ISTANBUL TECHNICAL UNIVERSITY ★ INSTITUTE OF SCIENCE AND TECHNOLOGY

TOPOLOGY AND BANDWIDTH ADAPTATION IN OPTICAL WDM BACKBONE NETWORKS WITH DYNAMIC TRAFFIC

Ph.D. Thesis by
Ayşegül GENÇATA, M.Sc.
(504952017)

Department : Control and Computer Engineering
Programme : Control and Computer Engineering

Supervisor : Prof. Dr. Bülent ÖRENCİK

APRIL 2003

PREFACE

Optical networks are widely accepted as one of the most promising technologies for the future of our knowledge society, and I feel very lucky to be a scientist working in this exciting field of research. I would like to thank to my advisor Prof. Dr. Bülent Örencik for encouraging me to start to work on WDM networks, and for his support throughout my study.

The research in this thesis has been supported by two research programmes: by Istanbul Technical University Grant for Long Term Research Activities from December 2000 to April 2001, and by US National Science Foundation Grant No. ANI-98-05285 from May 2001 to June 2002.

There is no doubt that my stay at University of California at Davis, between December 2000 and June 2002 had a big impact on this thesis, as well as on my whole career as a researcher and as an academician. I thank to Prof. Dr. Biswanath Mukherjee, for accepting an unknown student to his lab, for his time and all the effort he has put on this study, and for his advices that allowed me to exploit his expertise in the scientific world. I am honoured to be one of his students.

I thank UC Davis Optical Lab folks, Keyao Zhu, Shun Yao, Narendra Singhal, Laxman Sahasrabuddhe, and Jian Wang for valuable discussions that improved this research.

Finally, I thank my family for their continuous support that brought me to this point.

April 2003 Ayşegül GENÇATA, M.Sc.

CONTENTS

ABBREVIATIONS

ADM	: Add Drop Multiplexer
ATM	: Asynchronous Transfer Mode
CATZ	: Capacity Allocation with Time-Zones
CDM	: Code Division Multiplexing
CWDM	: Coarse Wavelength Division Multiplexing
DFC	: Difference Frequency Converter
DWDM	: Dense Wavelength Division Multiplexing
EDFA	: Erbium-Doped Fibre Amplifier
FDDI	: Fibre Distributed Data Interface
FTTB	: Fibre to the Building
FTTH	: Fibre to the Home
FWM	: Four-Wave Mixing
ILP	: Integer Linear Program
IP	: Internet Protocol
ISP	: Internet Service Provider
MEMS	: Micro-Electro-Mechanical Systems
MILP	: Mixed-Integer Linear Program
OC	: Optical Channel
OEO	: Optics-Electronics-Optics
OPS	: Optical Packet Switching
OXC	: Optical Cross-Connect
RWA	: Routing and Wavelength Assignment
SDH	: Synchronous Digital Hierarchy
SOA	: Semiconductor Optical Amplifier
SONET	: Synchronous Optical Network
TDM	: Time Division Multiplexing
WADM	: Wavelength Add Drop Multiplexer
WDM	: Wavelength Division Multiplexing
WDMUX	: Wavelength Demultiplexer
WMUX	: Wavelength Multiplexer
WWW	: World-wide Web

LIST OF TABLES

LIST OF FIGURES

Page No

TOPOLOGY AND BANDWIDTH ADAPTATION IN OPTICAL WDM BACKBONE NETWORKS WITH DYNAMIC TRAFFIC

SUMMARY

Communication networks reach more people everyday, providing new means of information exchange. Consequently, traffic demand is growing rapidly, mainly due to data-centric applications. Today, the major technology that is promising to meet this high bandwidth demand of networks is optical networking. Optical networks use optical fibres as transport medium and they are designed to exploit the fibre's unique properties. The bandwidth of a single fibre is nearly 50 Tb/s, which can be used by dividing it into smaller bands or channels, and using these channels concurrently. This is accomplished by wavelength division multiplexing (WDM), where each channel operates at a different wavelength and the electronic equipment should operate only at the bitrate of a wavelength channel (up to 40 Gb/s), which can be chosen arbitrarily.

It is becoming clear that IP-over-WDM (Internet Protocol over WDM) architecture will be the winning combination of the future, due to the success of IP as a convergence layer, and WDM as a bandwidth-rich physical layer. The challenge is to design an intelligent control plane that could realise the integration of the two layers, by solving interactive issues, such as, lightpath routing coupled with IP routing and resource management, reconfigurability, survivability, online service provisioning, monitoring.

The future's perspective for optical networking can be divided into two categories: Circuit switched, and packet switched. Packet switching technology is still in a very early stage, seen as a far future possibility by most researchers, because of several technical obstacles. Circuit switching is considered to be a realistic approach to design the optical networks for the near future. This thesis is based on a circuit switched optical layer, i.e., the main topic is wavelength-routed optical WDM networks.

Wavelength-routing is a major advantage of a WDM optical network. A wavelength-routed WDM network can provide end-to-end optical communication channels through optical fibres and intermediate nodes with optical cross-connects, even the source and destination nodes are not connected by a fibre directly. These optical channels (lightpaths) eliminate extra signal processing at intermediate nodes along the physical path through which the lightpath is routed. With the set up of a lightpath, two nodes become virtually neighbors. However, it may not be possible to establish a lightpath for every node pair, because of scalability and economic concerns. Hence, some traffic may need to be switched electronically from one lightpath to another at intermediate nodes until it reaches its destination; this approach is called multi-hopping. The processes of setting up individual lightpaths are clearly related to one another since a lightpath may carry multi-hop traffic besides the single-hop traffic between the two nodes it directly connects. For this reason, the design of the lightpath topology (virtual topology) is a combined problem of optimizing the use of network resources and network throughput, for a given traffic demand.

On the other hand, the traffic rates between node pairs of a network fluctuate distinguishably over time. The network resources can be arranged to accommodate a specific traffic matrix; however, the traffic matrix itself changes through time. In such case, a virtual topology which is optimized for a given traffic demand –which would be a snapshot of the changing traffic matrix– may not be able to respond with equal efficiency to a different traffic demand. Thus, the virtual topology should also be changed to match itself with the changing traffic. The dynamic structure of the optical cross-connects, i.e., the ability of switching the wavelengths from any input fibre to any output fibre dynamically, allows this change in the optical layer. This is known as reconfiguration of the virtual topology. Reconfiguration problem includes the virtual topology problem by its definition. Therefore, it is also an optimisation problem of several metrics, which needs to be solved on-line, in contrast to virtual topology design problem.

WDM networks have been a popular research area in the past decade, and many algorithms have been proposed to solve the design, routing, resource allocation, and reconfiguration problems. The objective of this thesis is to provide a new perspective for the reconfiguration of wide-area WDM networks. This objective is motivated by the problems originating from the usual view of the previous approaches. As a general assumption, reconfiguration studies are based on the idea that the decision of reconfiguration is sudden, triggered by an event, and consequently virtual topology change is an interrupted process. In this view a virtual topology V_1 designed for traffic matrix T_1 remains the same until the traffic changes to a matrix T_2, which in turn triggers the design of another virtual topology V_2. On the other hand, previous studies lack to design a proper triggering mechanisms, i.e., when and how to decide to reconfigure the virtual topology in a backbone network. Another problem with this two-step approach is that the future traffic is assumed to be known. In practise, a flexible reconfiguration method should deal with dynamic traffic, even when the future traffic is different than the expected matrix. Furthermore, since the reconfiguration is a complex problem, most of the methods proposed are useful for small-size networks, even though they are local search methods. For large networks, simple and effective algorithms are needed to provide on-line reconfiguration of the optical layer.

Observations on real networks show that the traffic rates between node pairs fluctuate distinguishably over time. The difference between the traffic intensity at highest point (busy hours) and the traffic intensity at lowest point (non-busy hours) is usually remarkable. Thus, reconfiguration may be very beneficial in terms of economical use of the bandwidths. One can also observe that there are two types of variations in traffic intensity, short-term variations (typically seconds or minutes) and long-term variations (typically hours). In long-term variations, which are the reason of reconfiguration, the amount of traffic between nodes changes in a smooth and continuous manner. Therefore, a virtual topology adaptation mechanism that changes the topology slowly, can be beneficial for backbone networks that have similar traffic behaviour.

The new perspective introduced in this thesis for reconfiguring the virtual topology in WDM networks under dynamic traffic is called virtual topology adaptation. Here, the problem is defined from a new point of view, different from the traditional two-step reconfiguration. The new approach sees the reconfiguration as a continuous process of observations and adjustments. This approach eliminates the need for traffic forecasts and critical triggering decisions, while providing robustness and speed.

The new approach introduced in this thesis can be viewed as a one-step reconfiguration approach in contrast to the previous methods. It is composed of continuous observation-adjustment cycles, where small changes may occur in the virtual topology. The traffic on each lightpath is observed continuously, and a new lightpath is added to or an existing lightpath is removed from the virtual topology when a change is necessary. The adaptation process is a continuous system where small adjustments are made, instead of waiting for a noticeable drop in system efficiency and changing the entire topology. With this new definition of the problem, the aim is to solve the issues of unpredictability of the traffic, critical triggering decisions, and the delay that is introduced by complicated reconfiguration algorithms. The ultimate goal is to obtain a scheme, which is practical for real backbone networks.

The basis of the new adaptation method is to apply simple adjustments to the virtual topology when it is necessary. The periodical measurements on network traffic provides information on the lightpath loads. This information is used to decide whether or not an adjustment is needed. Basically, a new lightpath is added when a congestion is encountered, i.e., a lightpath load is dangerously high, and a lightpath is deleted if it is being under-utilized, i.e., its load is very low. Therefore, this scheme tends to balance the traffic loads evenly to the lightpaths, and to modify the connectivity of nodes when the balance changes. Two system parameters are introduced to detect the lightpath-usage efficiencies: High watermark W_H and low watermark W_L. The traffic loads on lightpaths are measured at the end of each *observation period*, which is typically hundreds of seconds. At the end of an observation period, if the load of one or more lightpaths is higher than W_H, a new lightpath is established to decrease that load. When the load of a link drops below W_L, that link will be torn down with one condition: if alternate paths exist for diverting the traffic that uses the considered lightpath.

The approach proposed can be used by an ISP to minimise the operational cost of its virtual topology by leasing only the appropriate amount of lightpaths as it considers to be really necessary. Previous reconfiguration and virtual topology design studies usually maximises the number of lightpaths in a virtual topology. The aim of these methods is to use all possible network resources to carry maximum amount of traffic. On the other hand the fibre infrastructure may not belong to the ISP, in which case the ISP has to lease the lightpaths. This is a common business model in real world, hence, the topology design and reconfiguration methods are required to provide cost-efficient solutions. The new adaptation method is novel also in this sense, since it maintains only the necessary lightpaths. The new adaptation scheme is computationally simple when compared to previous studies. Simplicity is vital for a reconfiguration method, because the process should be on-line.

In this study, the adaptation method is first stated formally as a mixed-integer linear program (MILP). The performance of the new adaptation scheme is compared to an earlier reconfiguration method (optimal reconfiguration), by solving the MILP formulation. The results of this comparison shows that the adaptation scheme can reconfigure the topology using far fewer lightpath changes compared to the optimal reconfiguration.

Formulations aiming to find an optimal solution for the reconfiguration in large networks are intractable, hence heuristic approaches are needed. Therefore, in the next step, a heuristic algorithm is introduced for the virtual topology adaptation method proposed. Simulation experiments are conducted to investigate the fitness of the adaptation approach and to expose the effect of various system parameters. Various performance metrics are employed to reveal the different aspects of the adaptation scheme, such as, the average time between consecutive reconfiguration steps, the traffic-weighted average hop distance, number of lightpaths in the virtual topology, and the percentage of links with loads in the balanced region defined as $[W_L, W_H]$. Through these metrics, the effects of changing the system parameters are examined, specifically, high watermark, low watermark, and length of observation period. These experiments showed that the algorithm is reacting to the load changes as expected, and the two main system parameters, namely high and low watermarks, can be effectively used to regulate the operation of the proposed virtual-topology adaptation method.

As a second contribution, this thesis proposes a bandwidth reconfiguration method called CATZ (Capacity Allocation with Time Zones), that exploits the daily traffic fluctuations and time-zone based traffic shifts to cost-effectively operate a world-wide WDM backbone network. The main assumption of this method is that, in the near future the bandwidth markets will allow fast online bandwidth provisioning for service providers. Here, the reconfiguration problem in an IP-over-WDM network is investigated from the point of view of an ISP under the light of the new challenges posed by the emerging bandwidth markets, and different service types in leasing bandwidths.

Previous studies on virtual topology design concentrate on routing the IP traffic appropriately, so that the lightpath connectivity can carry a large amount of traffic. These studies consider a constant capacity for any lightpath and do not take into account the sub-wavelength bandwidth granularities. They also contravene the objective of the ISP, to minimise the network cost for high profit; because they try to set up as many lightpaths as possible, instead of an appropriate and cost-effective working set of lightpaths. Network reconfiguration studies are similar to design studies in both aspects, and their application to future real networks is difficult.

In this work, a cost-efficient method is investigated for bandwidth allocation and bandwidth reconfiguration for global IP-over-WDM backbone networks. IP traffic is very dynamic, but in a backbone network, due to traffic aggregation, this dynamism becomes smoother, following a well-shaped daily profile. This property can be exploited by WDM technology, since wavelength-routed WDM networks provide a flexible infrastructure that can adjust the network topology (connectivity and bandwidth) by reconfiguring the optical cross-connects. On the other hand, when a world-wide network traffic structure is examined, another dimension of dynamism can be seen: traffic shifts based on the time zones. A typical traffic profile between two nodes follow a sinusoidal shape, where the peak is around midday and the valley around early hours of the morning, e.g. 03.00. Since different locations in the world are in different time zones, each node will create a traffic whose intensity changes according to that node's time zone. When this new type of dynamism is considered, the bandwidth cost of the network can be reduced, even in the design step of the virtual topology. The study in this thesis show that a large reduction on bandwidth cost can be achieved by reconfiguring the lightpath capacities.

In CATZ, the channels are observed periodically to measure the traffic intensities flowing through them, and the bandwidth of the channel for the next period is set according to these intensities. In this method, the capacities of the channels are allowed to change, since the underlying transport system, optical switches, and the global carriers are assumed to provide on-line bandwidth services. The amount of step increase or decrease can be a multiple of OC-1, since the bandwidth markets provide optical channels at capacities of OC-3, OC-12, OC-48, or OC-192. Parallel channels are also possible between nodes, as long as there are enough transceivers. The new approach is compared with two static bandwidth allocation schemes using network bandwidth cost as the performance metric. The experiments show that an important cost improvement can be obtained by using a bandwidth reconfiguration scheme that takes into consideration the dynamism of the traffic.

In summary, this thesis brings the virtual-topology reconfiguration problem to a new perspective where it is considered as a continuous process, and by using this new approach, it proposes two new methods for fast online topology adaptation. The results obtained in this thesis are encouraging for building the backbone reconfigurable WDM networks, since they show that simple and efficient algorithms that could properly manage the network resources can be designed. The new view introduced in this work, namely *adaptation* instead of *reconfiguration*, may propel new methods that are more efficient for the network operators.

DEĞİŞKEN VERİ TRAFİKLİ OPTİK WDM OMURGA AĞLARINDA TOPOLOJİ VE BANT GENİŞLİĞİ UYARLAMA

ÖZET

İletişim ağları hergün daha çok kişiye erişiyor ve yeni bilgi değişim olanakları sağlıyor. Bunun sonucunda, özellikle veri temelli uygulamalara bağlı olarak trafik isteği hızla artıyor. Bugün, bilgisayar ağlarının bu yüksek bant genişliği isteğini karşılayabilecek başlıca teknoloji optik ağlardır. Optik ağlar iletişim ortamı olarak optik fiberleri kullanırlar ve fiberin benzersiz özelliklerinden yararlanmak için tasarlanmışlardır. Tek bir fiberin bant genişliği yaklaşık olarak 50 Tb/s'dir, bu kapasite daha dar bant ya da kanallara bölünerek ve bu kanallar eşzamanlı çalıştırılarak kullanılabilir. Bu, her kanalın farklı bir dalgaboyunda çalıştığı ve elektronik parçaların yalnızca bir dalgaboyu kanalının istendiği gibi seçilebilen bit hızında (40 Gb/s'e kadar) çalışabildiği, dalgaboyu bölmeli çoğullama (WDM) ile gerçeklenebilir.

IP'nin (Internet Protokolu) yakınsama katmanı olarak başarısına ve WDM'in bantgenişliği açısından zengin bir fiziksel katman olmasına bağlı olarak, IP/WDM (Internet Protokolu/WDM) mimarisinin geleceğin kazanan kombinasyonu olacağı giderek belirginleşiyor. Bu noktada zorluk; IP yönlendirme ve kaynak yönetimi ile bağlantılı olan ışıkyolu yönlendirme, yeniden konfigüre edilebilirlik, bağışıklık, anında hizmet sağlama, izleme gibi etkileşim gerektiren konuları çözerek, iki katmanın entegre edilmesini sağlayabilecek, akıllı bir denetim ortamı tasarlamaktır.

Optik ağlar için geleceğe yönelik bakış açısı iki bölüme ayrılabilir: Devre anahtarlamalı ve paket anahtarlamalı. Paket anahtarlama teknolojisi hala oldukça erken aşamadadır ve birçok teknik engeller nedeniyle, çoğu araştırmacı tarafından uzak gelecekte bir olasılık olarak görülmektedir. Devre anahtarlama yakın geleceğin optik ağlarını tasarlamak için gerçekçi bir yaklaşım olarak değerlendirilmektedir. Bu tez devre anahtarlamalı optik katmanı temel almaktadır, bu nedenle ana konu dalgaboyu yönlendirmeli optik WDM ağlarıdır.

Dalgaboyu yönlendirme optik WDM ağlarının önemli bir avantajıdır. Dalgaboyu yönlendirmeli bir WDM ağı, kaynak ve varış düğümleri bir fiber ile doğrudan bağlı olmasa bile, optik fiberler ve ara düğümlerdeki optik anahtarlar aracılığı ile uçtan uca optik iletişim kanalları sağlar. Bu optik kanallar (ışıkyolları) fiziksel olarak ışıkyolunun üzerinden yönlendirildiği ara düğümlerde fazladan işlemi ortadan kaldırır. Bir ışıkyolunun kurulması ile iki düğüm görüntü olarak komşu durumuna gelir. Bununla birlikte, ölçeklenebilirlik ve ekonomik kaygılar nedeniyle, her düğüm çifti için bir ışıkyolu kurmak mümkün olmayabilir. Böylece, trafiğin bir kısmının, varışa erişene kadar bir ışıkyolundan diğerine, ara düğümlerde elektronik olarak yönlendirilmesi gerekebilir; bu yaklaşım çok sıçramalı olarak adlandırılır. Işıkyollarının ayrı ayrı kurulması işlemi, açıkça görüleceği gibi birbirlerine bağlıdır, çünkü bir ışıkyolu doğrudan bağladığı iki uç düğüm arasındaki tek hop trafiğin yanısıra çok hoplu trafik de taşıyabilir. Bu nedenle, ışıkyolu topolojisinin (sanal topoloji)

tasarımı, verilen bir trafik isteğine göre ağ kaynaklarının kullanımını ve ağ verimini eniyilemeyi hedefleyen birleşik bir problemdir.

Diğer taraftan, ağdaki düğümlerin arasındaki trafik hızları zaman içerisinde belirgin biçimde dalgalanmaktadır. Ağ kaynakları belirli bir trafik matrisine hizmet verecek şekilde düzenlenebilir; ancak, trafik matrisinin kendisi zaman içinde değişmektedir. Bu durumda, verilen trafik isteğine –bu istek değişen bir trafik matrisinin zamanda bir kesitidir– göre eniyilenen bir sanal topoloji, farklı bir trafik isteğini aynı etkinlikle karşılayamayabilir. Bu nedenle, sanal topoloji de değişen trafiğe uyum sağlamak için değiştirilmelidir. Optik anahtarların dinamik yapısı, bir başka ifadeyle, dalgaboylarını herhangi bir giriş fiberinden herhangi bir çıkış fiberine dinamik olarak anahtarlama yeteneği, optik katmanda böyle bir değişikliğe olanak verir. Bu özellik, sanal topolojinin yeniden konfigürasyonu olarak bilinir. Yeniden konfigürasyon problemi tanım gereği sanal topolojinin tasarımı problemini de içermektedir. Bundan dolayı birçok ölçüt içeren bir eniyileme problemidir, ve sanal topoloji tasarımına karşıt olarak çevrimiçi çözülmesi gerekmektedir.

