sets A, E, and F[3]

[3] The graphs of experiment sets A and E overlap.

Table 4-8 shows the results of solving the TE-OS model for experiment set G. This set uses the same base traffic matrix and β_f values, but has significantly higher drop probabilities (φ_f) than experiment set A. The results show general trends similar to those in Table 4-2, but with some significant differences. As expected, Table 4-8 shows higher penalties and penalty-to-revenue ratio for traffic increments up to 40% (at higher increments the percent of demand is higher so the penalty is lower). In experiment set G, some demand is rejected when the traffic increases by 30% or more, whereas in set A (Table 4-2), similar rejections do not occur until 60% is reached. Also, Table 4-8 shows that the first critical-range overflows occur when demand increases by 100%, whereas Table 4-2 shows that only a 50% increase is required to trigger this in set A. Experiment set G identifies (3, 18) as the first critical-range link, whereas in experiment set A, link (5, 13) is first link and link (3, 18) is next.

Table 4-8 Solution results of the TE-OS model for experiment set G

Demand Growth %	OD Pairs Rejected %	Net Revenue $	Net Rev. Inc. %	Penalty $	Penalty to Rev. %	Links with no Overflow	Links with overflow in the Range (100%-110%]	(110% - 120%]	(120% - 140%]	> 140%
0	0	18,193	0	45	0.2	57	5	0	0	0
10	0	20,011	10.0	52	0.3	56	6	0	0	0
20	0	21,768	19.6	118	0.5	53	9	0	0	0
30	0.3	23,016	26.5	600	2.5	44	15	3	0	0
40	2.9	23,965	31.7	758	3.1	41	17	3	1	0
50	5.0	24,751	36.0	1,086	4.2	37	20	4	1	0
60	7.9	25,622	40.8	936	3.5	39	15	6	2	0
70	11.6	26,458	45.4	983	3.6	35	22	2	3	0
80	13.2	27,157	49.3	1,189	4.2	35	18	8	1	0
90	15.5	27,994	53.9	1,087	3.7	32	21	7	2	0
100	16.3	28,650	57.5	1,679	5.5	31	20	7	3	1

4.3.4 Analysis of Results

The solutions of the TE-OS model for the different experiment sets provide useful information for two major capacity-planning decisions. First, they yield an estimate of the percent of traffic growth that the network can accommodate with existing capacity (hereafter referred to as the *growth potential*) and, second, they identify the links that require capacity augmentation. To determine the growth potential, some criteria with metrics need to be defined and then applied to the solution results. The following are some proposed criteria. The growth potential is the highest percent increase in demand that meets one or some combination of the following criteria:

1. No demand is rejected.
2. No link has traffic overflow in the critical range.
3. The penalty-to-revenue ratio is less than a given predetermined value (e.g., 1%).
4. The portion of links with no traffic overflow is at least some predetermined value (e.g., 50%).
5. The net revenue relative increase is no less than some predetermined percent (e.g., 3%) from the corresponding demand increase.

Applying each criterion separately to experiment set A's solution results (Table 4-2) yields the following growth potential for the example network studied here. Criteria 1 to 5 estimate the growth potential as 50%, 40%, 30%, 40%, and 40%, respectively. Based on the most conservative measure in this case, the growth potential is estimated as 30%. Comparing the solution results of experiment set A with those of set G shows that the values of the over-subscription parameters significantly affect the results with respect to growth potential. Applying each criterion separately to experiment set G's solution results (Table 4-4) yields the following growth potential. Criteria 1 to 5 estimate the growth potential as 20%, 90%, 20%, 100%, and 20%, respectively. Based on the most conservative measure in this case, the growth potential is estimated as 20%. Considering the most conservative measure of both experiment sets, the growth potential is estimated as 20%.

The model can also assist management by identifying links that require capacity augmentation. Defining the critical range and using it as a criterion for identifying the severely congested links proved to be effective and robust. Despite the differences in demand and over-subscription parameter values, the experiment sets identified the same set of congested links, but not in the same order. Considering the results of all experiment sets, backbone link (5, 13) and access link (3, 18) are the most-congested links and should be considered first for capacity augmentation.

The TE-OS model also provides additional insight into bandwidth-utilization differences between full-mesh and half-mesh demand matrices. Comparing the solution results of experiment sets A (full mesh) and D (half mesh) shows that the network is more efficient in its handling of the full mesh of demands. With experiment set A, less demand is rejected, penalty-to-revenue ratios are lower, relative net-revenue increase is higher, and traffic overflow on links is lower. This behavior is attributed to the fact that, in the half-mesh case, the higher $\bar{d}$ tends to congest the lower-speed access links. A detailed examination of the solution results reveals that access links overflow into the critical range before the backbone links do. These results are consistent with the results of Chapter 2, which shows that for the same total demand, the traffic delivery ratio is higher with a larger number of OD pairs with correspondingly smaller individual demand values.