WDM ağları geçen on yılda popüler bir araştırma alanı olmuş ve tasarım, yönlendirme, kaynak ayırma ve yeniden konfigürasyon problemlerini çözmek için birçok algoritma önerilmiştir. Bu tezin amacı geniş alanlı WDM ağlarının yeniden konfigürasyonu için yeni bir bakış açısı ortaya koymaktır. Bu amaç, önceki yaklaşımların geleneksel bakış noktasından kaynaklanan problemler ile motive edilmektedir. Genel bir kabul olarak, yeniden konfigürasyon konusundaki çalışmalar, konfigürasyon kararının ani verildiği, bir olay ile tetiklendiği, ve bunun sonucunda sanal topoloji değişikliğinin kesintili bir süreç olduğu kanısına dayanmaktadır. Bu bakışta T_1 trafik matrisi için tasarlanmış bir V_1 sanal topolojisi, trafik T_2 matrisine değişene kadar aynı kalır, bu yeni matris yeni bir V_2 sanal topolojisinin tasarlanmasını tetikler. Diğer yandan, önceki çalışmalar uygun bir tetikleme mekanizması, yani bir omurga ağında sanal topolojinin konfigürasyonu kararının ne zaman ve nasıl verileceğine ilişkin bir yöntem geliştirme konusunda yetersiz kalmıştır. Bu iki aşamalı yaklaşımda bir diğer problem ise, gelecekteki trafiğin bilindiğinin varsayılmasıdır. Pratikte, esnek bir konfigürasyon yöntemi dinamik trafik ile, gelecekteki trafik beklenen trafik matrisinden farklı bile olsa, başa çıkabilmelidir. Buna ek olarak, yeniden konfigürasyon karmaşık bir problem olduğundan, önerilmiş olan yöntemlerin çoğu yerel arama yöntemi olmalarına rağmen, küçük boyda ağlar için kullanışlıdır. Büyük ağlar için, optik katmanın çevrimiçi konfigürasyonunu sağlayacak, kolay ve etkin algoritmalara gereksinim duyulmaktadır.

Gerçek ağlarda yapılan gözlemler, düğümler arasındaki trafiğin zaman içinde belirgin biçimde değiştiğini göstermektedir. Trafiğin en yüksek olduğu andaki (yoğun saatler) hız ile en düşük olduğu andaki (yoğun olmayan saatler) hızın farkı genelde oldukça farklıdır. Bundan dolayı, yeniden konfigürasyon bant genişliğinin ekonomik kullanımı açısından oldukça faydalı olabilir. Trafik hızında iki tür değişkenlik olduğu gözlenebilir, kısa vadeli değişimler (genelde saniye ya da dakika mertebesinde) ve uzun vadeli değişimler (genelde saatler). Yeniden konfigürasyonun nedeni olan uzun vadeli değişimlerde, düğümler arasındaki trafik miktarı sürekli ve yumuşak bir şekilde değişir. Bu yüzden, topolojiyi yumuşak bir biçimde değiştiren bir sanal topoloji yeniden konfigürasyon mekanizması, bu tür bir trafik davranışı olan omurga ağlarında yararlı olabilir.

Dinamik trafik altındaki WDM ağlarında sanal topolojiyi yeniden konfigüre etmek için bu tezde önerilen yeni bakış açısı sanal topoloji uyarlama olarak adlandırılmıştır. Burada, problem önceki geleneksel iki aşamalı yeniden konfigürasyondan farklı olan yeni bir yaklaşımla tanımlanmıştır. Yeni yaklaşım yeniden konfigürasyonu gözleme ve ayarlamaların kesintisiz bir süreci olarak görmektedir. Bu yaklaşım trafik tahminleri ve kritik tetikleme kararlarına olan gereksinimi ortadan kaldırmakta, aynı zamanda gürbüzlük ve hız sağlamaktadır.

Bu tezde tanıtılan yeni yaklaşım öncekilere karşıt olarak tek aşamalı konfigürasyon olarak görülebilir. Yeni yaklaşım sanal topolojide küçük değişikliklerin olabildiği sürekli gözlem-ayar çevriminden oluşmaktadır. Her ışıkyolundaki trafik sürekli olarak gözlenmekte ve bir değişiklik gerektiğinde sanal topolojiye yeni bir ışıkyolu eklenmekte ya da burada olan bir ışıkyolu çıkarılmaktadır. Uyarlama süreci, sistem verimliliğinde gözle görülür bir düşüş olmasını beklemek ve bundan sonra tüm sanal topolojiyi değiştirmek yerine küçük ayarlamaların sürekli yapıldığı bir sistemdir. Problemin bu yeni tanımı ile amaç trafiğin tahmin edilebilirliği, kritik tetikleme kararları ve karmaşık yeniden konfigürasyon algoritmalarından kaynaklanan gecikme sorunlarının çözümüne olanak vermektir. Temel hedef gerçek omurga ağları için pratik olan bir yöntem elde etmektir.

Yeni uyarlama yönteminin temeli sanal topolojiye gerektiği zaman basit ayarlamalar uygulamaktır. Ağ trafiği üzerinde yapılan periyodik ölçümler ışıkyollarının yükleri hakkında bilgi sağlamaktadır. Bu bilgi ayarlamanın gerekli olup olmadığına karar vermek için kullanılır. Temel olarak, bir sıkışıklığa rastlandığında, yani bir ışıkyolunun yükü tehlikeli biçimde yükseldiğinde yeni bir ışıkyolu eklenmekte, bir ışıkyolu etkin olarak kullanılmıyorsa yani yükü çok düşükse o ışıkyolu kaldırılmaktadır. Böylece, bu yöntem ışıkyollarına trafik yükünü uygun şekilde dağıtmaya, ve yük dengesi bozuldukça düğümler arasındaki bağlılığı düzenlemeye çalışmaktadır. Işıkyolu kullanımının etkinliğini saptamak için iki sistem parametresi tanımlanmıştır: Yüksek eşik W_H ve alçak eşik W_L. Işıkyollarındaki trafik yükleri, tipik olarak birkaç yüz saniye uzunlukta olan her m gözlem aralığı sonunda ölçülür. Bir gözlem aralığı sonunda bir ya da daha fazla ışıkyolunun yükü W_H'dan büyükse yükü azaltmak için yeni bir ışıkyolu kurulur. Bir ışıkyolundaki yük W_L'nin altına düştüğünde bu yol bir koşula bağlı olarak kaldırılır: Sözkonusu ışıkyolunu kullanan trafiğin yönlendirileceği başka bir yolun bulunması.

Önerilen yöntem bir ISP tarafından sanal topolojisinin işletim masrafını, yalnızca tam anlamıyla gerekli gördüğü miktarda ışıkyolunu kiralayarak, enaza indirgemek için kullanılabilir. Önceki yeniden konfigürasyon ve sanal topoloji tasarlama çalışmaları, genellikle sanal topolojideki ışıkyolu sayısını maksimize etmektedir. Bu yöntemlerin amacı en fazla miktarda trafik taşıyabilmek için mümkün olan tüm ağ kaynaklarını kullanmaktır. Diğer yandan fiber altyapısı ISP'e ait olmayabilir, bu durumda ISP ışıkyollarını kiralamalıdır. Bu gerçek dünyada sık rastlanan bir iş modelidir, bu nedenle topoloji tasarım ve yeniden konfigürasyon yöntemlerinin ekonomik çözümler sağlamasına gerek duyulmaktadır. Yeni uyarlama yöntemi bu anlamda da yenilik getirmektedir, çünkü yalnızca gerekli olan ışıkyollarını canlı tutmaktadır. Yeni uyarlama yöntemi önceki yöntemlerle kıyaslandığında işlem karmaşıklığı açısından daha basittir. Bir yeniden konfigürasyon yöntemi için basitlik yaşamsal bir öneme sahiptir çünkü sürecin çevrimiçi işlemesi gerekmektedir.

Bu çalışmada uyarlama yöntemi önce formel olarak bir karışık tamsayılı lineer program (MILP) olarak ortaya koyulmuştur. Bu MILP formülasyonu çözülerek yeni uyarlama yönteminin başarımı daha önceki bir yeniden konfigürasyon yöntemi (optimal rekonfigürasyon) ile karşılaştırılmıştır. Bu karşılaştırmanın sonuçları, uyarlama yönteminin topolojiyi optimal rekonfigürasyona kıyasla çok daha az ışıkyolu değişikliği yaparak uyarlayabildiğini göstermektedir.

Büyük ağlarda, yeniden konfigürasyon için optimal çözümü bulmayı hedefleyen formülasyonları çözmek çoğunlukla olanaksızdır, bu nedenle sezgisel (heuristic) yaklaşımlara gerek vardır. Bundan dolayı, bir sonraki aşamada, önerilen sanal topoloji uyarlama yöntemi için sezgisel bir algoritma geliştirilmiştir. Uyarlama yaklaşımının uygunluğunu araştırmak ve çeşitli sistem parametrelerinin etkisini ortaya çıkarmak için benzetim deneyleri yapılmıştır. Uyarlama yönteminin değişik yönlerini ortaya çıkarmak için, ardışıl uyarlama adımları arasındaki ortalama süre, trafik ağırlıklı ortalama hop uzaklığı, sanal topolojideki ışıkyolu sayısı ve $[W_L, W_H]$ aralığı olarak tanımlanan denge bölgesindeki yüklerin yüzdesi gibi çeşitli başarım ölçütleri uygulanmıştır. Bu ölçütler aracılığıyla sistem parametrelerini, yani yüksek eşik, alçak eşik ve gözlem aralığının uzunluğunu değiştirmenin etkileri incelenmiştir. Bu deneyler, algoritmanın yük değişimlerine beklendiği gibi tepki verdiğini, ve iki ana sistem parametresi yüksek ve alçak eşiklerin, önerilen uyarlama yönteminin çalışmasını ayarlamada etkin olarak kullanılabileceğini göstermiştir.

İkinci bir katkı olarak bu tez CATZ (Capacity Allocation with Time Zones) adı verilen bir bant genişliği uyarlama yöntemi önermektedir. Bu yöntem dünya çapında bir WDM omurga ağını ekonomik bir şekilde işletebilmek için günlük trafik dalgalanmaları ve zaman dilimi temelli trafik kaymalarından yararlanmaktadır. Bu yöntemde ana varsayım yakın gelecekte bant genişliği pazarlarının hizmet sağlayıcılara hızlı çevrimiçi bant genişliği atamaya izin vereceğidir. Burada, bir IP/WDM ağındaki yeniden konfigürasyon problemi, gelişmekte olan bant genişliği pazarları ve bant genişliği kiralamada değişik hizmet tipleri ile ortaya çıkan yeni zorlukların ışığı altında, bir ISP'nin bakış açısından araştırılmaktadır.

Sanal topoloji tasarımı üzerindeki önceki çalışmalar ışıkyolu bağlılığının yüksek miktarda trafik taşıyabilmesi için IP trafiğini uygun olarak yönlendirme konusunda yoğunlaşmaktadır. Bu çalışmalar bir ışıkyolu için sabit bir kapasite varsaymakta ve dalgaboyu altı bant genişliklerini gözönüne almamaktadır. Bunlar ayrıca ISP'nin amacı, yani ağ masrafını enaza indirerek yüksek kar sağlamak ile de çelişmektedir; çünkü uygun ve ekonomik bir çalışan ışıkyolu kümesi yerine mümkün oldukça çok sayıda ışıkyolu kurmaya çalışmaktadırlar. Ağ uyarlama çalışmaları tasarım çalışmalarına bu her iki konuda da benzerlik gösterirler ve bu tip çalışmaların gelecekteki gerçek ağlara uygulanmaları zordur.

Bu çalışmada IP/WDM omurga ağlarında bant genişliği atama ve bant genişliğini yeniden konfigüre etmek için ekonomik bir yöntem araştırılmaktadır. IP trafiği oldukça dinamiktir, ancak bir omurga ağında, trafiğin üstüste binmesine bağlı olarak, bu dinamiklik yumuşar ve belirli bir şekle sahip günlük bir profil izler. WDM teknolojisi ile bu özellikten faydalanılabilir, çünkü dalgaboyu yönlendirmeli WDM ağları optik anahtarları düzenleyerek ağ topolojisini (bağlılık ve bant genişliği) ayarlayabilen esnek bir altyapı sunmaktadır. Diğer yandan, dünya çapında bir ağın trafik yapısı incelendiğinde, dinamikliğin başka bir boyutu görülebilir: trafik zaman dilimlerine göre kayar. İki düğüm arasındaki tipik bir trafik profili, tepenin gün ortası civarında olduğu ve vadinin sabahın erken saatlerine rastladığı (örneğin 03.00) sinüzoidal bir şekil izler. Dünyanın değişik bölgeleri farklı zaman dilimlerinde bulunduğundan her düğüm bulunduğu bölgenin zaman dilimine uygun olarak değişen bir trafik yaratır. Bu yeni dinamik yapı göz önüne alındığında ağın bant genişliği masrafı sanal topolojinin tasarım aşamasında bile azaltılabilir. Bu tezdeki çalışma, ışıkyolu kapasitelerini yeniden konfigüre ederek bant genişliği masrafında önemli düşüşler elde edilebileceğini göstermektedir.

CATZ yönteminde kanallar üzerlerinden akan trafik miktarını ölçmek amacıyla periyodik olarak gözlenir, ve kanalların bir sonraki zaman aralığındaki bant genişlikleri bu yüklere göre belirlenir. Bu yöntemde kanal kapasitelerinin değişmesine izin verilmektedir, çünkü alttaki iletim sistemi, optik anahtarlar ve küresel taşıyıcıların çevrimiçi bant genişliği hizmeti sağladığı varsayılmaktadır. Bant genişliği pazarları OC-3, OC-12, OC-48 ya da OC-192 kapasitelerinde optik kanallar sağlayabildiklerinden kanal kapasitelerinde adım artırma ve azaltma miktarları OC-1'in katları olabilir. yeterli sayıda alıcı ve verici olduğu sürece düğümler arası paralel kanallar da kurulabilir. Yeni yaklaşım, ağın bant genişliği masrafı başarım ölçütü olarak kullanılarak, iki statik bant genişliği atama yöntemiyle karşılaştırılmıştır. Deneyler trafiğin dinamikliğini dikkate alan bir bant genişliği uyarlama yöntemi kullanılarak masrafta önemli iyileştirmeler elde edilebileceğini göstermektedir.

Sonuç olarak bu tez, sanal topolojinin yeniden konfigürasyonu problemini, kesintisiz bir süreç olarak ele alındığı yeni bir perspektife taşımaktadır ve bu yeni yaklaşımı kullanarak hızlı çevrimiçi topoloji uyarlama için iki yeni yöntem önermektedir. Bu tezde elde edilen sonuçlar uyarlanabilir WDM omurga ağlarını kurmak için cesaret vericidir, çünkü ağ kaynaklarını uygun şekilde yönetebilen basit ve etkin algoritmaların tasarlanabileceğini göstermektedirler. Bu çalışmada tanıtılan yeni bakış açısı olan *yeniden konfigürasyon* yerine *uyarlama*, ağ yöneticileri için daha etkin yeni yöntemlerin tasarımına öncülük yapabilir.

1. INTRODUCTION

Communication networks reach more people everyday, providing new means of information exchange. Consequently, traffic demand is growing rapidly (exponentially), mainly due to data-centric applications, as reported by many service providers [2]. This growth is triggering the development of new technologies, and the advance of existing technologies. Methods to exploit these advances need to be developed to build our future knowledge society.

The major technology that is promising to meet the high bandwidth demand of data networks is optical networking. Optical fibre has been used as physical medium by many networking technologies, such as FDDI, SONET, ATM, due to its high bandwidth and low signal degradation. Optical networks, on the other hand, are designed to exploit the fibre's unique properties, and they are separated from the other fibre-based technologies with this characteristic. The potential bandwidth of a fibre is nearly 50 Tb/s, a speed that today's electronic processing capacity cannot match. Such a large bandwidth can be used by dividing it into smaller bands or channels, and using these channels concurrently. This can be accomplished by a multiplexing technique. Among time division multiplexing (TDM), code division multiplexing (CDM), wavelength (or frequency) division multiplexing (WDM), optical TDM and CDM are futuristic technologies today [3]. WDM is currently the favourite multiplexing method in optical networks since the electronic equipment should operate only at the bitrate of a wavelength channel, which can be chosen arbitrarily. Optical networks with WDM are the choice of the network service providers, and are being deployed around the world. Typically, a backbone optical network infrastructure has tens (sometimes hundreds) of fibres between core network nodes, each fibre capable of carrying hundreds of wavelength channels, and each channel operating at a transmission rate up to 40 Gb/s.

At present, WDM technology is used mostly as point-to-point links between electronic switching nodes. The reason is that today's optical switch architectures lack the proper control mechanisms that would interact with the upper layers to provide efficient and flexible services at the optical layer. It is becoming clear that IP-over-WDM (Internet Protocol over WDM) will be the winning combination of the future, due to the success of IP as a convergence layer, and WDM as a bandwidth-rich transport layer [4]. The challenge is to design an intelligent control plane that could realise the integration of the two layers, by solving interactive issues, such as, lightpath routing coupled with IP routing and resource management, reconfigurability, survivability, online service provisioning, monitoring. Efforts are on their way to design a common control plane that would allow the realisation of next generation IP-over-WDM networks [5, 6, 7, 8].

The future's perspective for optical networking can be divided into two categories: Circuit switched, and packet switched. The future's circuit switched WDM network will consist of the WDM fiber-optic links and WDM-aware nodes connecting them. The optical nodes in such a network include optical cross-connects (OXC), transmitters and receivers, and these nodes are capable of separately routing the various wavelength channels in the optical domain. In the long term, optical packet switching [9, 10, 11] could become a viable candidate as a networking technology, because of its fine-granularity switching, and flexibility, but the technology is still in a very early stage, seen as a far future possibility, because of several technical obstacles. For now, circuit switched approach seems to be a realistic approach to design the optical networks. This thesis is based on a circuit switched optical layer, i.e., our main topic is wavelength-routed optical WDM networks.

There has been long a debate on the design of optical cross-connects that affect the design of the whole network: all-optical networks versus optics-electronics-optics (OEO). In the former, all transport layer activities are in optical domain, which means that a signal that enters the network from the source node remains in optical domain until it reaches its destination. In the latter, the optical signals are converted into electronic signals at the switching stage of a cross-connect where they can be demultiplexed, routed, and multiplexed at electronic domain, and they are converted back to optical signals before they leave the switch. Both approaches have their supporters and investors; OEO approach seems easier to implement, but all-optical approach is gaining speed in the last years. For the objective of this thesis both approaches are suitable, since they do not affect directly the performance of the algorithms proposed.

Wavelength-routing is a major advantage of a WDM optical network. A wavelength-routed WDM network can provide end-to-end optical communication channels through optical fibres and intermediate nodes with optical cross-connects [12], even the source and destination nodes are not connected by a fibre directly. These optical channels (lightpaths) [13] eliminate extra signal processing at intermediate nodes along the physical path through which the lightpath is routed. With the set up of a lightpath, two nodes become virtually neighbors. However, it may not be possible to establish a lightpath for every node pair, because of scalability and economic concerns. Hence, some traffic may need to be switched electronically from one lightpath to another at intermediate nodes until it reaches its destination; this approach is called multi-hopping. The processes of setting up individual lightpaths are clearly related to one another since a lightpath may carry multi-hop traffic besides the single-hop traffic between the two nodes it directly connects. For this reason, the design of the lightpath topology (virtual topology) is a combined problem of optimizing the use of network resources and network throughput, for a given traffic demand.

On the other hand, another issue arises when real network traffic is observed: the traffic rates between node pairs fluctuate distinguishably over time. The network resources can be arranged to accommodate a specific traffic matrix, however, the traffic matrix itself changes through time. In such case, a virtual topology which is optimized for a given traffic demand–which would be a snapshot of the changing traffic matrix–may not be able to respond with equal efficiency to a different traffic demand. Thus, the virtual topology should also be changed to match it with the changing traffic. The dynamic structure of the optical cross-connects, i.e., the ability of switching the wavelengths from any input fibre to any output fibre dynamically, allows this change in the optical layer. This is known as reconfiguration of the virtual topology. Reconfiguration problem includes the virtual topology problem by its definition. Therefore, it is also an optimisation problem of several metrics, which needs to be solved on-line, in contrast to virtual topology design problem.

Reconfiguration of wavelength-routed networks has been studied, and several virtual-topology reconfiguration schemes have been proposed [14, 15, 16, 17, 18, 19, 20]. These studies have a general approach to the problem, they divide it into two sub-problems: Designing the new virtual topology, and changing the lightpaths so that the new topology is established. Consequently, the two subproblems are usually solved separately, and some studies are concentrated in one of the sub-problems.

This traditional approach of solving the reconfiguration problem has several drawbacks. A major difficulty is that the future traffic demand is assumed to be known. With this information in hand, the design of a new virtual topology is practicable, and in many studies considerable effort was spent for the design of the next virtual topology so as to reduce the topology changes during the transition phase [14, 17, 19]. When the future traffic can be predicted, the computation of the new topology can be started well in advance, giving enough time to complete complex computations. Such an approach is similar to off-line topology design, and very complicated algorithms were proposed in early studies.

In practice, the assumption that the future traffic is known, or predictable, may be too optimistic. When we are dealing with real data networks, the algorithms and design methods need to be robust enough to deal with unpredictable changes in the traffic. Since the value of the virtual topology is highly related to the traffic it is carrying, the reconfiguration method should be flexible to spontaneously reacting to the new conditions. This thesis departs of the view that a reconfiguration method should not make any assumptions on the future traffic matrices, and should be able to adapt the virtual topology without relying on traffic forecasts.

Reconfiguration of virtual topology in a network with dynamically changing traffic is an on-line process, i.e., it should be completed in a small time interval compared to the time scale in which considerable traffic changes occur. This implies that we need fast algorithms that can update the virtual topology quickly. Considering the complexity of the problem, and that optimal solutions are difficult to obtain even for small-size networks, we have to turn to heuristic methods that can provide good sub-optimal solutions. Second basic point of this study is that we need fast and simple reconfiguration schemes usable for large networks.

The two sub-problem or two step approach of earlier studies brings another decision question: *When should the reconfiguration occur?* The changes in traffic intensities in a network are more likely to be smooth, increasing or decreasing slowly. How much change in traffic would be enough to trigger a reconfiguration event? This is a difficult question to answer, since it is an optimisation problem itself, including the calculation of trade-off between the amount of inefficiency in using network resources if the virtual topology does not change and the cost of reconfiguration. Reconfiguration process in two step approach is a costly operation, so the triggering decision is critical. Our third basic point is eliminating the critical and costly decisions as much as possible, shaping the solution approach to enable self corrections.

Given these bases, the aim of this thesis is to provide a new perspective for reconfiguration of WDM networks under dynamic traffic. The reconfiguration methods proposed in this study are designed for wide-area WDM mesh networks with IP-over-WDM architecture. A new one-step reconfiguration approach is derived and two reconfiguration methods are proposed for two different network cases. The performance of the reconfiguration methods are shown through simulations, and in one case, by solving also the optimisation problem with an optimiser.

In our new approach, the problem is defined differently: As traffic fluctuates over time, it will be monitored systematically, and the virtual topology will be changed accordingly, but we do not make any other assumptions on future traffic pattern. The reconfiguration process is seen as a continuous measurement-adaptation system where small adjustments are made, instead of waiting for a noticeable drop in system efficiency and changing the entire topology. With this new definition of the problem, the unpredictability of the traffic is solved by continuous monitoring and we aim to obtain a scheme which is practical for real backbone networks. The new method is called adaptation rather than reconfiguration, since the network is allowed to adapt itself as the traffic changes slowly.

The second method developed in this study provides a new type of reconfiguration for virtual topology: bandwidth reconfiguration. Although a lightpath is usually considered to have a fixed bandwidth in many previous studies, in real networks the service providers may obtain optical channels at different capacities. Such flexible bandwidth channels might be expected to become commonplace in the future with the development of traffic grooming architectures and bandwidth markets. Our aim is to investigate cost efficient methods for bandwidth allocation and reconfiguration, by exploiting a second type of dynamism in backbone traffic: in a world-wide network traffic intensities shift according to the time-zones. We propose to increase or decrease the bandwidth of a connection when total traffic flowing through it increases or decreases. This method exploits the online bandwidth provisioning capability of the network to minimise the bandwidth cost. We show that by using our proposed method, CATZ (Capacity Allocation with Time Zones), the bandwidth cost for operating a world-wide WDM backbone network can be dramatically reduced if the daily traffic dynamism and time-zone effect are considered.

The content of the chapters of this thesis are as follows:

- Chapter 2 summarises today's communication network structure, by briefly describing access, metro and backbone networks from a point of view of optical networks. Optical technology is then introduced, including transmission of optical signal, optical fibres, impairments, and basic optical backbone network elements. This chapter concludes with general terminology of optical networks, and an analysis of the state of the art.
- Chapter 3 introduces the new virtual topology adaptation method. The formal definition of the problem is given and the proposed solution approach is formulated using mixed-integer linear program. The MILP is solved to compare the performance of the new method with an optimal reconfiguration scheme.
- Chapter 4 provides a new heuristic algorithm for the adaptation method introduced in the previous chapter. The algorithm given in this chapter is practical for large networks. Simulation results are used to discuss the performance of the adaptation algorithm.
- Chapter 5 introduces the new method bandwidth adaptation method in a global network. This chapter discusses the emerging network model for the future, and it provides a comparison in terms of bandwidth cost between static and dynamic bandwidth assignment schemes.
- Conclusion summarises the outcome of the thesis, and discusses the possible future research directions.

2. OPTICAL WDM NETWORKS

In this chapter today's communication network structure is summarised, by describing different scale of networks: access, metro and backbone. The chapter develops on backbone networks' point of view, and introduces the different technologies including transmission of optical signal, optical fibres, and basic optical backbone network elements. Next, the general terminology of optical networks is explained, and the research on optical networks is summarised. A detailed review of virtual topology reconfiguration studies follows the general state-of-the-art discussion.

2.1 Network Structure

Fibre-optic technology is described in the following section, emphasising its use in networks of different scale: access, metropolitan, and backbone. Then, the wavelength-division multiplexing technique is introduced.

2.1.1 Today's Communication Networks

Computer networks' aim is to provide communication services to every person in our society. With the dissemination of Internet we gain the possibility of accessing any information we desire from any corner of the world. We want to reach this information easily at any time, and with high speed. The demand for data traffic grows exponentially both with the growth of the population having access to the Internet, and with the bandwidth-intensive applications such as data browsing on the world-wide web (WWW), java applications, video conferencing, interactive distant learning, on-line games. This exponential growth in the data traffic needs to be supported by high-bandwidth networks whose capabilities are much beyond those of current high-speed networks.

Fibre-optic technology can meet the above mentioned need, because of its potentially limitless capabilities [12, 21]: huge bandwidth (nearly 50 terabits per second), low signal attenuation, low signal distortion, low power requirements, low cost. Because of all of these qualities, several networking technologies use fibres as their physical layer media, such as ATM, SONET, FDDI, Gigabit Ethernet. On the other hand, optical networks that are the scope of this thesis are described as follows [22]: a telecommunication network with transmission links that are optical fibers, and with an architecture that is designed to exploit the unique features of fibres. Clearly, optical networks are very different from the other networking technologies in terms of their use of the fibre, and consequently the bandwidth they can provide to their users.

Today's communication networks can be considered to consist of three subnetworks: access (spanning about 1-10 km), metropolitan (covering about 10-100 km), and backbone (extending to 100s or 1000s of km) as shown in Fig. 2.1. Each of these subnetworks has a different set of functions to perform; hence, each has a different set of challenges, technological requirements, and research problems.

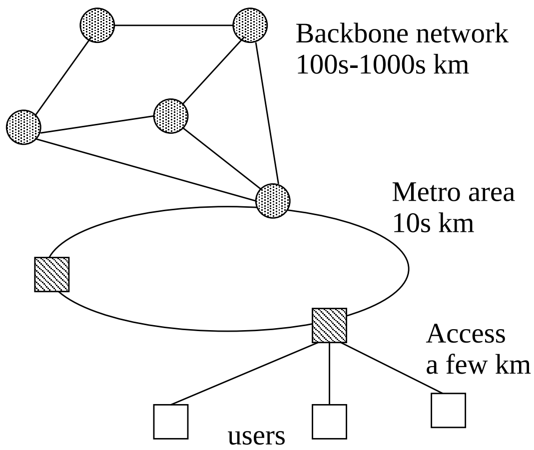

Figure 2.1: Today's network structure in backbone, metro and access areas.

The access network connects the subscriber (home or business) to the service providers; in other words, it serves as the *last mile* (as well as the *first mile*) of the information flow. It is a general opinion that the last mile is becoming a bottleneck in today's network infrastructure [23]. Optical technology can solve the bandwidth problem in access networks since it can provide at least 10 to 100 times more bandwidth over a larger coverage area, compared to its counterpart technologies. The next wave in access network deployment will bring the fibre to the building (FTTB) or to the home (FTTH), enabling Gb/s (gigabit per second) speeds at costs comparable to other technologies.

Metropolitan-area (or metro) networks serve geographic regions spanning several hundred kilometers, typically covering large metropolitan areas. They interconnect access networks to backbone service providers. Currently, the physical layer infrastructure in metro networks is formed by SONET/SDH-based rings.

The backbone network has optical network nodes interconnected by a mesh of fibre links. Traffic from the end users (which could be an aggregate activity from a collection of terminals) is collected by the access networks and fed into the backbone networks through metro networks. This high-bandwidth traffic is carried on a backbone network from one end to the other by optical communication channels on fibres. Design and management of optical backbone networks has been an exciting research area comprising several issues. Since the scope of this thesis is backbone networks, the rest of the chapter will cover several aspects of optical communication with a point of view of the backbone networks.

2.1.2 Wavelength Division Multiplexing

A fibre's potential bandwidth (50 Tb/s) is nearly three order of magnitude higher than electronic data rates of a few tens of gigabit per second. Because the maximum rate at which an end user -which can be a workstation or a gateway that interfaces with lower speed subnetworks- can access the network is limited by electronic speed, to exploit the fibre's huge bandwidth concurrency among multiple transmission channels should be introduced.

Wavelength-division multiplexing (WDM) is a favourite multiplexing technology in optical communication networks because it supports a cost-effective method to provide concurrency among multiple transmissions in the wavelength domain. Several communication channels, each carried by a different wavelength, are multiplexed into a single fibre strand at one end and demultiplexed at the other end, thereby enabling multiple simultaneous transmissions. Each communication channel (wavelength) can operate at chosen electronic processing speeds, e.g., OC-192[1] (i.e. 10 Gb/s) or OC-768 (40 Gb/s). For example, a fibre strand that supports 160 communication channels, i.e., 160 wavelengths, each operating at 40 Gb/s, would yield an aggregate capacity of 6.4 Tb/s.

[1]OC-n is an optical channel with data rate of $n \times 51.84$ Mb/s approximately.

2.2 Optical Technology in Backbone Networks

This section provides basic knowledge on the physical layer technology in optical networks. It includes the fibre, optical transceivers, amplifier, wavelength converter and switches.

2.2.1 Optical Media: Fibre

Optical fibre is the physical medium for transporting the signals in optical networks. The transmission geometry is explained in the following subsection, and then the properties of the fibre, and transmission impairments are summarised.

Transmission of Optical Signal

Optical fibres are made of silica. A typical cross section of a fibre is shown in Fig. 2.2. The light travels through the core, which is surrounded by the cladding to keep the rays inside the core. The size of the optical fiber core determines how the light travels through it. Each optical signal can actually generate many different lightwaves, which can all travel through the fiber at the same time. This is allowed in *multimode* fibres, but can cause problems with each wave arriving at the end of the fiber slightly out of synchronisation. Most modern optical networks use *singlemode* fibre, which has a much smaller core than that of multimode. The core size is small enough to ensure that only one lightwave from each optical signal can travel in the fiber, so there are no problems at the receiving end. To give an idea of sizes, in a singlemode fiber the core is usually around 10 micrometers in diameter, and the cladding is over 10 times thicker at a diameter of 125 micrometers. Once the polymer coatings are added for protection, the whole package may be around 0.25 millimeters in diameter.

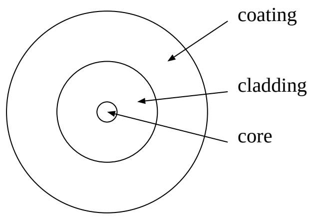

Figure 2.2: Cross section of a fibre.

The optical signals are carried by fibres thousands of kilometers. The principle of the propagation of the signals can be explained by simple optics. Figure 2.3 illustrates the reflection principle on a boundry between two material with different refractive index. Refractive index is the ratio of the speed of light in free space to the speed of light in the medium. The refractive index of the Medium 1 is $n1$ and the refractive index of the Medium 2 is $n2$. Relations between the angle of incidence Θ_i, the angle of reflection Θ_r, and the angle of transmitted ray Θ_t are:

$$\Theta_r = \Theta_i \tag{2.1}$$

$$n_1 \cdot \sin(\Theta_i) = n_2 \cdot \sin(\Theta_t) \tag{2.2}$$

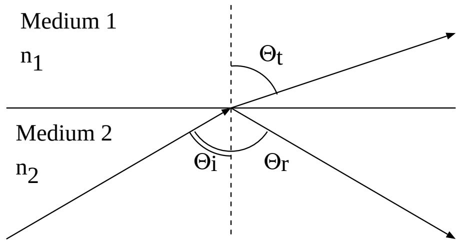

Figure 2.3: Refraction and reflection of a ray.

Equation 2.2 is called Snell's law, and it describes the refraction angle of a ray that crosses the boundary of two different media. There is a critical angle Θ_c for the incident ray, where the transmitted ray lies on the boundary:

$$\Theta_c = \sin^{-1}\left(\frac{n_2}{n_1}\right) \tag{2.3}$$

At angles larger than Θ_c the energy of the incident ray is totally reflected resulting in a *guided ray*. Guided rays are used to carry the optical signals.

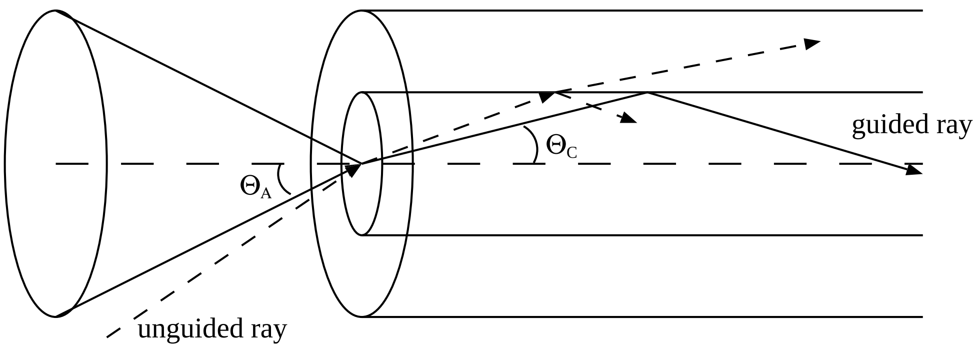

Figure 2.4: Propagation of guided rays in fibre.

The core of a fibre has a refractive index n_1, and the surrounding cladding has a slightly lower refractive index n_2 ($n_2 < n_1$). Figure 2.4 illustrates the propagation of rays in a fibre. A ray that enters the fibre at a sufficiently small angle of incidence propagates through the fibre bouncing back and forth between the walls of the core. An unguided ray that enters the fibre at a larger angle is partially refracted at each bounce. The largest angle for a guided ray is the *acceptance angle* Θ_A, and can be calculated according to the refractive index of the environment. If the ray is entering the fibre from air the acceptance angle is:

$$\Theta_A = \sin^{-1}(\sqrt{n_1^2 - n_2^2}) \tag{2.4}$$

Optical Fibres

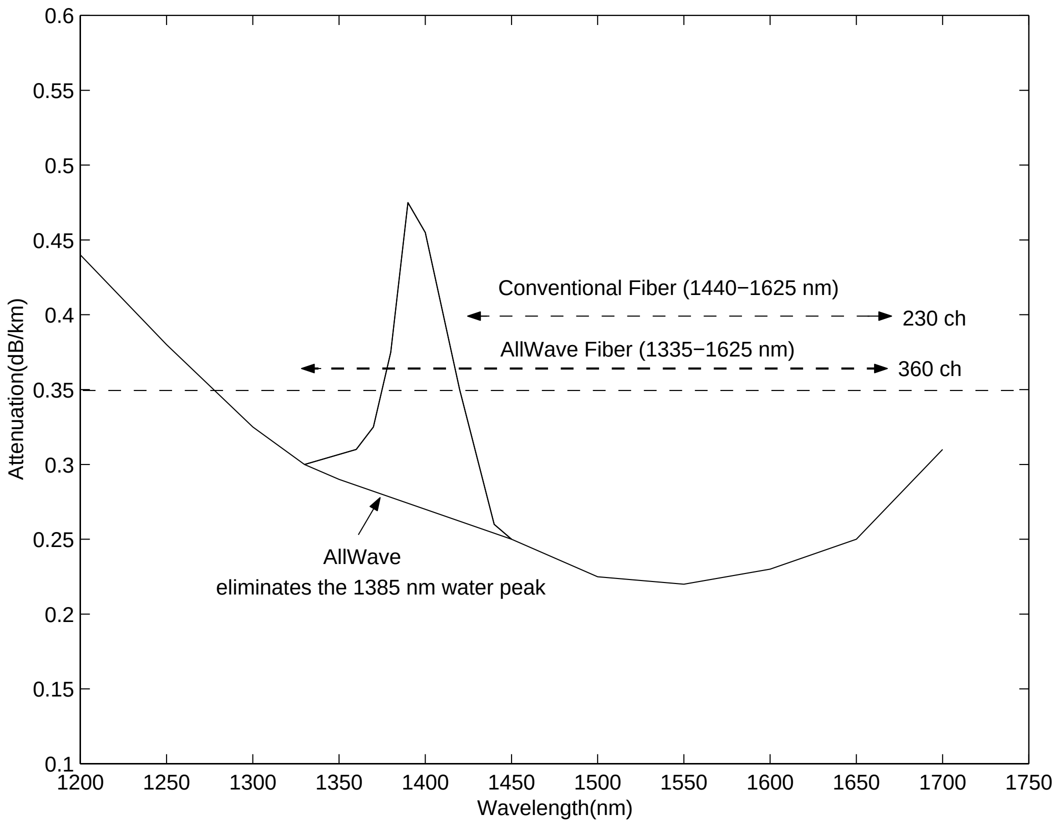

Figure 2.5: The low-attenuation regions of a conventional and all-wave fibres.

In WDM, the bandwidth of a fibre is divided into several low-bandwidth channels, each operating at a different wavelength. The number of channels that can exist in a single fibre is limited primarily by the total available bandwidth (or spectral range) of the fibre and spacing between channels [22, 24]. Conventional fibres have low-attenuation region between 1335 nm - 1625 nm with a *water-peak window* at 1385 nm. New *all-wave* fibres do not have this water-peak and hence can use a larger spectrum as shown in Fig. 2.5. Channel spacing is affected by several factors such as channel bit rates, nonlinearities in the fiber, and the resolution of the optical components. The wavelength channels must be spaced sufficiently far apart to keep neighboring signal spectra from overlapping, and consequently from interchannel cross-talk. In dense wavelength-division multiplexing (DWDM) a large number of wavelengths (above 160) is packed into the fibre. An alternative WDM technology with a smaller number of wavelengths (less than 10), large channel width, large channel spacing is termed as coarse WDM (CWDM).

Transmission Impairments

If we shine a light through a glass rod, very small part of the light would reach the far end. The reason is that, no matter how clear the glass appears, it is full of impurities that absorbes the light. The optical fibre is made from extremely pure silica, but a part of light is still lost on its way through the fibre.

Signal power has losses through the fibre due to material absorption, Rayleigh scattering, and waveguide imperfections. Material absorption occurs because of the resonance of the silicon molecules at short and long wavelengths and impurities in the fibre, such as OH ions effect (from traces of water in the fibre). Since the refractive index of the fibre is not absolutely uniform, the light is scattered attenuating the propagating wave. Waveguide imperfections are caused by non-ideal fibre geometries, bends and distortions in the medium. Fibre loss is greatly compensated with the introduction of erbium-doped fibre amplifiers (EDFA) in 1989.

A narrow beam that enters the fibre tends to broaden as it progress in the medium. This effect is called *dispersion*. Chromatic dispersion and polarisation mode dispersion are two types of this effect (for multi-mode fibres there is also intermodal dispersion). *Chromatic dispersion* is due to different velocities of each spectral component in a signal. Asymmetries in the fibre causes also different spectral components travel at different velocities, known as *polarisation mode dispersion*.

The fibre has also the following nonlinear effects:

- Stimulated Raman scattering

- Stimulated Brillouin scattering
- Four-wave mixing
- Self- and cross-phase modulation

In *Stimulated Raman scattering* the light interacts with the medium producing collisions with the molecules. The result of the process is energy loss. In *Stimulated Brillouin scattering* the power lost in the process is transferred to an acoustic wave. *Four-Wave Mixing* (FWM) occurs if three signals are present at neighboring optical frequencies. FWM causes cross-talk between channels and can accumulate along a long fiber line. Refractive index nonlinearities can also produce phase shifts that are accumulated over a length of a fibre, known as *self-phase modulation*. When more than one signal is present, the nonlinear interactions between the signals produce a related effect *cross-phase modulation*.

2.2.2 Signal Amplifiers

The optical signal loses its energy because of the transmission impairments in a fibre. Long distance optical transmission is possible only by signal amplifiers that provide a power boost to the signal. The three types of amplifiers are Semiconductor Optical Amplifiers (SOA), Erbium-Doped Fibre Ampliers (EDFA), and Raman Amplifiers. Among them EDFAs can be used for long distance links, such as submarine fibres, placing them every 80 km along the link. Each of the type has specific spectral range and amplification gain. The gain of EDFA is between 25-51 dB, and its spectral range is 35 nm around 1550 nm.

Optical amplifiers can be placed in three different positions in a link. Power amplifier is placed at the beginning of a link to give a power boost to the clear signal. Pre-amplifier is placed before the receiver, and the line amplifier is placed in the fibre. A recent technology uses a circuit of EDFAs to fully exploit the spectrum of all-wave fibre and it is called *ultrawide-band* EDFA[3].

2.2.3 Optical Transmitter: Laser

The optical signal is produced by the optical transmitter which is a laser (light amplification by stimulated emission of radiation). The laser is a device that converts electrical energy to monochromatic light. The lasers used in optical networks are semiconductor lasers, and they are very small (size of a grain of salt). They are designed specifically to give a precise and intense light, since in a WDM network it is very important to obtain the light at a precise wavelength. Different materials are used to obtain different wavelengths from the laser. In optical networks the wavelength range is in infrared region, between 1300 nm - 1550 nm.