4.4 Summary and Conclusions

A new formulation of the problem of MPLS traffic engineering with over-subscription is presented. The problem is formulated as an OD integer multi-commodity network flow problem with side constraints. The model maximizes revenue and minimizes total dropped traffic and associated penalty. The computational experiments demonstrate the usefulness of the model for capacity-planning purposes. The results show that the model can be of use in two major capacity planning decisions. It helps estimate network growth potential and identify the links that require capacity augmentation. Estimating the growth potential requires defining some criteria to be applied to the model solution results; proposed criteria are defined and recommended for use by network administrators.

The results also show that estimating growth potential depends on the values of the over-subscription parameters (which are also estimates). Therefore, it is recommended that the model be run with different sets of parameter values to obtain multiple estimates. With a conservative approach, the lowest estimate may be selected and used for planning purposes. The results also show that identifying the severely congested links requiring capacity augmentation is more robust and less dependent on the over-subscription parameter values. However, it is still recommended that the model be run with different sets of parameter values to ensure more reliable identification.

Chapter 5

SUMMARY AND CONCLUSIONS

Multi-Protocol Label Switching (MPLS) is an evolving switching technology that is being integrated into Internet Protocol (IP) networks to overcome IP-routing deficiencies. MPLS facilitates traffic engineering (TE) by providing the mechanisms needed to control traffic flows in IP networks. Combined with differentiated services (Diffserv) capabilities, MPLS enables the implementation and support of multiple classes-of-service (CoS) types, each with specific quality-of-service (QoS) guarantees. Thus, MPLS facilitates network optimization to maximize resource utilization and enables the convergence of data, voice, and video applications over a common network infrastructure.

5.1 Contributions

This praxis contributes new models for multiple fundamental problems related to MBLS-based TE in IP networks supporting single or multiple CoS types. The models focus on revenue maximization, which is one of the primary goals of MPLS deployment by service providers. The praxis also contributes new methodologies for evaluating the benefits and effectiveness of different strategies and capabilities of MPLS-based TE. Comprehensive computational studies are conducted on realistic networks and provide a greater insight into design considerations and factors influencing the effectiveness of the models. The models have practical applications and can be used by network administrators and managers in the TE design process.

Chapter 2 presents an optimization model for the basic TE problem of constraint-based routing and admission control, which involves the design of label switched paths (LSPs) to route traffic efficiently and to maximize revenue. The model assumes a single CoS type and that demand between each origin-destination (OD) pair is routed along a single, unique path. A computational study compares the performance of an offline strategy utilizing the optimization model with an online strategy, which implements a first-come-first-served (FCFS) algorithm. The study also assesses the impact of different factors on the performance of the two strategies.

The study demonstrates an improvement of up to 13.77% in revenue increase by the offline strategy under some scenarios. Generally, optimization provided improvement as the load increased and the average link utilization approached or exceeded 60%. Both strategies perform better with a larger number of OD pairs, assuming that the average demand per OD pair is lower. This is a favorable result for actual implementation in real networks, where demand is likely to exist between any pair of nodes and a full mesh of paths is required. The two design factors, network topology and node degree have a significant impact on performance. The results demonstrate that the doubling of the node degree, while reducing the link bandwidth by half (so that the total network capacity is the same), results in a more efficient network. This conclusion needs to be seriously considered and explored during actual network design processes.

Chapter 2 also extends the basic model for load balancing using multiple paths per OD pair. A computational study evaluates the benefits of load balancing (with different split-ratios) and assesses the affects of different factors on performance improvement. The results show that load balancing provides low to moderate improvement. An improvement of up to 11.75% is achieved on the realistic network under some scenarios. While the characterization of the impact of the different factors on performance is not conclusive, some general observations are made. The improvement depends on the demand per OD pair but is not monotonic as the demand increases. The improvement depends on the ratio of link capacity to the demand per OD pair (i.e., the number of OD pairs that can be packed within a link's bandwidth). The larger the number of OD pairs (and assuming lower demand per OD pair), the less significant the improvement. The results also show that neither split-ratio tested dominates the other for all demands. Finally, network topology does not seem to affect the general behavior but can impact the magnitude of improvement itself.