A relatively young technology permits producing tunable lasers. These are light sources that can be adjusted to emit light at different wavelengths. They are also semiconductor-based and their design principle is similar to non-tunable lasers. Current products can cover around four wavelengths in their tuning range. Some new technologies allow aroud 40 nm range that can cover up to 50 wavelengths.

2.2.4 Optical Receiver

Optical receivers detect incoming lightwave signals and convert them to an appropriate signal for processing by the receiving device. The optical signal is converted into photocurrent by photodetectors which are semiconductor photodiodes. The carried signal is detected from the photocurrent by detection circuit. The signal is processed for clock recovery, sampling and threshold detection to extract the digital bitstream from the received signal. As tunable transmitters, there are also tunable receivers which range can be as large as 500 nm.

2.2.5 Wavelength Converter

A wavelength converter converts an optical signal at wavelength λ_i to a signal at wavelength λ_j. The conversion may take place totally in optical domain, as in difference frequency converter (DFC). Another method to build a converter is to use optical-to-electronic then electronic-to-optical conversions, by reproducing the optical signal. Wavelength converters may be used in optical switches to facilitate the switching and prevent blocking due to wavelength unavailability at the output fiber.

2.2.6 Optical Switches

In a mesh structured backbone network, the nodes are equipped by switches that route the incoming data to their destination nodes. Statistically aroud 80% of the traffic arriving to a node has a destination different from that node. In a packet switched network, each packet arriving to a node has to be processed by the electronic switches to be routed to their proper destination. The electronic switches used in backbones have to work at very high data rates to accomodate all traffic from all its incoming transmission lines.

WDM technology provided a very high transmission bandwidth within a single fibre (nearly 50 Tb/s). A backbone node connected to several other nodes potentially may have hundreds of incident fibres. Using the electronic packet switching method in such a network is not practical because of the processing load and today's electronic processing speed. It is also very inefficient since most of the traffic is transit.

Optical switching technology provides a natural solution for optical mesh networks to provide an efficient and flexible traffic routing. In a typical optical switch, each wavelength channel can be routed without data level processing. Only the traffic arriving to that node need to be converted into electronic domain. This section covers basic switching technology and switch types briefly.

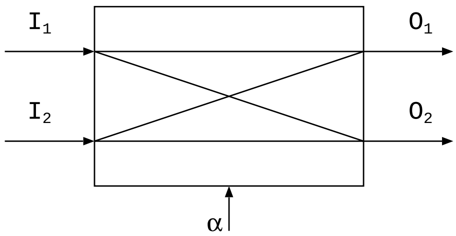

Figure 2.6: A 2×2 directional coupler.

The simplest optical switch is a 2×2 optical coupler shown in Fig. 2.6, and it can be used as a building block of larger switches. The optical signals enter the coupler from the two input ports, I_1 and I_2. They are combined and divided linearly according to the control signal α and the output signals leave the coupler from the two output ports, O_1 and O_2. When it is used as a 2×2 switch, the coupler may have two states: bar and cross. The control signal α becomes a digital signal defining one of the two states. In bar state, the signal entering from port I_1 is directed to the output port O_1, and the signal entering from port I_2 is directed to the output port O_2. In cross state, the signal entering from port I_1 is directed to the output port O_2, and the signal entering from port I_2 is directed to the output port O_1.

Classification

Optical switches may have many different structures and properties, and they can be classified in many ways accordingly. One classification is related to the transparency of the switch. If an optical signal travels through the switch without being converted into electronic domain, the switch is called to be *transparent*. Transparent switches are one of the building blocks of all-optical networks. If a transit optical signal is converted to electronic and back to optical domain, to be routed in an optical switch, the switch is called to be *opaque*.

Switches can be classified also according to their routing abilities. Some simple architectures have *static* routing patterns, and they are called static switches. Among them are star coupler and arrayed waveguide gratings. Other switches can be configured to route any wavelength channel coming from an input port to any output port. These switches are called *dynamic* switches. Some important dynamic switches will be studied in the following paragraphs, since they are of interest for this study.

Crossbar Switches

One of the simplest switch structures is the crossbar that can be seen in Fig. 2.7. An $n \times m$ crossbar switch has n input lines, m output lines, and $n \cdot m$ cross points. Each cross point has an optical coupler that can be in bar or cross state. Only one input can be connected to an output, hence such a switch architecture allows only point-to-point connections. A drawback of cross-bar switches is that $n \cdot m$ optical couplers are needed in total to build the fabric. For large size switches, multi-stage architectures are more popular since they can be realised with fewer cross points.

Multi-stage Switches

The 2×2 optical coupler can be used to build $n \times n$ multi-stage switch fabrics. The structure of an optical switch fabric is similar to those in electronic switches (e.g. ATM).

An 8×8 Benes switch fabric is shown in Fig. 2.8. Larger switches can be built using the same recursive pattern, and a general $n \times n$ Benes fabric can be built using $n \log_2 n - n/2$ optical couplers. Another popular switch fabric is Clos network which is composed of three stages.

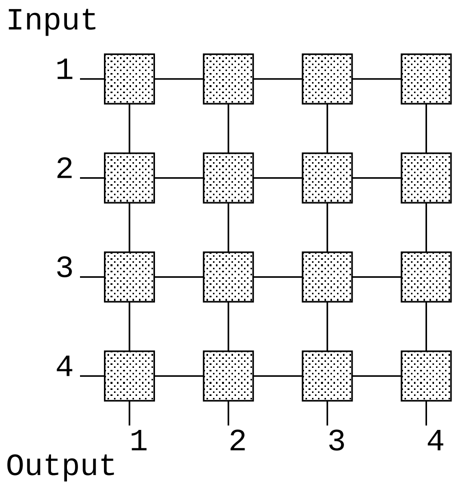

Figure 2.7: A 4×4 optical crossbar switch.

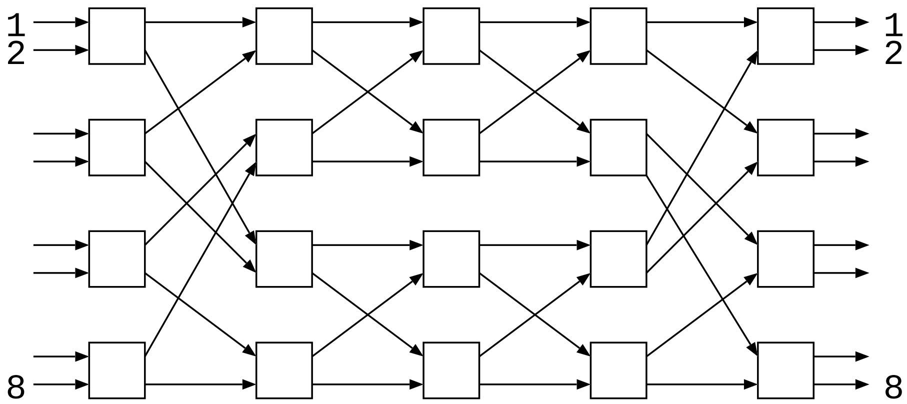

Figure 2.8: An 8×8 Benes switch.

MEMS Switches

Among several other technologies (e.g., bubble, liquid crystal, thermo-optic, holographic, etc.) used for building optical switches, MEMS (micro-electro -mechanical systems) are becoming popular because of their compact design, low power consumption, and promise for high port count [25, 26]. Two architectures for MEMS-based switches have emerged. The first architecture, known as 2D switch, is a square array of $N \times N$ mirrors (each mirror having a diameter on the order of 1 mm) used to couple light from a linear array of N fibres on an adjacent side of the square. The (i, j) mirror is raised up to direct light from the i^{th} input fibre to the j^{th} output fibre. Mirror control for these switches is digital but the optical power loss grows rapidly due to path length inside the switch. The second architecture, known as 3D switch, makes use of the three dimensional space. In this architecture, there is a dedicated movable mirror for each input and output port. Each mirror operates in analog mode rather than binary, tilting freely about two axes. Port counts of several thousands are achievable in this architecture.

Structure of an Optical Cross-Connect (OXC)

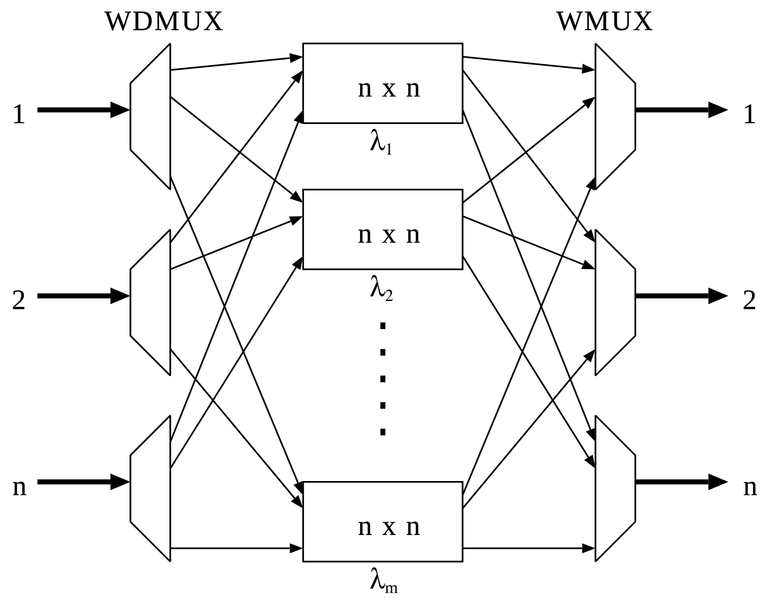

Figure 2.9: An optical cross-connect with n input and output fibres, and m wavelengths.

A typical $n \times n$ optical cross-connect is shown in Fig. 2.9. The n input fibers, each carrying optical signals in m wavelengths ($\lambda_1, \lambda_2, ..., \lambda_m$), are connected to wavelength demultiplexers (WDMUX) which spatially separate the wavelengths on the fibers. Same wavelengths from all input fibers are sent to a wavelength specific $n \times n$ switch that routes the input signals to appropriate output ports. Hence, the optical cross-connect has m $n \times n$ optical switches, one for each wavelength. Once the signals are routed properly to their output ports, they are combined into a fibre by wavelength multiplexers (WMUX) situated at each output port.

Switch Properties

An important property for a switch is its blocking characteristic. An $n \times n$ switch may be required to realise $n!$ different connection states (permutations of n). Such a switch is *rearrangeably nonblocking* if there exists a set of paths through the switch fabric that can realise each of the $n!$ connection states. The term rearrangeable comes from the fact that it may be necessary to rearrange ongoing connections in the switch, to route a new connection between a pair of idle input and output ports.

Wide-sense nonblocking switches can realise any connection state without rearranging active connections, provided that the correct rule is used to route each new connection. *Strict-sense nonblocking* switches do not require any rearrangement or complex routing rules, and new connections can use any free path in the switch. As might be expected, the complexity of the hardware grows according to the nonblocking property.

Among the several properties that can be added to a switch structure, wavelength conversion provides flexibility in routing. Wavelength converting switches allow a wavelength channel λ_i from input port p_{in} to be routed to a channel λ_j at output port p_{out}. Wavelength conversion capability of a switch can be partial or full according to the design of the fabric. Opaque switches have full wavelength conversion property since they regenerate the optical signal every time it enters the switch.

2.3 Terminology of Optical Networks

In a wavelength-routed WDM network, end users communicate with one another via end-to-end WDM channels, which are referred to as *lightpaths* [13]. A lightpath may span several fibre links and optical cross-connects from its source node to the destination. It provides an optical channel, virtually one-hop length in optical domain, to the connection it supports. Figure 2.10 shows several node pairs in a network communicating through lightpaths, e.g., a lightpath from A to F spans accross the physical links AD, DE, and EF. A lightpath is a unidirectional channel, and a bidirectional (or directionless) path shown in the figure denotes two lightpaths between the end nodes in both directions. If the network nodes do not have wavelength converters, a lightpath must occupy the same wavelength on the fiber links through which it traverses; this property is known as the *wavelength continuity constraint*. All lightpaths shown in Fig. 2.10 obey this constraint, e.g., lightpath between C and G occupies λ_2 along the fibres CE, ED, and DG.

Assuming that only two wavelengths are available, and wavelength conversion is not allowed in the network shown in Fig. 2.10, a new ligthpath from B to E cannot be established due to wavelength continuity constraint. Wavelength λ_1 is used on link BC, hence a lightpath starting at node B going toward node C can use only

λ_2 on that link. However, the only available wavelength is λ_1 on link CE, where λ_2 is occupied by lightpath CG. If wavelength conversion was possible at node C, the new lightpath could be established using λ_2 on BC, and λ_1 on CE. As seen in this example, routing and assigning wavelength to a lightpath is a complex issue. This problem is known as *routing and wavelength assignment* (RWA). Given a set of connection requests, the problem of setting up lightpaths is known as *static lightpath establishment*. In dynamic traffic case, a lightpath is set up for each connection request as it arrives, and the ligthpath is released after some finite amount of connection time. The *dynamic lightpath establishment* problem involves routing the ligthpaths and assigning wavelengths to them while minimising the connection blocking probability.

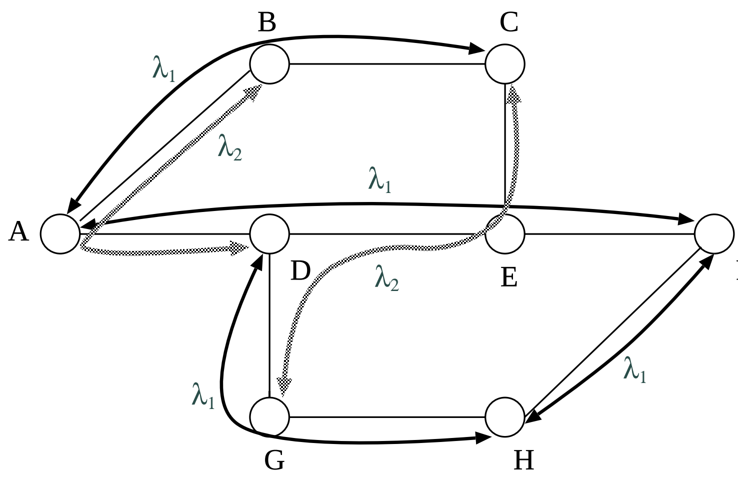

Figure 2.10: An example optical network with lightpaths established in two wavelengths.

2.3.1 IP-over-WDM

Internet is emerging as the dominant data network, and it is becoming clear that Internet Protocol (IP) will be the common traffic layer for communication networks [4]. New network applications appear, becoming more and more bandwidth-hungry, such as real-time video and multimedia applications, or online games. Network operators foresee the exponentially growing bandwidth demand, more fibre cables are deployed, and the standard bodies [7, 27] are working on putting together a high-speed Internet that can accomodate future's traffic demand. Fibre is coming closer to the end users to overcome the first mile [23] bottleneck, and it is fair to expect that the first mile also will be covered with fibre as optical technology becomes cheaper. When this happens, it will encourage, no doubt, the end user to use even more bandwidth.

On the transport side, WDM is seen as the most promising technology that can support such bandwidth growth. Optical technology is developing rapidly, yet it is still considered to be at its infancy, comparable to the technological level of computer industry in 1940s [28]. All-optical technology is expected to improve in speed, channel number, and gain flexibility with the advances in optical component technology to provide an intelligent transport network.

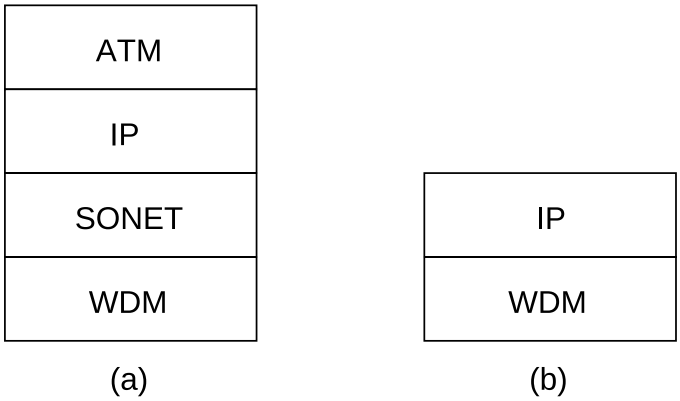

Figure 2.11: (a) Today's network model in WDM backbone. (b) Simplified two-layer model.

Today's most long-distance IP network architectures are based on synchronous optical network/synchronous digital hierarchy (SONET/SDH) which encapsulate IP packets (or ATM cells that carry IP packets) in SONET/SDH frames (Figure 2.11a). This multi-layered structure is highly inefficient and increases the operational cost, therefore networking world converges toward a simpler two-layer model (Figure 2.11b) called IP-over-WDM.

In an IP-over-WDM architecture, routers at each node are connected to the ports of optical crossconnects (Figure 2.12). The connectivity of the routers can be arranged by appropriately configuring the OXCs. The IP layer sees the optical layer as a reconfigurable topology, since the neighboring relation of a router can be changed by reconfiguring the OXCs. The IP traffic between a source and a destination can be carried by lightpaths established in optical layer. Nodes that are not connected directly by lightpaths can still communicate with one another using the multi-hop approach, i.e., by using electronic packet switching at the intermediate nodes. Interaction between the optical layer and the electronic layer (IP in this case) is a major issue including several functions, such as bandwidth provisioning, fault management, performance monitoring.

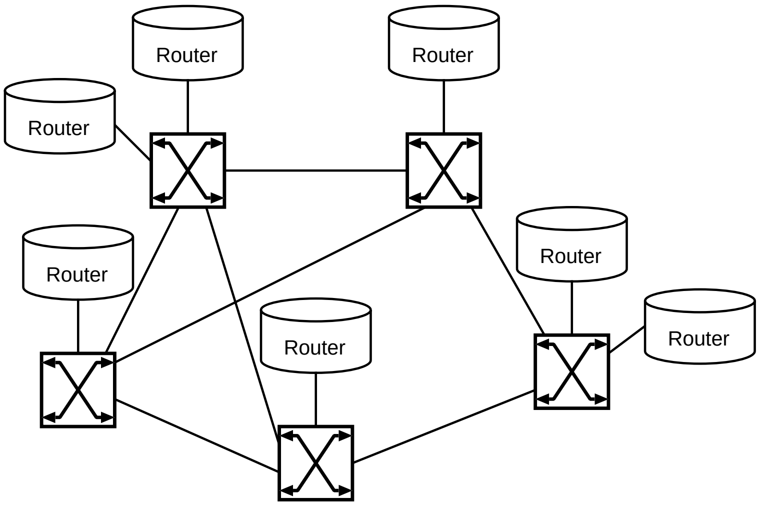

Figure 2.12: IP-over-WDM network architecture.

2.4 State of the Art

Optical transmission dates back to 1970s when low-loss fibres were produced and semiconductor lasers were invented [22]. Transmission speed and distance grew at exponential rate until late 1980s, and the first trans-Atlantic optical fibre cable, using electronic repeaters, was laid in 1988. The invention of EDFA in late 1980s made long distance optical transmission possible.

Interest in optical networks began in mid 1980s, and the first laboratory prototype, Lambdanet, appeared in 1990 [29]. The first deployed network was Rainbow-1 (1991), and the first commercial WDM products appeared in 1995. The technological advances were leading the improvements of the last decade: Development of all-optical products, more channel per fibre, more bitrate per channel, longer transmission range, new kinds of fibres and amplifiers, lower cost.

Research on optical network architectures and protocols started to appear in 1990, and since then optical networks became an increasingly popular area. This section focuses on past research on wavelength-routed WDM networks, putting emphasis on virtual topology design and reconfiguration studies.

2.4.1 Studies on Optical Networks

The studies on optical networks may be classified according the scale of the network considered, and Figure 2.13 shows the several research topics in optical networking field.

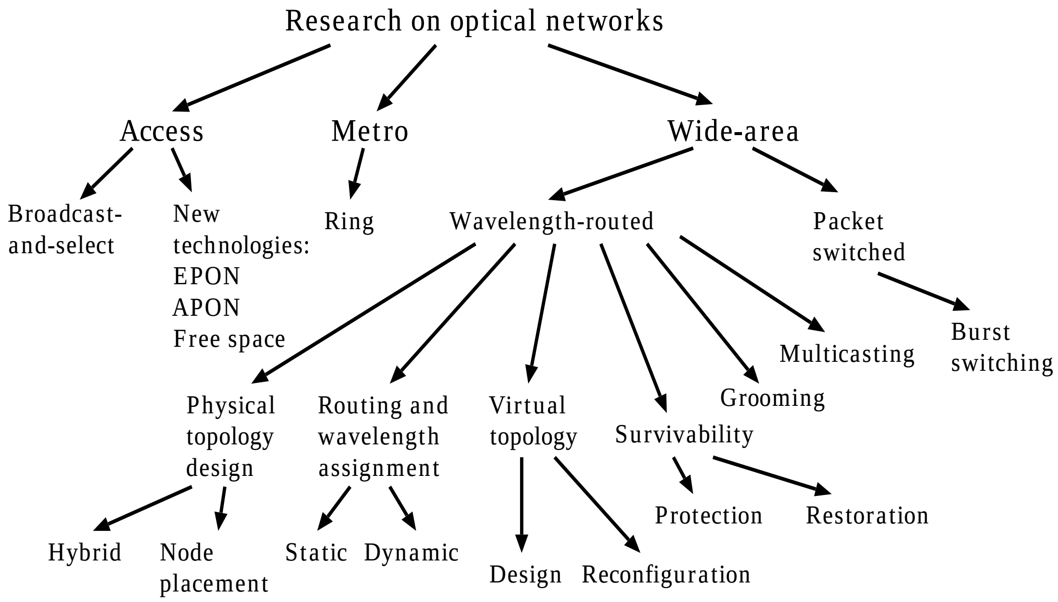

Figure 2.13: Classification of optical network studies.

Research on access networks dates back to early stages of optical networks, starting at late 1980s. Single-hop and multi-hop systems were developed, and different networks architectures were proposed. A survey of the different local lightwave networks is given in [30, 31]. Regular and irregular topology structures including shufflenet, deBruijn, Manhattan Street, or hypercube networks, topology optimisation formulations are examined; time-slot scheduling, dedicated channel models are developed through these studies [32, 33, 34, 35, 36, 37].

There are three optical technologies that are promising candidates for the next-generation access networks: Point-to-point topologies, passive optical networks, and free-space optics. In point-to-point topologies dedicated fibre links connect each subscriber to the central office. This architecture is simple but but expensive due to the extensive fibre deployment. An alternative approach is to use an active star topology where a curb switch is placed close to the subscribers to multiplex signals between the subscribers and the central office. Passive optical networks (PONs) replace the curb switch with a passive optical component such as an optical splitter. Several topologies are suitable with PONs including tree-and-branch, ring, and bus [23, 38]. A third option is to use infrared lasers to transmit data via point-to-point topologies. An optical data connection can be established through the air via lasers sitting on rooftops aimed at a receiver [39, 40].

Metropolitan-area networks interconnect access networks to backbone service providers, and their physical infrastructure is formed by SONET/SDH based rings. SONET rings utilise a single channel with a TDM technique. With TDM, a high bandwidth channel can be divided into several low bandwidth sub-channels and each sub-channel can carry a different low rate traffic streams. With the emergence of WDM technology, a logical step is to upgrade the one-channel SONET ring to a multiple-channel WDM/SONET ring, where each wavelength can operate similar to a SONET TDM channel. However in a simple-minded solution a SONET add/drop multiplexer is needed (ADM) for each wavelength at each node, increasing the total number of ADMs in a network W times where W is the number of wavelengths. Fortunately, it may be possible to have some nodes on some wavelengths where no add/drop operation is needed on any time slot. Tho total number of ADMs can be reduced by carefully packing the low-bandwidth connections into wavelengths. Packing low-speed traffic streams into high-speed traffic streams to minimise the resource usage is called *traffic grooming*, and it is a research topic that has received a lot of attention [41, 42, 43, 44, 45, 46].

2.4.2 Studies on Wide-area Optical Networks

Wide-area WDM networks have been a popular research area in the past decade, and many algorithms and methods have been proposed to solve the design, routing, resource allocation, reconfiguration, survivability, grooming, multicasting, and packet switching problems. The studies related to routing and wavelength assignment (RWA) and virtual topology are examined in detail in the next section. Here, the other wide-area networking studies (Figure 2.13) are summarised.

Physical layer properties of a network affects the design of its higher layers in terms of cost, efficiency, and survivability. Many studies consider the physical layer design with different aspects: Topology planning [47], Node placement [48], routing based on physical layer properties and impairments [49], routing based on duct layer [50], etc.

In a WDM network the failure of a network element, e.g., fibre link, crossconnect, may result in the failure of several optical channels, thereby leading to a large data and revenue losses. There are several approaches to ensure network survivability against fibre-link failures [51]. Survivable network architectures are based either on reserving backup resources in advance called *protection* [52], or on discovering spare backup resources in an online manner called *restoration* [53]. In protection the disrupted service is restored by utilising the precomputed and reserved network resources. In restoration, the spare capacity, if any, available within the network is utilised for restoring services affected by the failure. Design of survivable optical networks is a major topic that has been studied thoroughly [54, 55, 56, 57, 58, 59, 50].

The traffic grooming problem has been studied for WDM/SONET ring networks with the objective of minimising the total network cost measured in terms of the number of SONET ADMs. As today's optical backbone networks are evolving from interconnected rings to mesh topology, traffic grooming in WDM mesh networks has become an important problem [60, 61, 62, 63]. In order to support traffic grooming, each node in a WDM mesh network is equipped with an OXC which is able to switch traffic at wavelength granularity as well as finer granularity. A grooming fabric in the OXC performs multiplexing, demultiplexing, and switching of low speed traffic streams.

In a static grooming problem, all connection requests of different bandwidth granularities are known a priori. In a dynamic grooming problem, connection requests arrive randomly, hold for a finite duration, and require provisioning in real time and tearing down when they are over. The grooming problem in wide-area networks is an interesting and challenging problem that need and attract growing attention [64, 65, 66].

Another interesting problem is providing multicast connections in optical networks. A *light-tree* [67] is a point-to-multipoint generalisation of a lightpath and provides single-hop communication between a source and a set of destinations. To realise multicast-capable WDM wide-area networks, multicast-capable switch architectures and effective routing algorithms need to be developed [68, 69].

An alternative to the wavelength-routed networks is optical packet switching. The optical packet switching (OPS) technology is in a very early stage, and attracts growing attention by the researchers. In the long term OPS could become a viable candidate because of its high speed, fine-granularity switching, flexibility, and its ability to use the resources economically [9, 10, 70, 11, 71]. As a midway solution between circuit switching and packet switching, optical burst switching was proposed [72, 73].

Since the focus of this study is virtual topologies, the following section examines the past research on this fiels in detail.

2.4.3 Survey of Virtual Topology Design and Reconfiguration Studies

The major studies related to the subject of this thesis can be categorised in three fields:

- Routing and wavelength assignment (RWA)
- Virtual topology design
- Virtual topology reconfiguration

There are two different types of the RWA problem depending on the nature of the traffic in a network: static and dynamic. Typically, connection requests in a network are either static, where a traffic matrix is given and does not change after the connections are set up; or dynamic, where connection requests arrive randomly in time, and they are required to be set up one by one. The static RWA problem can be formulated mathematically as an optimisation problem (integer linear program – ILP) [74, 75, 76]. Solving this formulation is a very complex problem, because the search space grows very quickly with the network parameters such as number of nodes and number of wavelengths. In fact, RWA is an NP-complete problem, hence the studies on this topic concentrate on relaxation techniques for solving the ILP [74], or alternative heuristic algorithms [75, 76]. Usually, RWA problem is decomposed into sub-problems: Routing sub-problem and wavelength assignment sub-problem. This technique leads to sub-optimal solutions, but it simplifies the structure of the problem and algorithms to be used. Some well-known routing algorithms and wavelength assignment algorithms are given in Appendix A.

In dynamic RWA problem, connection requests arrive one by one and the lightpaths are set up dynamically, according to the actual state of the network resources. Similar to the static case, the two sub-problems can be considered separately by combining different routing schemes and wavelength assignment schemes in tandem.

Several studies consider the RWA problem simultaneously to obtain better blocking performance [77, 78, 79]. A thorough review of routing and wavelength assignment approaches can be found in [80].

Virtual topology design is the problem of selecting and setting up (by routing and assigning wavelength) a group of lightpaths, when a traffic request matrix is given. The objective of the design may be maximising the packet traffic, balancing the lightpath loads, minimising some of the network resources or minimising the average packet delay, etc. This problem contains several sub-problems [12]:

1. Determining a good virtual topology, i.e., a good set of lightpaths
2. Routing the lightpaths over the physical topology
3. Assigning wavelengths optimally to the lightpaths
4. Routing packet traffic on the virtual topology

These sub-problems are not independent, hence, solving them separately may lead to sub-optimal solutions.

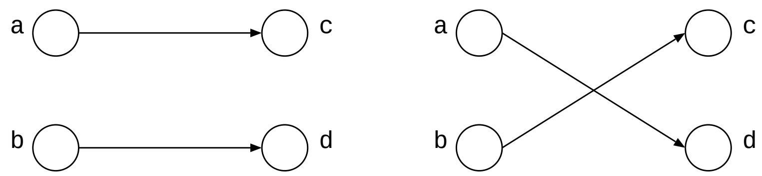

Figure 2.14: Branch-exchange operation.

Virtual topology design was often formulated as a Mixed-Integer Linear Program (MILP) [14, 81, 82, 17, 83]. Since solving these various formulations are not tractable for networks having more than few nodes, several heuristic algorithms were proposed. In [14] a two-step iterative algorithm is presented. This method starts with a heuristic initial solution, and then performs a local search by applying branch-exchange operations (Figure 2.14) on the least utilised links of the connectivity graph. The objective of the formulation in [81] is to minimise the network congestion while constraining the average delay seen by a source-destination pair. The authors propose to split the problem into sub-problems to solve it for large networks. They propose several heuristic algorithms, one that aims maximising one-hop traffic, one minimising the packet delay, one being a traffic independent design, and a method based on LP-relaxation (linear program relaxation). Another ILP formulation is given in [82] that accommodate many physical connectivity constraints. The authors concentrate on the design of the virtual connectivity (sub-problem 1) and routing of packet traffic (sub-problem 4), and they provide heuristic solutions based on simulated annealing and flow deviation. An exact ILP formulation for the complete virtual topology design is given in [17] with the objective of minimising average packet hop distance. The authors propose a method that terminate the optimisation within the first few iterations of the branch-and-bound method to provide high-quality solutions. In [83] a general ILP is presented, which takes into account the maximum number of hops a lightpath is permitted to take, multiple logical links in the virtual topology, multiple physical links in the physical topology, and symmetry restrictions. The authors also propose a virtual topology design method based on rounding the solutions obtained by solving the relaxed ILP problem. A comparative study of virtual topology design algorithms can be found in [84, 85].

Reconfigurability of optical layer is one of the most valued features of WDM networks. Reconfiguration provides flexibility against the changing network conditions, such as traffic intensities, node and fibre failures. Rearranging or reconfiguring a virtual topology means designing the connections to be set up given the network conditions, by considering the previously established connections. An important point is to design the new topology as close as possible to the previous (existing) one, to minimise the disruption of the traffic being routed.

Studies on virtual topology reconfiguration generally have a two-step approach. First step is to design a new topology considering several input information including the new traffic matrix, current virtual topology, and the primary objectives of the design (e.g. minimising the average hop distance, maximising network throughput). Second step is the transition phase where the optical switches, transmitters and receivers are reconfigured and the current virtual topology is changed to the new one. This approach is taken in [15], where the authors' aim is to provide a method for the second step, i.e., the transition problem. Once the new virtual topology is decided, the problem is to find the shortest sequence of operations to go from the initial to the target topology. This way, the duration of the transition can be minimised. The paper proposes three polynomial-time algorithms and compares their performance.

A reconfiguration methodology to adapt to changing traffic requests at the ATM layer is given in [16]. All traffic patterns are assumed to be known in advance at the design phase, and the virtual topologies are pre-designed according to those changes. The reconfiguration of the optical layer occurs in a cycle, $T_0, T_1, ..., T_{N-1}, T_0, ...$ where T_i is the i^{th} pre-designed virtual topology. Once the topologies are designed, their ordering and configuration patterns are fixed. Having this information, the aim is to provide hitless reconfiguration by appropriately designing the transition phase. For a hitless reconfiguration, first the new lightpaths existing in the new topology are established, and then the old lightpaths that do not exist in the new topology are torn down.

An ILP formulation is used to solve the reconfiguration problem in [17]. The method proposed computes a new virtual topology from an existing one, such that the new virtual topology is optimal with respect to the changing traffic patterns. Among all such virtual topologies, the algorithm selects the topology which is closest in structure to the current virtual topology. This method minimises the number of switch retunings during the transition period and the number of disrupted lightpaths. A modification of this method is proposed in [18] where the ILP optimises various metrics, number of hops number of physical links, number of lightpaths and various combinations of these variables. The aim of the paper is to include trade-offs between the amount of reconfiguration and the objective. The methods in both studies are applicable only to small-size networks because of their complexity.

Two recent studies provide a combination of a fixed virtual topology and dynamic lightpath set up when new connections are needed. In [86] a loose virtual topology is constructed and it remains static. Some wavelengths are reserved to cope with the dynamic traffic. In [87] the virtual topology has two logical layers. The cross-connect network is managed on long time scale, while the switched-path network is managed on short time scale. The latter allows dynamically adding or deleting lightpaths. In this method, at least one path is always kept between every pair of routers that communicate. Since routing and wavelength assignment problems are not considered, the method is applicable only to local area networks.

A reconfiguration method to reduce the maximum link load in the network is proposed in [88]. The paper first deals with ring virtual topologies, and it provides a three-branch-exchange method to reconfigure the topology step by step. For general topologies two-branch-exchange operation is used to converge to a desired topology. This study performs small changes to the virtual topology to minimise the traffic disruption. The selection of nodes for branch-exchange is based on the idea of selecting a node group that would result in best reduction in the maximum link load. This selection is made by exhaustive search, resulting in high time-complexity.

An on-line reconfiguration study is given in [19]. The authors propose a heuristic method based on a two-stage approach. First stage is reconfiguration where the changes to be made in virtual topology are determined with the aim of minimising the objective function value. At this stage the number of changes are limited within a pre-specified value. Second stage is optimisation aiming to bring the objective function value closer to the optimal point.

As a general assumption, reconfiguration studies are based on the idea that the decision of reconfiguration is sudden, triggered by an event, and consequently virtual topology change is an interrupted process. In this view a virtual topology V_1 designed for traffic matrix T_1 remains the same until the traffic changes to a matrix T_2, which in turn triggers the design of another virtual topology V_2. On the other hand, previous studies lack to design a proper triggering mechanisms, i.e., when and how to decide to reconfigure the virtual topology in a backbone network. Another problem with this two-step approach is that the future traffic is assumed to be known. In practise, a flexible reconfiguration method should deal with dynamic traffic, even when the future traffic is different than the expected matrix.

Since the reconfiguration is a complex problem, most of the methods proposed are useful for small-size networks, even though they are local search methods. For large networks, simple and effective algorithms are needed to provide on-line reconfiguration of the optical layer.

3. VIRTUAL TOPOLOGY ADAPTATION

This chapter introduces the new virtual topology adaptation method. After an introduction to the problem, the formal definition is given and the solution approach proposed in this thesis is formulated using linear programming method. This formulation is solved for an example network to compare the performance of the new adaptation method with an earlier optimal reconfiguration scheme.

3.1 Introduction

As today's networks evolve towards the IP-over-WDM (Internet Protocol over Wavelength-Division Multiplexing) networking paradigm, we need to invent new methods to design and efficiently operate them. An IP-over-WDM network consists of the WDM-aware nodes and WDM fiber-optic links connecting these nodes, possibly in an arbitrary mesh topology. The optical nodes in such a network include optical cross-connects (OXC), transmitters and receivers, and these nodes are capable of separately routing the various wavelength channels in the optical domain. Optical nodes can be configured to set-up lightpaths to provide single "electronic" hop communication channels between any two nodes, which may be geographically far apart in the physical network, and to eliminate extra signal processing at intermediate nodes along that path. (Section 2.3).

A virtual topology is defined to be the set of all lightpaths in a network. Such a virtual topology can be employed by an Internet Service Provider (ISP) or a large institutional user of bandwidth, hereafter referred to as ISP, to connect its end equipment (e.g., IP routers) by leasing bandwidths (wavelength channels) from the network operator who owns the fiber plant and OXCs. In fact, multiple virtual topologies, possibly belonging to different ISPs, may coexist on the same fiber plant. However, for the sake of clarity, this work focuses on only one such virtual topology and its adaptation in the rest of this chapter.

The aim of this work is to provide a new perspective, which is called in this thesis *virtual topology adaptation*, for reconfiguration of the virtual topology in WDM networks under dynamic traffic. This study defines the problem from a new point of view, different from the traditional two-step reconfiguration. The new approach sees the reconfiguration as a continuous process of observations and adjustments. This approach eliminates the need for traffic forecasts and critical triggering decisions, while providing robustness and speed.

Observations on real networks show that the traffic rates between node pairs fluctuate distinguishably over time [1, 89]. An examplary traffic measurement can be seen in Figure 3.1 [1]. In this example, the measurements for both directions of a link on the Abilene network are displayed as two profiles (black and gray) over a 33-hour period, starting at 9:00 AM on one day and ending a little after 6:00 PM the next day. An important characteristic in this example is that the difference between the traffic intensity at highest point (busy hours) and the traffic intensity at lowest point (non-busy hours) is remarkable, their ratio is 4:1. Thus, reconfiguration may be very beneficial in terms of economical use of the bandwidths. A second observations is that there are two types of variations in traffic intensity, short-term variations (typically seconds or minutes) and long-term variations (typically hours). In long-term variations, which are the reason of reconfiguration, the amount of traffic between nodes changes in a smooth and continuous manner. Therefore, a virtual topology adaptation mechanism that changes the topology slowly, can be beneficial for backbone networks that have similar traffic behaviour. The example in Figure 3.1 shows a typical snapshot from a backbone link, i.e., many different backbone links have similar traffic characteristics, although the shape of the profiles may change.

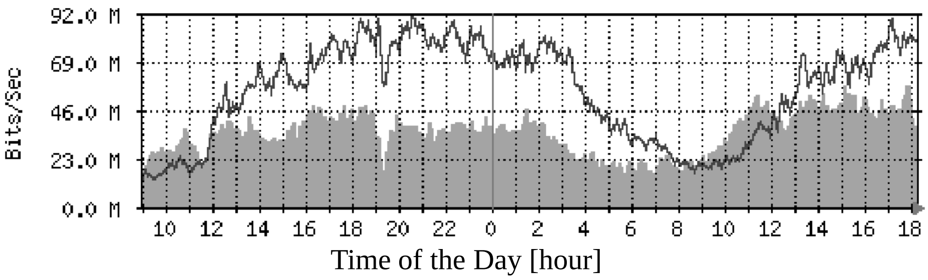

Figure 3.1: Traffic measurements on a link in the Abilene network [1] during a 33-hour period from 9:00 AM on Day 1 to 6:00 PM on Day 2. The two profiles correspond to the two directions of traffic on a link.

As summarised in Chapter 2, reconfiguration of optical networks has been studied extensively by researchers. In these studies, the problem is generally treated as a two-phase operation, where the first phase is virtual-topology design for the new traffic conditions and the second phase is the transition operation from the old virtual topology to the newly designed one. A major difficulty with this view is that the future traffic demand is assumed to be known. With this information in hand, the design of a new virtual topology is practicable, hence, in many studies considerable effort was spent for the design of the next virtual topology so as to reduce the topology changes during the transition phase [17, 19]. In practice, the assumption on the future traffic may decrease the value of the newly designed virtual topology if the traffic changes are predicted inaccurately. Although the daily fluctuations in traffic intensities have similarities in consecutive days, the traffic profiles are not the same, and they show considerable deviations from weekdays to weekends and holidays. Therefore, in this study, there is no assumption on future traffic, i.e., the reconfiguration method developed does not rely on traffic forecasts. This property makes the new approach robust and flexible.

The new approach can be viewed as a one-step reconfiguration approach in contrast to the previous methods. It is composed of continuous observation-adjustment cycles, where small changes may occur in the virtual topology. The traffic on each lightpath is observed continuously, and a new lightpath is added to or an existing lightpath is removed from the virtual topology when a change is necessary.

Regarding this new approach, we define the problem of reconfiguration differently: As traffic fluctuates over time, it will be monitored periodically, and the virtual topology will be changed accordingly. This new view of the problem is called virtual topology adaptation. The adaptation process is a continuous system where small adjustments are made, instead of waiting for a noticeable drop in system efficiency and changing the entire topology. With this new definition of the problem, the aim is to solve the issues of unpredictability of the traffic, critical triggering decisions, and the delay that is introduced by complicated reconfiguration algorithms. The ultimate goal is to obtain a scheme, which is practical for real backbone networks.

When the reconfiguration is performed on a working network, the changes made on the underlying virtual topology may disrupt the ongoing traffic. In two-step reconfiguration approach the changes on lightpaths occur during the transition phase.

When several lightpaths are added to and removed from the topology the traffic disruption would be unavoidable. For this reason, most earlier studies developed techniques to minimize the disruption to the ongoing traffic [15, 90, 88]. Some studies dealt with this problem by performing the reconfiguration on all network elements concurrently [90], thus stopping majority of the traffic flows on the network for a short transition phase. Others applied step-by-step changes, such as one branch-exchange operation per step, until the new virtual topology is settled upon [15, 88]. In this second method traffic disruption is local (small compared to the whole network), but the transition phase takes longer time. In both cases, the traffic disruption because of the transition phase cannot be eliminated completely.

In an earlier study, hitless reconfiguration was defined as a reconfiguration process without the loss of any data [16]. According to the methodology proposed in [16], the transition between topologies was achieved by first establishing all new links without removing any link. The links of the old topology were removed only when the traffic was rerouted through the links of the new topology. Similarly, the traffic-disruption problem due to the unavailability of network resources can also be solved by our new approach if the step adjustments on the virtual topology are chosen carefully. In this study, the step adjustments are designed in a hitless manner, so that a change can be either a lightpath addition or a lightpath deletion. Before a lightpath is deleted, the traffic flows using this lightpath is rerouted through an alternate path.