Chapter 3 enhances the basic TE model to deal with multiple CoS types. Each CoS type is assigned a priority, given its own QoS performance requirements, and service for each type is priced differently. A computational study for the case of two CoS types (A and B) compares the performance of the two-class model with the single-class model for a realistic network. Subsequently, a breakeven revenue analysis is conducted to help determine the value of the relative increase in revenue per unit of demand for the higher-class traffic (class A).

The results show that both models achieve higher traffic delivery ratios with larger numbers of OD pairs but smaller demand sizes. The results also demonstrate that delivering the higher-class traffic may reduce the total traffic delivered and increase bandwidth utilization depending on the size of demand per OD pair. To compensate for this in resource utilization and protect against potential loss in revenue, the revenue per unit of demand of class-A traffic would need to be increased. The relative increase depends on the number of OD pairs, the demand per OD pair, and the effect of load balancing. The results show that the relative increase can be as high as 18.62%.

The results also show that partitioning an aggregate demand into multiple demands for different classes and routing them separately indirectly realizes the benefits associated with load balancing and results in increasing the traffic delivery ratio. This result is of major practical significance since it suggests that service providers may exploit a revenue-generating service feature as a vehicle to increase network efficiency and promotes the adoption of Diffserv-aware traffic engineering (DS-TE).

Chapter 4 extends the basic TE model to deal with over-subscription of link capacities. Over-subscription allows to admit demand that otherwise would have been rejected but a penalty may be assessed on traffic that exceeds link capacities. Over-subscription exploits fluctuations in demand and provides a statistical-multiplexing gain. The model approximates a complex stochastic problem using a parametric deterministic approach.

The computational experiments demonstrate the usefulness of the model for capacity-planning purposes. The results show that the model can be used to make two major capacity-planning decisions. It helps estimate the network growth potential and identify the links that require capacity augmentation. Estimating the growth potential requires defining some criteria to be applied to the model solution results. Some such criteria are defined and recommended for use by network administrators. The results also show that estimating the growth potential depends on the values of the over-subscription parameters, which by themselves are also estimates. Therefore, it is recommended that the model be run with different sets of parameter values to obtain multiple estimates. With a conservative approach, the lowest estimate would be selected and used for planning purposes. The results also show that identifying the severely congested links that require capacity augmentation is more robust and less dependent on the over-subscription parameters values.

5.2 Future Research Topics

The models and computational results presented in this praxis can provide a basis for future research in several areas. The models can be extended and computational testing can be conducted to evaluate several enhancements. The models can be enhanced to model additional MPLS TE capabilities such as resource class affinity and resilience attributes [7], which can be associated with OD pairs and would require additional constraints on path selection. Resources (e.g., nodes and links) can be grouped into different classes, with the resource class affinity associated with an OD pair being used to exclude or include specified resource classes from the path for that pair. Incorporating resiliency and investigating different schemes for path, link, or node recovery [15, 16] would be of interest to the telecommunications industry. Specifically, the load-balancing model in Chapter 2 can be enhanced so that the two paths per OD pair are required to be link-disjoint and possibly node-disjoint. Another modification to the models involves enhancing the delay (and number of hops) constraints to have an upper bound specified per LSP instead of one common upper bound for all LSPs.

The other area of suggested future work is the validation or generalization of some of the important observations or conclusions made. The conclusion that the different models perform better with larger number of OD pairs but smaller demands needs to be further explored. Clearly the demand matrix is an input and not a control variable but the above conclusion suggests that splitting the demand through natural mechanisms, such as load balancing or multiple CoS types, should be exploited to increase network efficiency. Similarly, the observation that higher network efficiency can be achieved by doubling the number of links rather than doubling the capacity of the existing links can have major implications for network design and needs to be investigated further.

For that purpose the computational testing needs to be expanded and supporting analytical work conducted. The testing conducted used uniform or constant distributions of OD pair demands. More elaborate traffic models and distributions need to be investigated. Also, larger problem instances involving larger numbers of nodes, links, and OD pairs, and experiments with different types of network topologies should be investigated. The testing of multiple CoS types in Chapter 3 should be expanded to include more classes (e.g., four classes) with different proportions of the total demand per OD pair among the different classes.

Finally, the setting of the over-subscription parameter values in the model in Chapter 4, to better approximate the traffic fluctuations and queuing behavior in real networks, can be explored in greater detail. For that purpose, it is suggested that a simulation study or in-depth analysis of empirical data be conducted.

BIBLIOGRAPHY

[1] AHUJA, R. K., MAGNANTI, T. L., AND ORLIN, J. B. *Network Flows: Theory, Algorithms, and Applications*, Prentice-Hall, Upper Saddle River, NJ, 1993.

[2] AMIRI, A., ROLLAND, E., AND BARKHI, R. Bandwidth packing with queuing delay costs: Bounding and heuristic solution procedures. *European Journal of Operations Research 112* (1999), 635-645.

[3] AMIRI, A., AND BARKHI, R. The multi-hour bandwidth packing problem. *Computers & Operations Research 27* (2000), 1-14.

[4] AUKIA, P., KODIALAM, M., KOPPOL, P. V. N., LAKSHMAN, T. V., SARIN, H., AND SUTER, B. RATES: a server for MPLS traffic engineering. *IEEE Network* (2000), 34-41.

[5] AWDUCHE, D., BERGER, L., GAN, D., LI, T., SWALLOW, G., AND SRINIVASAN, V. RSVP-TE, extensions to RSVP for LSP tunnels. RFC 3209, Internet Engineering Task Force, http://www.ietf.org, 2001.

[6] AWDUCHE, D. O. MPLS and traffic engineering in IP networks. *IEEE Communications Magazine* 37:12 (1999), 42-47.

[7] AWDUCHE, D., MALCOLM, J., AGOGBUA, J., O'DELL, M., AND MCMANUS, J. Requirements for traffic engineering over MPLS. RFC 2702, Internet Engineering Task Force, http://www.ietf.org, 1999.

[8] AWDUCHE, D. O., AND JABBARI, B. Internet traffic engineering using multi-protocol label switching (MPLS). *Computer Networks 40* (2002), 111-129

[9] AYYANGAR, A., AND SIDHU, D. Analysis of MPLS based traffic engineering solution. *IEEE International Conference on ATM and High Speed Intelligent Internet Symposium* (2001), 21-27.

[10] BANERJEE, G., AND SIDHU, D. Comparative analysis of path computation techniques for MPLS traffic engineering. *Computer Networks 40* (2002), 149-165.

[11] BARR, R. S., GOLDEN, B. L., KELLY, J. P., RESENDE, M G. C., AND STEWART, W. R. Designing and reporting on computational experiments with heuristic methods. *Journal of Heuristics 1* (1995), 9-32.

[12] BERTSEKAS, D., AND GALLAGER, R. *Data Networks*, Prentice-Hall, Englewood Cliffs, NJ, 1987.

[13] BLAKE, S., BLACK, D., CARLSON, M., DAVIES, E., WANG, Z., AND WEISS, W. An architecture for differentiated services. RFC 2475, Internet Engineering Task Force, http://www.ietf.org, 1998.

[14] BROOKE, A., KENDRICK, D., MEERAUS, A., AND RAMAN, R. *GAMS: A User's Guide*, GAMS Development Corporation, Washington, DC, 1998.

[15] CHEN, T., KENNINGTON, J. L., AND OH, T. Fault restoration techniques for MPLS with QoS constraints. Tech. Rep. 00-CSE-7, Southern Methodist University, Dallas, TX, 2000.

[16] CHEN, T. M., AND OH, T. H. Reliable services in MPLS. *IEEE Communications Magazine* 37:12 (1999), 58-62.

[17] DAVIDSON, J., AND PETERS, J. *Voice over IP Fundamentals*, Cisco Press, Indianapolis, IN, 2000.

[18] DAVIE, B., CHARNEY, A., BENNET, J. C. R., BENSON, K., LE BOUDEC, J. Y., DAVARI, S., COURTNEY, W., FIRIOIU, V., AND D. STILIADIS, D. An expedited forwarding PHB (per-hop behavior). RFC 3246, Internet Engineering Task Force, http://www.ietf.org, 2002.

[19] ELWALID, A., JIN, C., LOW, S., AND WIDJAJA, I. MATE: MPLS adaptive traffic engineering. *IEEE INFOCOM* (2001), 1300-1309.

[20] FORTZ, B., REXFORD, J., AND THORUP, M. Traffic engineering with traditional IP routing protocols. *IEEE Communications Magazine 40:10* (2002), 118-124.

[21] FORTZ, B., AND THORUP, M. Internet traffic engineering by optimizing OSPF weights. *IEEE INFOCOM* (2000), 519-528.

[22] GIRISH, M. K., ZHOU, B., AND HU, J. Q. Formulation of the traffic engineering problems in MPLS based IP networks. In *Proceedings of the Fifth IEEE Symposium on Computers and Communications* (2000), 214-219.