The basis of the new adaptation method is to apply simple adjustments to the virtual topology when it is necessary. The periodical measurements on network traffic provides information on the lightpath loads. During each observation period the traffic on lightpaths are measured and at the end of the period the average lightpath loads are calculated based on the observed loads -therefore, throughout the text, the term lightpath load is used to describe the average lightpath load measured in the last observation period-. This information is used to decide whether or not an adjustment is needed. Basically, a new lightpath is added when a congestion is encountered, i.e., a lightpath load is dangerously high, and a lightpath is deleted if it is being under-utilized, i.e., its load is very low. Therefore, this scheme tends to balance the traffic loads evenly to the lightpaths, and to modify the connectivity of nodes when the balance changes. We introduce two system parameters to detect the lightpath-usage (in)efficiencies: High watermark W_H and low watermark W_L. The traffic loads on lightpaths are measured at the end of each *observation period*, which is typically hundreds of seconds. At the end of an observation period, if the load of one or more lightpaths is higher than W_H, a new lightpath is established to decrease that load. When the load of a link drops below W_L, that link will be torn down with one condition: if alternate paths exist for diverting the traffic that uses the considered lightpath.

The approach proposed can be used by an ISP to minimise the operational cost of its virtual topology by leasing only the appropriate amount of lightpaths as it considers to be really necessary. Previous reconfiguration and virtual topology design studies usually maximises the number of lightpaths in a virtual topology. The aim of these methods is to use all possible network resources to carry maximum amount of traffic. On the other hand the fibre infrastructure may not belong to the ISP, in which case the ISP has to lease the lightpaths. This is a common business model in real world, hence, the topology desing and reconfiguration methods are required to provide cost-efficient solutions. The new adaptation method is novel also in this sense, since it maintains only the necessary lightpaths.

The new adaptation scheme is computationally simple when compared to previous studies, e.g., [88]. Simplicity is vital for a reconfiguration method, because the process should be on-line. The system parameters W_H and W_L provide flexibility and control over the adaptation process: the frequency of the virtual-topology adjustments, average hop distance, and network resource usage can be changed by modifying the watermark values.

The following two chapters introduce the new adaptation method. This chapter provides a formal definition of the problem and the solution approach using mixed-integer linear program (MILP). Next chapter includes a heuristic method that can be used for large networks, and its performance study.

The rest of the chapter is organized as follows. Section 3.2 explains the network model and provides an informal statement of the problem. In Section 3.3, the new adaptation method, i.e., the selection of the best candidate lightpath to add or to drop at each step, is formulated as an optimization problem (hereafter called as MILP-A). The MILP-A formulation is solved for a simple network, and the solutions obtained are used to compare this method to an earlier study [17]. This early study proposed an MILP formulation (MILP-B) for solving the reconfiguration problem optimally, by minimising the number of changes for transition. In this sense, it provides a good comparison basis for the performance of the new adaptation approach.

3.2 Problem Definition

This section defines the network model assumed, and provides a formal definition of the problem to be solved.

3.2.1 Network Model

We consider a network of N nodes connected by bidirectional optical links forming an arbitrary physical topology. Each optical link supports W wavelengths, and any node i is assumed to have T_i transmitters and R_i receivers. We assume that each node is equipped with an optical cross-connect (OXC) with full wavelength-conversion capability, so that a lightpath can be established between any node pair if the resources (an optical transmitter at source, an optical receiver at destination, and at least a wavelength on each fiber link) are available along the path. Mechanisms to accommodate no wavelength conversion and different numbers of wavelengths on different links are straightforward. Unidirectional lightpaths are considered, since the traffic between two nodes is not necessarily symmetric (as can be seen in Fig. 3.1).

Each OXC is connected to an edge device, e.g., an IP router, which can be a source or a destination of a traffic flow and which can provide routing for multi-hop traffic passing by that node. We assume that each router is capable of processing all packet traffic flowing through it and of observing the amount of traffic on its outgoing lightpaths. In this study, for ease of explanation, a centralized approach is considered to the virtual-topology reconfiguration problem. A central manager will collect the virtual-link usage information from routers at the end of every observation period. Specifically, the link-usage information needed to make a reconfiguration decision consists of which links are overloaded, which links are underloaded, and what are the end-to-end packet-traffic intensities flowing through the overloaded links. The decision for a topology change will then be made by the central manager, and a signaling mechanism will be started if a lightpath addition or deletion is required as a result of the decision algorithm. An implicit assumption here is that the observation period is much longer (typically hundreds of seconds or longer) than the time it takes for control signals to propagate from various nodes to the central manager. One may expect that it is possible to design a decentralized protocol to do this job as well, but this is outside the scope of this thesis.

In the optical layer, shortest-path routing is used for routing lightpaths on the physical topology and the first-fit scheme for wavelength assignment [80]. For packet routing, a shortest-path (minimum-hop) routing scheme is considered, since it provides better usage of network links and is frequently used by existing routing protocols.

3.2.2 Problem Statement

Essentially, our aim is to provide a very good virtual topology under dynamic traffic conditions, by keeping the lightpath loads balanced, and by changing the virtual connectivity only when it is necessary. The adaptation process should be quick enough to match the long-term traffic fluctuations, but should not disturb the traffic unnecessarily. The problem to be solved can be stated as follows:
Given:

- A network graph $P(\mathcal{V}, E_P)$ where $\mathcal{V}$ is the set of nodes and E_P is the set of links connecting the nodes. Graph nodes correspond to network nodes with OXCs and links correspond to the fibers between nodes.
- Number of wavelength channels carried by each fiber.
- Number of transmitters and receivers at each node.
- Current virtual topology $V(\mathcal{V}, E_V)$ as another graph where the nodes correspond to the nodes in the physical topology. Each link in E_V corresponds to a direct optical lightpath between the nodes.
- Current traffic load carried by each lightpath.

Determine:

- Whether the current virtual topology of the network is efficient for the current traffic.
- Whether a change in the virtual topology should be made.
- If a change is necessary, which lightpaths should be added and/or deleted?

One can identify the following important steps in solving such a problem:

- Traffic should be monitored continuously to provide adequate information to the reconfiguration system.
- A decision mechanism is needed to trigger a virtual-topology change if the current topology is not convenient.
- Finally, the exact modification to the topology should be determined.

3.3 The Adaptation Method

The new approach adapts the virtual topology to the changing traffic by using the following idea: adding or deleting only one lightpath at a time. In this scheme, a new lightpath is added when one or more existing lightpath experiences a heavy load; similarly, an existing lightpath is deleted if the traffic load on some lightpath becomes very light. Two parameters are used to decide whether a lightpath load is too heavy or too light: High watermark W_H and low watermark W_L. A lightpath's load is considered to be too high if this load is higher than W_H, and too low if it is lower than W_L.

To follow the changes in lightpath loads, and therefore changes in traffic intensities, the lightpath loads are measured periodically, based on a time interval called *observation period*. A change to the virtual topology is made only when a lightpath's load is higher than W_H or lower than W_L at the end of an observation period; otherwise, no change is made until the end of the next observation period.

The adaptation cycle works as follows:

```
Start with initial virtual topology.
Every T seconds do: // T = Observation period
   Find the highest and lowest lightpath loads: L_H and L_L
   If L_H > W_H , then add an appropriate lightpath.
   Else if L_L < W_L , then delete an appropriate lightpath.
       Else do not adapt the virtual topology.
```

This problem can be stated formally as a mixed-integer linear program (MILP). The adaptation process is composed of consecutive steps of small adaptations, and the formulation is defined for one step (or cycle) of the method. This formulation can be solved for each step, i.e., at the end of each observation period, to reconfigure the virtual topology.

3.3.1 Formulation of Adaptation as an MILP (MILP-A)

This work assumes full wavelength-conversion capability at each node in this formulation. The following notations are used throughout the formulation:

- s and d denote *source* and *destination* of a traffic flow when used as a superscript or subscript.
- i and j denote *originating* and *terminating* nodes of a lightpath, respectively.
- m and n denote the end points of a physical link.

At any step of the adaptation, one of these three decisions can be made: Addition of a lightpath, deletion of a lightpath, or no change to the virtual topology. The choice of the decision is related to the highest and the lowest lightpath loads, and the watermark values. Selecting the proper action and the best lightpath which keeps the maximum link load as low as possible is a local optimization problem.

Given:

- Number of nodes in the network $= N$.
- Physical topology of the network $P = \{P_{mn}\}$, where P_{mn} indicates the number of fibers between physically adjacent nodes m and n, and $P_{mn} = P_{nm}$ for $m = 1, 2, ..., N$ and $n = 1, 2, ..., N$.
- Current traffic matrix $\Lambda = \{\Lambda_{sd}\}$ denotes the average traffic rate (in bits/s) claimed during the last observation period between every node pair, with $\Lambda_{ss} = 0$ for $s = 1, 2, ..., N$.
- Current virtual topology $V = \{V_{ij,q}\}$, where $V_{ij,q}$ is a binary value denoting the q^{th} lightpath between nodes i and j, and $V_{ii,q} = 0$. $V_{ij,0} = 0$ if there is no lightpath from node i to node j. $V_{ij,0} = V_{ij,1} = ... = V_{ij,k-1} = 1$ and $V_{ij,k} = 0$ if there are k lightpaths from node i to node j. Since lightpaths are not necessarily assumed to be bidirectional, $V_{ij,q} = 0 \not\Rightarrow V_{ji,q} = 0$.
- Number of wavelengths on each fiber $= W$.
- Capacity of each wavelength channel $= C$ bits/s.

- Number of transmitters and receivers at node i: T_i and R_i respectively.
- High watermark value $= W_H$, where $W_H \epsilon (0, 1)$, e.g., W_H = 0.8 implies that a lightpath is considered to be overloaded when its load exceeds $0.8 * C$.
- Low watermark value $= W_L$, where $W_L \epsilon (0, 1)$, and $W_L < W_H$.
- Highest and lowest lightpath loads measured during the last observation period: L^P_{Max} bits/s and L^P_{Min} bits/s, respectively.

Variables:

- Physical routing binary variable $p^{ij,q}_{mn} = 1$ if the q^{th} lightpath from node i to node j is routed through the physical link (m, n).
- New virtual topology: $V' = \{V'_{ij,q}\}$, where $V'_{ij,q}$ is defined similar to $V_{ij,q}$. Note that V' is at most one lightpath different from V.
- Traffic routing: The binary variable $\Upsilon^{sd}_{ij,q}$ is 1 when the traffic flowing from node s to node d traverses lightpath $V'_{ij,q}$, and 0 otherwise. $\Upsilon^{sd}_{ii,q} = 0$ by definition. The traffic from s to d is not bifurcated, i.e., all traffic between s and d will flow through the same path.
- Load of maximally-loaded lightpath in the network: L_{Max}.

Objective:

$$\text{Minimize} \quad L_{Max}$$

The objective function minimizes the load of the maximally-loaded lightpath in the network. This objective allows us to balance the network load in the new virtual topology, by addition or deletion of the best possible lightpath.

Constraints:

- On physical topology:

$$\forall i, j, q, \quad \sum_n p^{ij,q}_{in} = V'_{ij,q} \tag{3.1}$$

$$\forall i, j, q, \quad \sum_n p^{ij,q}_{nj} = V'_{ij,q} \tag{3.2}$$

$$\forall i, j, q, l, \quad \sum_n p^{ij,q}_{nl} - \sum_n p^{ij,q}_{ln} = 0 \ , \quad i \neq l \ \ and \ \ j \neq l \tag{3.3}$$

$$\forall m, n, \quad \sum_i \sum_j \sum_q p_{mn}^{ij,q} \leq W \cdot P_{mn} \tag{3.4}$$

$$\forall m, n, i, j, q, \quad p_{mn}^{ij,q} \leq V'_{ij,q} \tag{3.5}$$

Equation (3.1) ensures that only one outgoing physical link of the source node will be assigned to a lightpath. Equation (3.2) ensures that only one incoming physical link at the destination node will be assigned to a lightpath. Equation (3.3) guarantees that the number of incoming and outgoing links reserved for a lightpath at any intermediate node will be equal. The total number of wavelengths used between two nodes is limited to (the number of fiber links) * W by Equation (3.4). Note that we assume wavelength conversion capability on network nodes, and we use the wavelength channels on different fibers as non-distinguishable entities. To capture wavelength continuity in a non-wavelength-conversion network, some additional constraints will be needed. Equation (3.5) states that a physical link is assigned only if the lightpath exists.

- On virtual-topology connections:

$$\sum_i \sum_j \sum_q V'_{ij,q} = \sum_i \sum_j \sum_q V_{ij,q} + k_H - (1 - k_H) \cdot k_L \tag{3.6}$$

where

$$k_H = \left\lceil \frac{L^P_{Max}}{C} - W_H \right\rceil \text{ and } k_L = \left\lceil W_L - \frac{L^P_{Min}}{C} \right\rceil$$

$$\forall i, j, q, \; \left[1 + 2 \cdot (k_H - 1) \cdot k_L\right] \cdot (V'_{ij,q} - V_{ij,q}) \geq 0 \tag{3.7}$$

Note that the values of k_H and k_L are binary and they are calculated by using the maximum and the minimum lightpath loads measured in the last observation period, watermark values, and channel capacity. Therefore, these values are constant for the formulation. Note that k_H is unity when one or more lightpaths are experiencing heavy load. This will ensure that a new lightpath will be added to the virtual topology. k_L is unity when one or more lightpaths has a load below the low watermark. In this case if $k_H = 0$ (i.e., none of the lightpaths in the virtual topology is heavily loaded), a lightpath will be deleted. Thus, we are giving a higher priority to a lightpath addition than to a lightpath deletion to better accommodate the traffic. Equation (3.6) specifies the total number of lightpaths in the new virtual topology, i.e., it decides whether a change should be made, and if the answer is affirmative, whether a lightpath should be added or deleted. Equation (3.7) guarantees that the new virtual topology consists of the same set of lightpaths of the old virtual topology except that one lightpath is added or deleted.

- On virtual-topology traffic variables:

$$\forall s,d,l,q \;\; \sum_i \Upsilon_{il,q}^{sd} - \sum_i \Upsilon_{li,q}^{sd} = \begin{cases} 1 & l = d \\ 0 & l \neq s \text{ and } l \neq d \\ -1 & l = s \end{cases} \tag{3.8}$$

$$\forall s,d,i,j,q, \quad \Upsilon_{ij,q}^{sd} \leq V'_{ij,q} \tag{3.9}$$

$$L_{Max} \leq C \cdot W_H \tag{3.10}$$

$$\forall i,j,q, \quad \sum_s \sum_d \Lambda_{sd} \cdot \Upsilon_{ij,q}^{sd} \leq L_{Max} \tag{3.11}$$

Equation (3.8) is a multicommodity-flow equation controlling the routing of packet traffic on virtual links. Equation (3.9) ensures that traffic can only flow through an existing lightpath, and Equation (3.10) specifies the capacity constraint for any lightpath. Equation (3.11) constrains the load on any lightpath to be lower than or equal to the maximum load L_{Max}.

- On transceivers:

$$\forall i, \quad \sum_j \sum_q V'_{ij,q} \leq T_i \tag{3.12}$$

$$\forall j, \quad \sum_i \sum_q V'_{ij,q} \leq R_j \tag{3.13}$$

Equations (3.12)–(3.13) limit the total number of lightpaths originating from and terminated at a node to the total number of transmitters and receivers at that node.

The above formulation gives the best selection for a virtual-topology adjustment of one lightpath. Now, this new adaptation scheme will be compared, by solving the MILP-A, to an earlier reconfiguration method proposed in [17], i.e., MILP-B (MILP-B is not repeated in this thesis, since this formulation can be found in the cited paper).

3.3.2 Adaptation with Minimal Lightpath Change

This section focuses on the performance measurement of the new adaptation scheme, by solving the MILP-A formulation introduced above. The solutions obtained from the optimisation of the MILP-A are used to compare the adaptation with an earlier reconfiguration study [17], which will be called hereafter as optimal reconfiguration. The comparison is based on solving the formulations of both methods by using a standard solver, CPLEX [91].

The objective of the adaptation scheme is to minimise the maximum lightpath load at each step. Note that this minimisation may trigger a virtual topology change, only if a lightpath load is outside of the region defined by the high and low watermarks. In general, minimising the maximum link load in a network is equivalent to balancing the load distribution over the links. When the load is balanced less changes will be needed on the virtual topology. On the other hand, the use of low watermark limits the number of lightpaths in the topology, since it provides the deletions of unused or very lightly loaded lightpaths. Therefore, one can conclude that the adaptation scheme, undirectly, limits the number of changes on the virtual topology and the total number of lightpaths in the network. These two metrics are also the two important cost metrics in an optical network for an ISP: resource usage cost (number of lightpaths), and operation cost (number of changes during the reconfiguration). Hence, these two metrics are used for measuring the performance of the two methods.

The reconfiguration method proposed in [17] is based on a sequence of actions. Considering two snapshots of the traffic matrix at two time instants, first, the linear formulations $\mathcal{F}(1)$ and $\mathcal{F}(2)$ corresponding to these matrices are generated. Second, the solutions $\mathcal{S}(1)$ and $\mathcal{S}(2)$, corresponding to $\mathcal{F}(1)$ and $\mathcal{F}(2)$ are derived. Next action is to modify $\mathcal{F}(2)$ by adding a new constraint, which ensures that all the virtual topologies generated by the new formulation $\mathcal{F}'(2)$ would be optimal with regard to the objective function. The new objective function for $\mathcal{F}'(2)$ is designed to minimise the number of changes during the transition from the virtual topology found by $\mathcal{S}(1)$ to the virtual topology found by $\mathcal{S}'(2)$.

Here, since the comparison is based on minimising the two cost metrics, the objective function of the MILP-B [17], which minimises the average packet hop distance, is substituted with:

$$Minimize \sum_i \sum_j \sum_q V_{ij,q} \tag{3.14}$$

to minimize the total number of lightpaths in the network. Also the maximum load, L_{Max} is limited to the maximum value of W_H. When the optimum virtual topology is obtained from this formulation, this solution gives the minimum number of lightpaths that any virtual topology must contain to carry the given traffic demand. Let this number be η. This formulation can then be modified as follows:

1. The objective function is changed so that it guarantees that the new virtual topology will be as close as possible to the previous one.
2. A new constraint is added to the formulation to guarantee that the new virtual topology will have η lightpaths.

This modified formulation selects the closest virtual topology (to the current one) among the feasible topologies having exactly η lightpaths and able to carry the given traffic demand. With this property the solution obtained from the optimal reconfiguration is guaranteed to contain the minimum number of lightpaths in the virtual topology and to minimise the number of changes–number of lightpath additions and deletions–between consecutive topologies. The optimum reconfiguration can be summarised as follows:

```
Start with initial virtual topology.
Every Δ seconds do:
     Find the optimal virtual topology for the new traffic
     pattern by minimising the number of lightpaths.
     Let this number be η.
     Find a feasible virtual topology such that:
         It requires minimum number of changes from the
         previous topology.
         It contains η lightpaths.
```

Solving MILP-B is computationally tractable only for small networks, therefore, the performance comparison is based on a 6-node topology shown in Fig. 3.2. Four wavelengths per fibre and 4 transceiver pairs per node are assumed in the network. The traffic matrices are created using a realistic traffic model based on daily traffic fluctuations, as explained in Section 4.4.1. Two sets of results are obtained for two

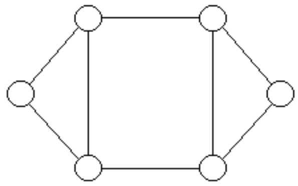

Figure 3.2: The 6-node network used in experiments.

different traffic scenarios. In the first case 33% of the node pairs have traffic flows between them, referred to as *low load* case. In the second case, 50% of the node pairs have non-zero traffic intensities, this case is referred to as *high load*.

The results of the adaptation scheme (MILP-A) are obtained with an observation period of 5 minutes, whereas the optimal topology is calculated every one hour. This difference in time intervals is necessary, because the optimal reconfiguration method calculates the new topology in one cycle, while the adaptation method performs the changes step-by-step, distributed over time.