[23] GUERIN, R., AHMADI, H., AND NAGHSHINEH, M. Equivalent capacity and its application to bandwidth allocation in high-speed networks. *IEEE Journal on Selected Areas in Communications 9* (1991), 968-981.

[24] GUICHARD, I., AND PEPELNJAK, I. *MPLS and VPN Architectures*, Cisco Press, Indianapolis, IN, 2001.

[25] GUICHARD, I., PEPELNJAK, I., AND APCAR, J. *MPLS and VPN Architectures,* Volume II, Cisco Press, Indianapolis, IN, 2003.

[26] HEINANEN, J., BAKER, F., WEISS, W., AND WROCLAWSKI, J. Assured forwarding PHB group. RFC 2597, Internet Engineering Task Force, http://www.ietf.org, 1999.

[27] HUNT, R. A review of quality of service mechanisms in IP-based networks — integrated and differentiated services, multi-layer switching, MPLS and traffic engineering. *Computer Communications 25* (2002), 100-108.

[28] KAR, K., KODIALAM, M., AND LAKSHMAN, T. V. Minimum interference routing of bandwidth guaranteed tunnels with MPLS traffic engineering applications. *IEEE Journal on Selected Areas in Communications 18* (2000), 2566-2579.

[29] KATZ, D., KOMPELLA, K., AND YEUNG, D. Traffic engineering (TE) extensions to OSPF version 2. RFC 3630, Internet Engineering Task Force, http://www.ietf.org, 2003.

[30] LE FAUCHEUR, F., WU, L., DAVIE, B., DAVARI, S., VAANANEN, P., KRISHNAN, R., CHEVAL, P., AND HEINANEN, J. Multi-protocol label switching (MPLS) support of differentiated services. RFC 3270, Internet Engineering Task Force, http://www.ietf.org, 2002.

[31] LE FAUCHEUR, F., AND LAI, W. Requirements for support of differentiated services-aware MPLS traffic engineering. RFC 3564, Internet Engineering Task Force, http://www.ietf.org, 2003.

[32] LI, T. MPLS and the evolving Internet architecture. *IEEE Communications Magazine* 37:12 (1999), 38-41.

[33] MCLOUD, T. W. *Constrained Integer Network Flows*. PhD thesis, Southern Methodist University, Dallas, TX, 2003.

[34] MITRA, D., AND RAMAKRISHNAN, K. G. A case study of multiservice, multipriority traffic engineering design for data networks. In *Proceedings of the IEEE Global Telecommunications Conference* (1999), 1077-1083.

[35] MOY, J. OSPF version 2. RFC 2328, Internet Engineering Task Force, http://www.ietf.org, 1998.

[36] OSBORNE, E., AND SIMHA, A. *Traffic Engineering with MPLS,* Cisco Press, Indianapolis, IN, 2003.

[37] POSTEL, J. DoD standard transmission control protocol. RFC 761, Internet Engineering Task Force, http://www.ietf.org, 1980.

[38] POSTEL, J. Internet protocol. RFC 791, Internet Engineering Task Force, http://www.ietf.org, 1981.

[39] ROSEN, E., VISWANASAN, A., AND CALLON, R., Multiprotocol Label Switching Architecture, RFC 3031, Internet Engineering Task Force, http://www.ietf.org, 2001.

[40] SAITO, H., MIYAO, Y., AND YOCHIDA, M. Traffic engineering using multiple multipoint-to-point LSPs. *IEEE INFOCOM* (2000), 894-901.

[41] SRIKITJA, A., AND TIPPER, D. Topological design of multiple VPNs over an MPLS network. In *Proceedings of the IEEE Global Telecommunications Conference* (2002), 2195-2199.

[42] SRIVASTAVA, S., KRITHIKAIVASAN, B., VENKATACHALAM, V., BEARD, C., MEDHI, D. VAN DE LIEFVOORT, A. ALANQAR, W., AND NAGARAJAN, A. A case study on evaluating the benefits of MPLS traffic engineering through constraint-based routing and network controls. In *Proceedings of the IEEE International Conference on Communications* (2002), 2437-2442.

[43] SWALLOW, G. MPLS advantages for traffic engineering. *IEEE Communications Magazine* (1999), 54-57.

[44] WALPOLE, R. E., MYERS, R. H., AND MYERS, S. L. *Probability and Statistics for Engineers and Scientists*, Prentice-Hall, Upper Saddle River, NJ, 1998

[45] XU, Y. X., AND ZHANG, G. D. Models and algorithms of QoS-based routing with MPLS traffic engineering. In *Proceedings of the Fifth IEEE International Conference on High Speed Networks and Multimedia Communications* (2002), 128-132.