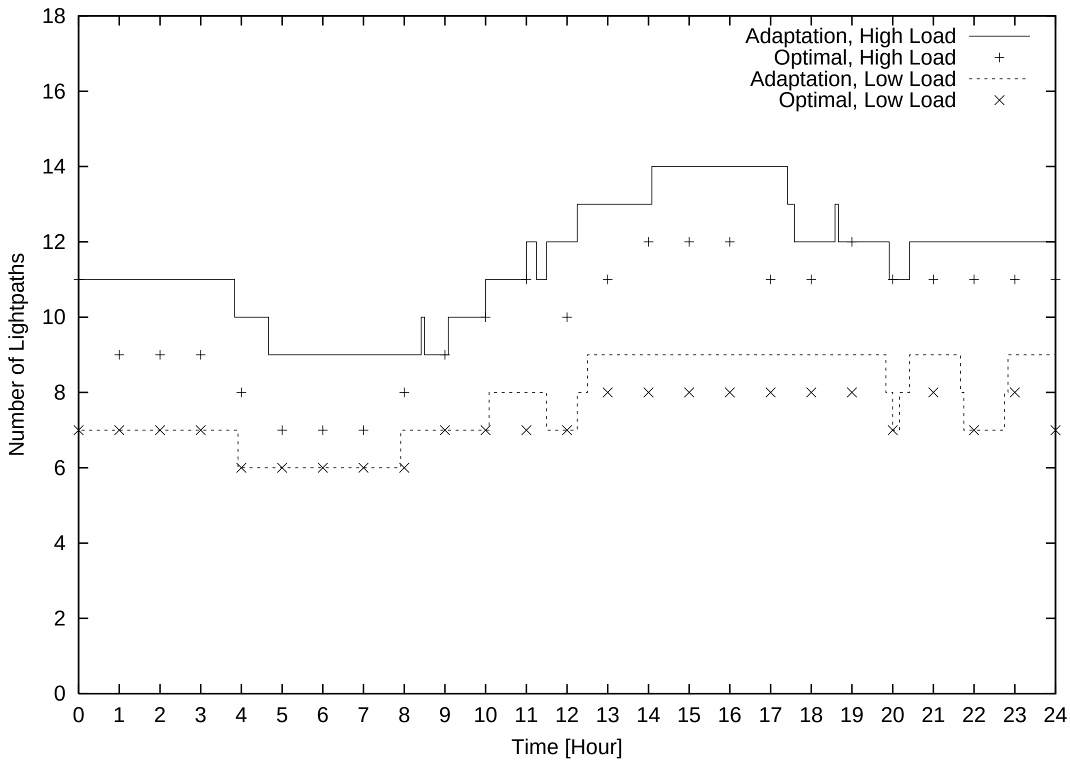

Figure 3.3: Number of lightpaths in virtual topologies obtained by solving our proposed adaptation MILP-A and the optimal reconfiguration MILP-B.

The two methods are compared for a 24-hour run. Figure 3.3 shows the number of lightpaths in virtual topologies obtained by the two methods. The results of the adaptation method are shown as continuous lines, and the results of the optimal reconfiguration–which occur every one hour–are shown as points. This plot shows that, as traffic load changes over time, the number of lightpaths in the virtual topology for both approaches changes in unison as well. (Compare the profiles in Fig. 3.1 with Fig. 3.3, noting that the measurements in Fig. 3.1 starts from 9:00 AM.) However, as also expected, the adaptation method establishes more lightpaths than those in the optimal topology, and the difference in the number of lightpaths increases slightly with the traffic load.

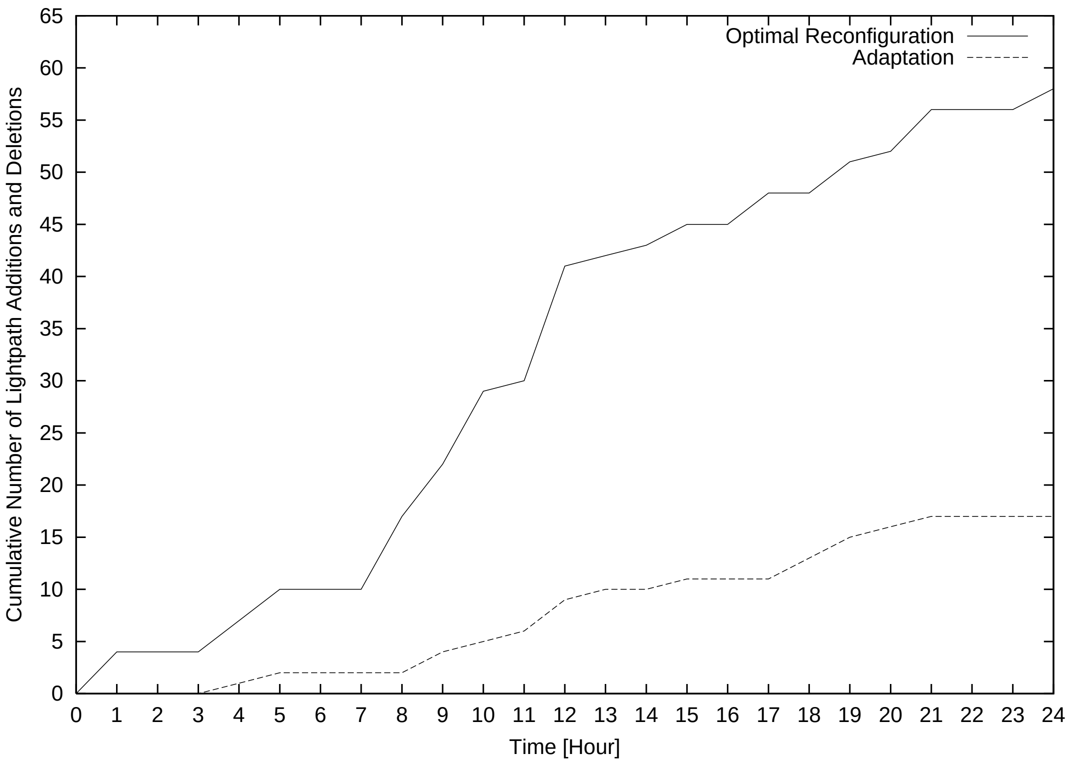

Figure 3.4: Comparison of the cumulative number of lightpath changes–sum of lightpath additions and deletions–between our adaptation scheme and optimal reconfiguration.

To show the number of lightpath changes during the reconfiguration, the cumulative number of lightpath additions and deletions for both methods during the same 24-hour period is plotted in Figure 3.4. In this figure only the results from the high-traffic-load experiment are shown. Here, one can observe that the adaptation scheme can reconfigure the topology using far fewer lightpath changes compared to the optimal reconfiguration. Considering the Figures 3.3 and 3.4, the adaptation scheme is very efficient, since it can follow the changes in traffic and adjust the topology with very small number of changes in the virtual topology.

One remark here is that the complexity of the MILP-A limits its use for large networks. The formulation of the adaptation method is much simpler compared to the optimal reconfiguration, and the solutions in the numerical examples are obtained much faster than the latter. However, it is still too complex to be useful for real-size networks. Thus, in the following chapter, an efficient heuristic method is designed to adapt the virtual topology in response to traffic changes in large networks.

4. A HEURISTIC ALGORITHM FOR VIRTUAL TOPOLOGY ADAPTATION

This chapter introduces a heuristic algorithm for the virtual topology adaptation method proposed in the Chapter 3. The key ideas of the adaptation algorithm are introduced in Section 4.1, and an outline of the algorithm is given in Section 4.2. In Section 4.4, representative numerical examples employing the new adaptation algorithm are discussed. This part also includes variations of the method and their numerical illustrations. Section 4.7 concludes the paper.

4.1 Key Ideas

In this study, a one-phase reconfiguration algorithm is proposed where the network traffic flows are continuously observed and simple updates to the virtual topology are made whenever necessary. The key idea of this approach is to set up new lightpaths when congestion occurs or tear down existing lightpaths when they are not used efficiently. Since frequent changes to the virtual topology are not desirable, the length of traffic-observation period and the instant to trigger a virtual-topology change (either a lightpath addition or a lightpath deletion) should be selected carefully.

The length of the observation period is a parameter which can be selected and updated according to the specific properties of the network traffic. It specifies how frequently the adaptation algorithm is activated. We use link loads of the lightpaths in the current topology as input to the triggering function of the adaptation algorithm. Since having perfectly balanced loads on every link may not be practical or may not be possible, the aim is to maintain quasi-balanced virtual links. We can define a quasi-balanced topology as follows:

$$V = \{V_1, V_2, \ldots, V_k\}$$
$$W_L \leq \delta(V_i) \leq W_H \quad i = 1, 2, \ldots, k$$

Here k is the number of lightpaths in the virtual topology whose lightpaths are denoted by V, $\delta(V_i)$ is the load of the i^{th} lightpath, and W_L and W_H are the low watermark and the high watermark values. The objective of the adaptation algorithm is to keep the load of every lightpath between these two thresholds. At the end of an observation period, if the load of one or more lightpaths is higher than W_H, the lightpath V_{max} having the maximum load $\delta_{max} = \delta(V_{max})$ will be considered, and a new lightpath will be established to decrease δ_{max}. The new lightpath V_{new} should carry at least one of the traffic flows that V_{max} was carrying, so that, after the new lightpath is established, the load of V_{max} can decrease. As an extreme case, one may set $W_H = 100$ and $W_L = 0$ minimizing the changes on the virtual topology but also decreasing the quick reaction ability of the network to dynamic changes in the traffic. The characteristic of the traffic should determine the proper choice for the watermark values.

There are several ways to select the source and the destination nodes of the new lightpath to be added. We choose to establish a lightpath between the end nodes of the multi-hop traffic with the highest load using lightpath V_{max}. By doing so, we expect to have a noticeable decrease in δ_{max}. If V_{max} is only carrying single-hop traffic, then a new lightpath between the end nodes of V_{max} is added.

When the load of a virtual link decreases to a value below W_L, either because of changes in the dynamic traffic or because of traffic rerouting due to previous topology changes, the algorithm will tear down that link, unless it is a part of the only path for a traffic flow. The main idea behind the deletion of underloaded lightpaths is to release the network resources whenever they are not used efficiently and to have them available when new lightpaths are to be established later. In some cases, it is possible for the deletion of a lightpath to lead to congestion on another lightpath and consequently to an addition of a new lightpath. However, the traffic rates are changing dynamically, and in such an environment this behavior should not be interpreted as instability. This dynamism is necessary for the method to find a very good set of lightpaths for the current traffic. Whenever required, this ping-pong effect can be controlled by adjusting the watermark values.

Whenever a lightpath is added or deleted, the routes for all traffic flows are re-calculated. Occasionally, new overloaded or underloaded lightpaths may occur after rerouting due to change of the routes of some traffic, but this is a part of the adaptation process and indicates that it may be beneficial to add or delete more lightpaths. So, a few more steps may be helpful for the virtual topology to reach a load-balanced state.

It should also be noted that variations to this approach exist, such as deletion or addition of multiple lightpaths in one step when the load of more than one lightpath is outside the region defined by watermarks. In this thesis, we study the basic case where only one lightpath is allowed to be added or deleted; and we also study a variation of the basic case where unlimited number of additions or deletions are allowed at every observation period.

4.2 Outline of the Algorithm

A pseudo-code of the heuristic adaptation algorithm is given in Fig. 4.1. The term λ_{sd} is used to denote the traffic rate from node s to node d normalized to the capacity of a virtual link. The notation (i, j) is used to express the originating and the terminating nodes of a lightpath, and G to represent the virtual-topology graph. $G - V_{ij}$ is a subgraph of G with one edge from node i to node j extracted.

The algorithm first checks for any lost end-to-end traffic because of a lack of connectivity in the virtual topology. Under normal operation, the adaptation process expects a connected network of already established lightpaths and it preserves the connectivity of node pairs as long as these nodes have a nonzero rate of traffic flowing between them. In some cases, traffic rate from a node s to a node d may cease, and if this silent period is longer than an observation period, a lightpath providing the connectivity may be torn down.

Input:

- Virtual topology, $V = \{V_{ij}\}$
- Current traffic rates, λ_{sd}
- Load of virtual links, δ_{ij}
- Available number of transmitters and receivers at node i, T_i and R_i

Output:
A decision which can be one of the following:

- Addition of a lightpath between nodes i_{new} and j_{new}
- Deletion of an existing lightpath between i_d and j_d
- No change to the virtual topology

Algorithm:
At the end of every observation period:

```
If λ_sd > 0 and there is no path from s to d then
    Establish V_sd for max_{s,d}(λ_sd)
else
    Find the maximum link load: δ_max = δ_{i_max,j_max} = max_{i,j}(δ_ij)
    if δ_max > W_H then
        max_load ← 0
        for every multi-hop traffic flow λ_sd using V_{i_max,j_max} do
            if (λ_sd > max_load) and (T_s > 0) and (R_d > 0) then
                max_load ← λ_sd
                s_m ← s,   d_m ← d
            endif
        endfor
        if max_load > 0 then
            if V_{s_m d_m} can be established then
                Establish V_{s_m d_m}
                T_{s_m} ← T_{s_m} − 1,   R_{d_m} ← R_{d_m} − 1
            endif
        else
            if V_{i_max j_max} can be established then
                Establish V_{i_max j_max}
                T_{i_max} ← T_{i_max} − 1,   R_{j_max} ← R_{j_max} − 1
            endif
        endif
    endif
    if no lightpath is added then
        Find the minimum link load: δ_min = δ_{i_min,j_min} = min_{i,j}(δ_ij)
        complete ← false
        while (δ_min < W_L) and not complete do
            if (i_min, j_min are connected in G − V_{i_min j_min}) or (δ_min = 0) then
                Tear down V_{i_min j_min}
                T_{i_min} ← T_{i_min} + 1,   R_{j_min} ← R_{j_min} + 1
                complete ← true
            else
                Find the next lowest link load: δ_nextlow
                δ_min ← δ_nextlow
                (i_min, j_min) ← (i_nextlow, j_nextlow)
            endif
        endwhile
    endif
endif
```

Figure 4.1: Pseudocode of heuristic adaptation algorithm.

If all of the communicating node pairs have at least one route, then the lightpath loads are considered next. Information on ligthpath loads is collected from routers which are source of a lightpath, by the central manager (which is running the adaptation algorithm), at the end of each observation period (please see Section 3.2.1). A change to the virtual topology is made only when a link load is higher than W_H or lower than W_L; otherwise, no change is made until the end of the next observation period. The maximum link load is considered first, since a highly-loaded link (and with a potentially increasing load) would cause more serious problems compared to an underloaded link. Lightpath deletion is considered only if no lightpath is added, because only one lightpath change at a time (i.e. in one observation period) is allowed. Lightpath deletion is performed only if that link is not on the unique path between a node pair i and j for which the traffic rate $\lambda_{ij} > 0$. All lightpaths with a load lower than W_L are considered in increasing order of their loads, until an appropriate lightpath is found.

The second variation that we implement in this study adds or deletes lightpaths until all loads are between the watermarks. In an observation period, only additions or deletions are allowed exclusively, but the number of changes is not limited. Although this method requires more processing compared to the basic algorithm explained above, it shows a very interesting aspect of the adaptation method, namely, even though multiple lightpath additions or deletions are allowed, the numerical examples reveal that, in the vast majority of adaptations, just a single lightpath addition or deletion suffices to bring the virtual topology to the quasi-balanced state.

4.3 Complexity of the Adaptation

The virtual topology reconfiguration problem is NP-hard, therefore the exact formulations given for this problem can be solved for only small-size networks. A reconfiguration method would be useful for large networks only if it could give a solution, although sub-optimal, in an acceptable amount of time. The adaptation algorithm proposed above is heuristic, i.e., the solutions it provides are not guaranteed to be optimal, but its complexity is low and can be used for large networks.

The algorithm in Figure 4.1 shows the computations performed at each observation period. There are two main program blocks that may be run, according to the lightpath loads and watermarks. If a new lightpath is to be added, first block starting with *if* $\delta_{max} > W_H$ will be executed. The complexity of the inner for loop is $O(N^2)$ where N is the number of nodes in the network. The if block following the for loop tries to establish a lightpath using a routing and wavelength assignment function whose complexity is $O(N^2 \log N)$. Hence, the complexity of the first block is $O(N^2 \log N)$. If no lightpath is added by the first part of the algorithm, then the second block is executed. Finding the minimum lightpath load and ordering the loads that are below the low watermark takes a time of $O(\tau N \log N)$ where τ is the number of transceivers at each node (τN is the maximum number of lightpaths that can be established in the network at a time, and typically $\tau < N$). The following while loop is executed at most τN times. Inside the while loop connectivity of the two nodes are tested using e.g. Dijksta's method, with complexity of $O(N \log N + E)$ (E is the number of edges, in a wide-area network $E \approx N$). This leads to a complexity of $O(\tau N^2 \log N)$ for the second block. Finally, the overall worst-case complexity of the adaptation cycle is $O(\tau N^2 \log N)$.

4.4 Illustrative Numerical Examples

This section includes several simulation results of the new adaptation algorithm. After introducing the simulation environment, the typical operation of the method is illustrated. Next, the effect of the watermarks and the observation period on the network are examined. A version of the algorithm where unlimited changes are allowed is also examined in this section.

4.4.1 Simulation Environment

Simulation experiments were conducted to investigate the fitness of the adaptation approach and to expose the effect of various system parameters. In the experiments, we implemented the heuristic algorithm shown in Fig. 4.1 in an event-driven simulation.

A telecommunication mesh network of 19 nodes interconnected by 31 bidirectional links was used as the physical topology for the numerical examples shown in this section for illustration purposes (Fig. 4.2). For the examples, we assume that the number of wavelengths is 16 per fiber link and equal for all links. Each node is assumed to have 8 transmitters and 8 receivers.

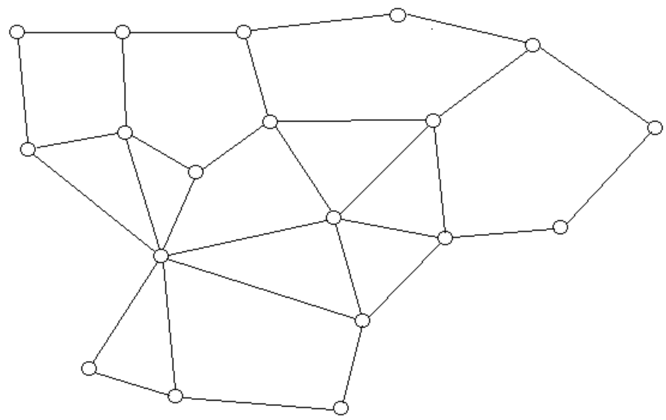

Figure 4.2: An example telco network used as the physical topology in this study.

Since this study is focused on backbone networks, a traffic model is derived based on the observations on several backbones' link loads. Backbone traffic is the aggregation of several end systems' traffic, and the aggregation process filters out the short-term variations. On the other hand, as can be seen in Fig. 3.1, long-term variations (on a scale of hours) remain and repeat their pattern in one-day periods. To obtain a realistic model, we sampled some representative link traffic rates (Fig. 3.1 shows one of them) from real networks over a 24-hour period. The sampling allows us to represent the traffic rate as a function of time. The average traffic rate at a point in time between two sampling points of the rate function was calculated by a linear interpolation to represent the continuous change in traffic. These traffic-rate functions are used to generate the simulated traffic between any node pair in the network. We expect that the effect of aggregation and smoothing is more intense at the edge routers with respect to the smoothing effect of aggregation of traffic from a few node-pairs on a network link. Therefore, we believe that it is reasonable to use the link traffic to simulate the traffic between the edge routers for simulation purposes. This assumption is not vital for the adaptation method proposed, and the method should work for any input traffic having long-term fluctuations, since the reaction time can be adjusted according the time scale of these fluctuations.

To create a simulation pattern for each node pair, a traffic rate function is randomly selected among 5 different patterns. The chosen rate function is then scaled by multiplying it by a random value. The random scaling allows to create differences between the traffic rates of different node pairs, even when the same pattern is chosen for different node pairs. Each random scaling factor α_{sd} (for a node pair $< s, d >$)

takes values in the range $[0.2, 1.2]$ in the examples to create appropriate variations in the traffic volumes between different node pairs. In this model, every element of the traffic matrix is a continuous function of time representing the traffic rates during one day, rather than a single value.

It should be also noted that no assumption on future traffic demand is being made in the adaptation scheme except that (without loss of generality) the traffic rates are expected to fluctuate slowly compared to the time interval between adaptation steps. Using this model, the traffic matrix was randomly generated to have nonzero rate functions for 60% of the node pairs (values on the main diagonal remain zero). Appendix D includes the technical details of the simulation program.

In this study, the following performance metrics are employed to reveal the different aspects of the adaptation scheme: The average time between consecutive reconfiguration steps, the traffic-weighted average hop distance [1], number of lightpaths in the virtual topology, and the percentage of links with loads in the balanced region $[W_L, W_H]$. Through these metrics, the effects of changing the system parameters are examined, specifically, high watermark, low watermark, and length of observation period.

4.4.2 Typical Operation

The goal of the first experiment is to demonstrate the operation of the system by measuring the maximal and minimal lightpath loads in the network at the end of every observation period. High and low watermarks will define a *balance region* in which the lightpath loads are allowed to fluctuate as long as they do not exceed the high watermark or drop below the low watermark. Every time a link load goes out of the balance region, the system will try to adapt to the new traffic conditions. In most of the examples that follow, we will allow a maximum of only one lightpath addition or deletion; finally, in Section 4.4.7, we will examine the system performance when an unlimited number of lightpath additions and deletions are allowed.

An exemplary operation of the system for a 3-day period with $W_H = 70$ and $W_L = 10$ is shown in Fig. 4.3. Note that, in this section the high and low watermark values are represented as $0 \le W_H, W_L \le 100$, normalized to 100, so that W_H and W_L

[1] It is calculated as the average value of the product of traffic intensity and hop distance for each node pair.

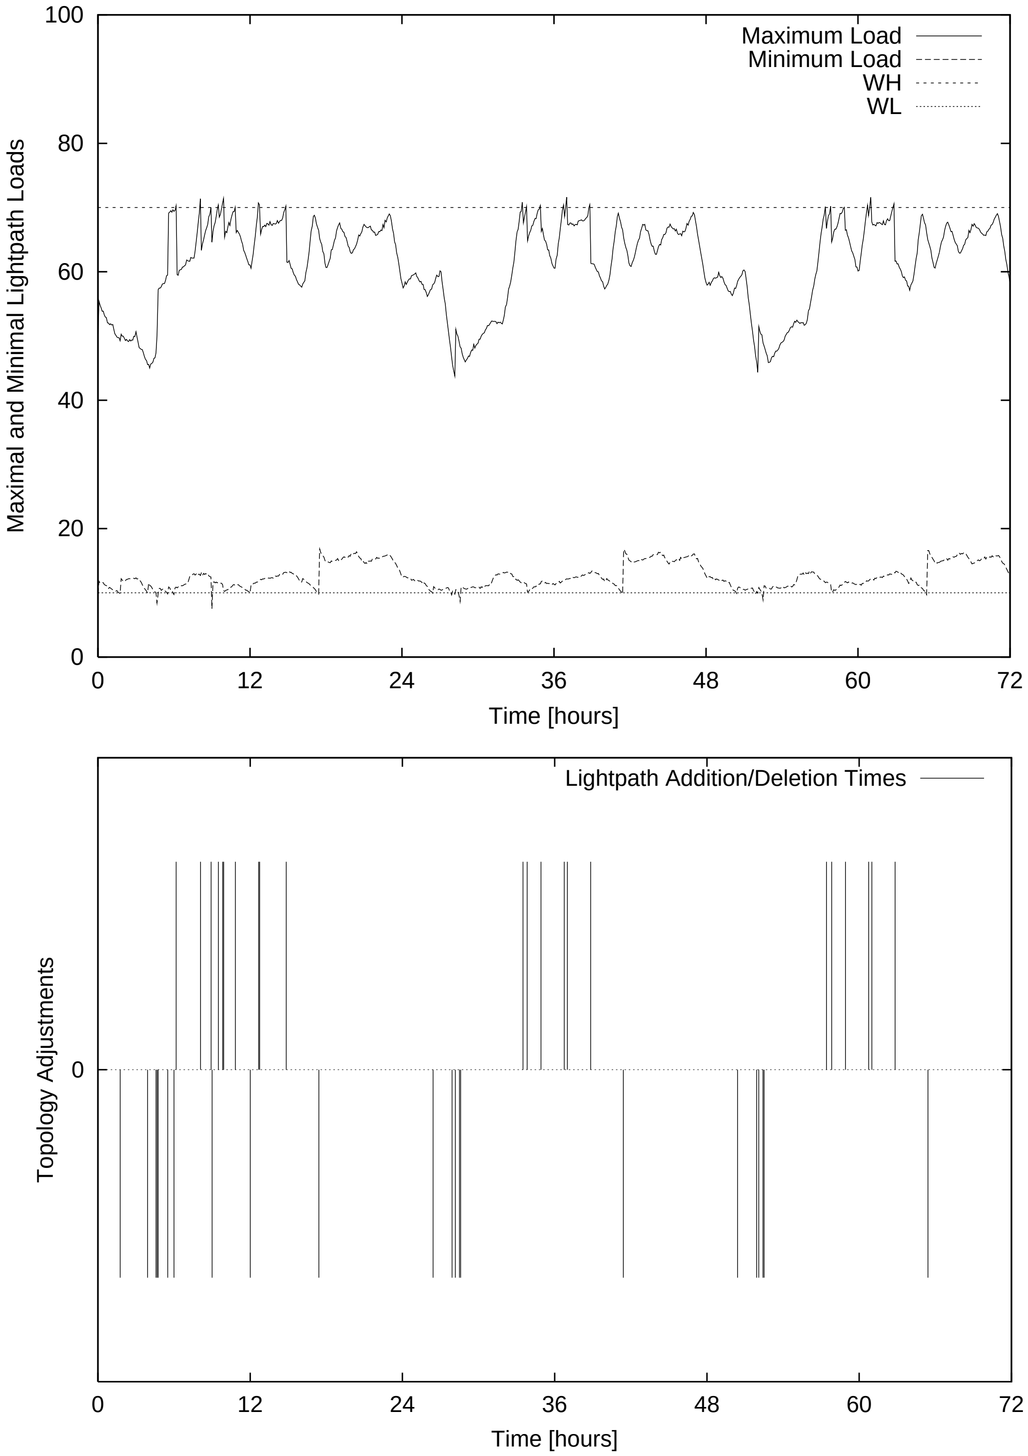

Figure 4.3: a. Maximal and minimal lightpath loads in the network during a 3-day run ($W_H = 70, W_L = 10$). b. Impulse graphic indicating times of lightpath addition or deletion. A positive impulse indicates a lightpath addition, while a negative impulse indicates a lightpath deletion.

can be interpreted as percentage loads. Figure 4.3(a) plots the maximal and minimal loads in the network observed at the end of each observation period (300 sec for this experiment). The watermark values are also plotted in the figure as lines, to expose the relation between the lightpath loads and the high and low watermarks. As can be seen in this figure, when a lightpath load reaches the high watermark, the algorithm reacts by adding a new lightpath, which provides a drop in the highest lightpath load. Similarly, when a lightpath load drops below the low watermark, another topology change–a lightpath deletion–occur to increase the lowest load in the network. Figure 4.3(b) shows the times of virtual topology adjustments for the same experiment. In this figure, a positive impulse indicates a lightpath addition and a negative impulse indicates a lightpath deletion. The same average traffic rate functions were repeated for each of the three days, but slight differences can be seen for the same time of different days related to randomly generated traffic and to the dynamic nature of the algorithm. Additional figures for different values of the watermarks can be seen in Appendix B.

4.4.3 Effect of Watermarks

The effect of changing the value of W_H on traffic-weighted average hop distance for a fixed value of $W_L = 10$ is shown in Fig. 4.4. The hours on the x axis correspond to the real hours of the day. During the night (until 9 or 10 in the morning), as the traffic load decreases all over the network, the lightpath loads decrease as well and the adaptation scheme deletes the lightly-loaded links causing an increase in average hop distance. During the daytime, the average hop distance decreases with the increasing traffic load. As can be seen in the figure, the average hop distance decreases when W_H takes smaller values, since the lightpaths are allowed to carry less load and the W_H limit can be reached more quickly. This effect results in new lightpath additions. For the case where $W_H = 60$, more lightpaths are deleted compared to the case where $W_H = 70$ during nighttime. This effect manifests itself in Fig. 4.4 as an aggressive climb in the average hop distance, and is related to the tightness of the watermarks: As the values of the watermarks approach each other, more lightpaths have loads closer to the bounds, and small changes in traffic may result in topology changes.

Figure 4.5 plots the effect of W_L on the average hop distance for a fixed value of $W_H = 70$. When W_L is small, lightpaths are deleted less frequently since fewer lightpaths will have their load drop below a small W_L, compared to the case with a higher W_L value. Therefore, the virtual topology keeps more lightpaths alive and has a lower hop-distance value.

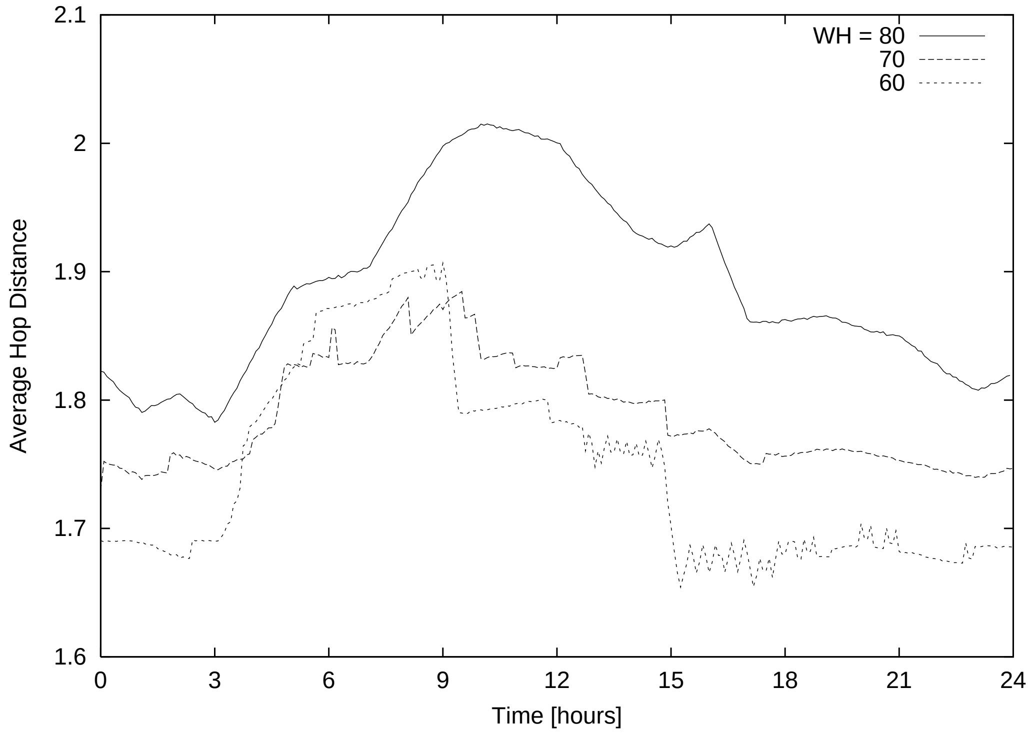

Figure 4.4: Traffic-weighted average hop distances during a day, for different values of W_H with $W_L = 10$.

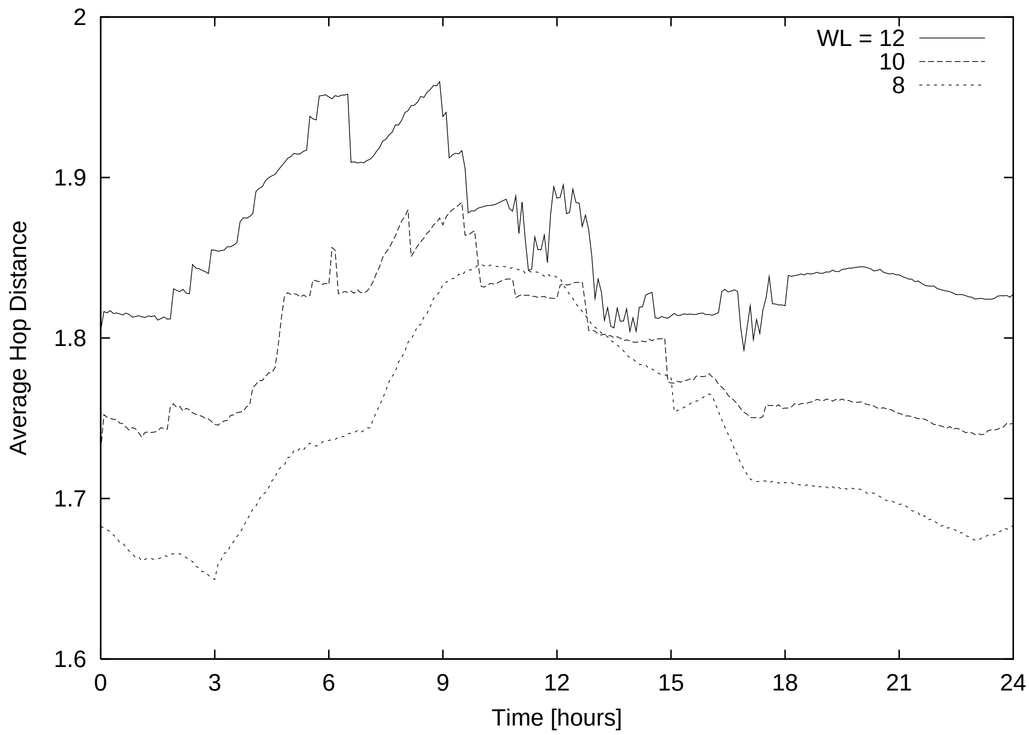

Figure 4.5: Traffic-weighted average hop distances during a day, for different values of W_L with $W_H = 70$.

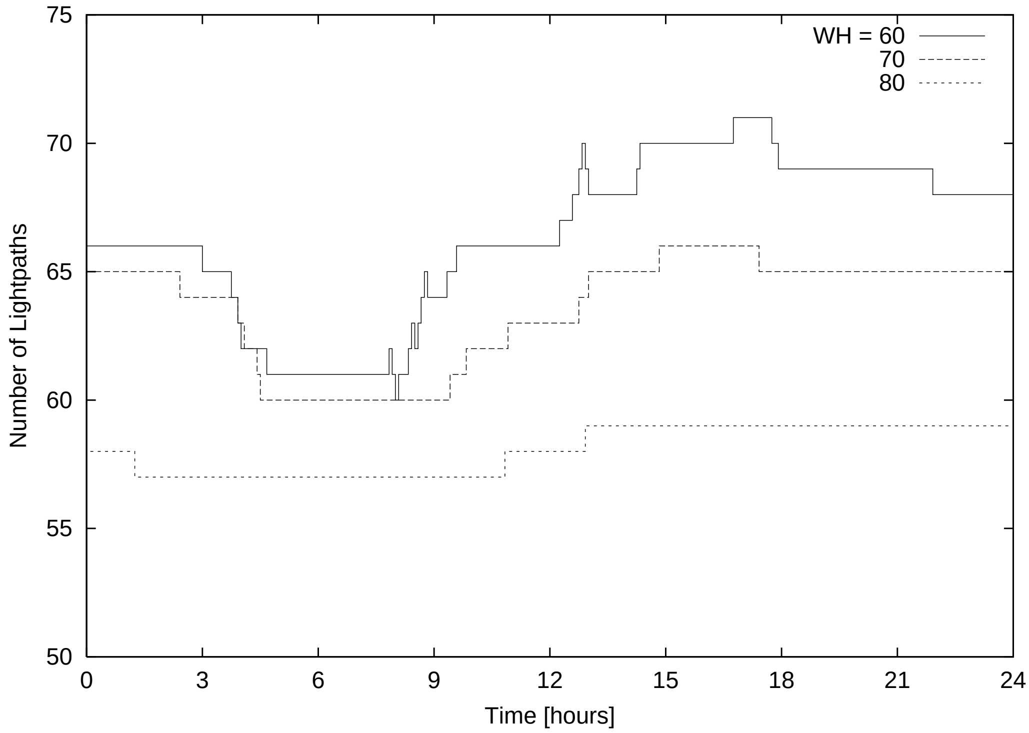

Figure 4.6: Number of lightpaths during a day, for different values of W_H with $W_L = 10$.

The total number of lightpaths in the virtual topology during a one-day period is plotted in Fig. 4.6 for different values of W_H with $W_L = 10$. Since the algorithm allows the addition or deletion of lightpaths at the end of each observation period, the total number of lightpaths in the virtual topology may change in time. The figure shows that higher values of W_H result in smaller number of lightpaths and more economic usage of network resources such as transceivers. The reason of this result is that, higher values of W_H allow the lightpaths to carry higher traffic loads without violating the high watermark. Consistent with the results obtained for average hop distance, the number of lightpaths is smaller when the network load decreases due to time-of-the-day effect.

Figure 4.7 plots the number of lightpaths for different values of W_L with $W_H = 70$. The number of lightpaths in the system decreases when W_L increases because the adaptation scheme tends to delete lightpaths more aggressively. The total number of lightpaths is smaller when W_L takes higher values. Since a larger number of lightpaths requires larger amount of network resources, one may prefer to keep this number small for economic use of resources. These examples show that the two system parameters W_H and W_L can be used to balance the use of resources and the average hop distance.

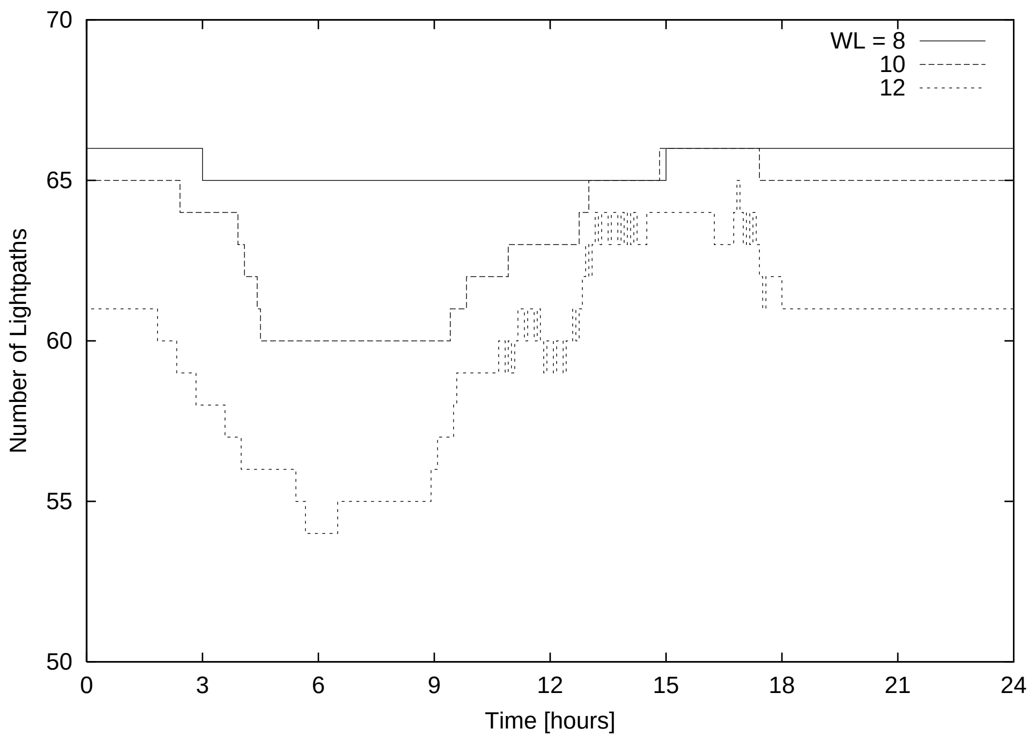

Figure 4.7: Number of lightpaths during a day, for different values of W_L with $W_H = 70$.

Figures 4.8 and 4.9 plot the average time between two consecutive adjustments in the virtual topology, as a function of W_H and W_L, respectively. These values are obtained by averaging the time of consecutive adjustments during a 5-day simulation. We observe in Fig. 4.8 that the virtual topology is adjusted less frequently when W_H takes higher values. This effect is related to the difference between the two watermarks; i.e., when the difference between W_H and W_L is high, fewer adjustments to the virtual topology are triggered. Similarly, when W_L takes higher values, the virtual topology is adjusted more frequently as can be seen in Figure 4.9. These two figures show that the average number of lightpath additions and deletions can be controlled by appropriately choosing the watermarks.

To give an idea on the stability of the lightpaths that are established by the adaptation method, the Tables 4.1 and 4.2 are given. The average lifetime of a lightpath in the network is shown in Table 4.1 for different high and low watermark values. The average is calculated for a simulation period of five days, and in terms of seconds. In this table, all lightpaths that exist in the network for some period of time is included in the calculation. The table shows that a lightpath remains alive around 21 hours even in the case where the watermarks are close to each other ($W_H = 60$ and $W_L = 14$).

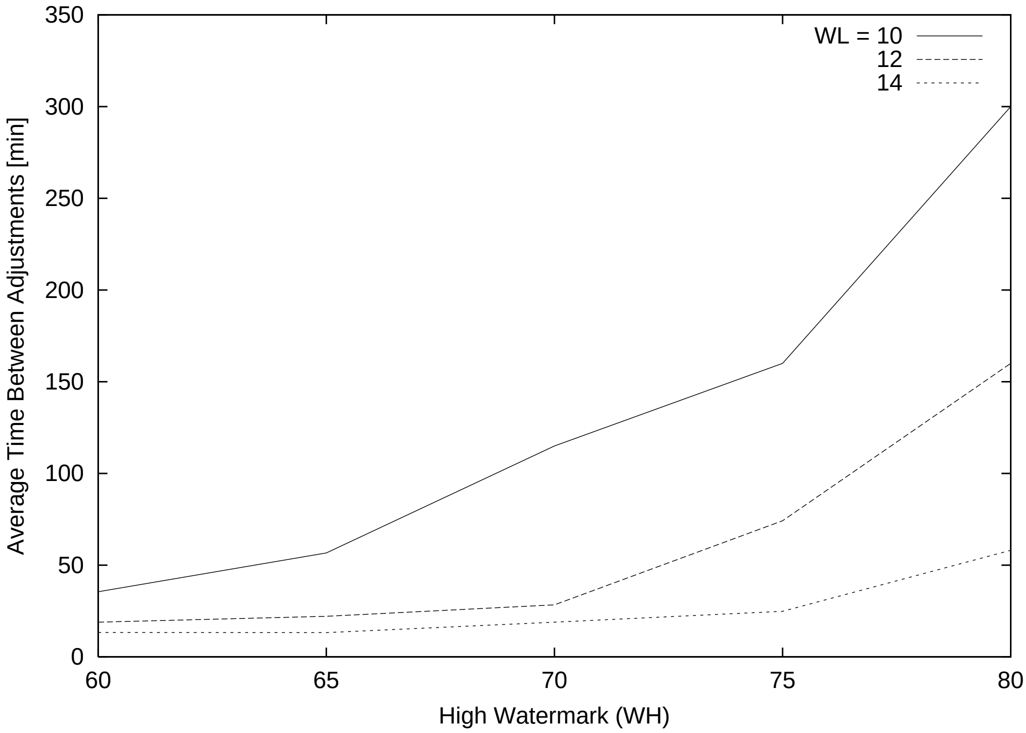

Figure 4.8: Average time between two consecutive reconfiguration steps, as a function of W_H.

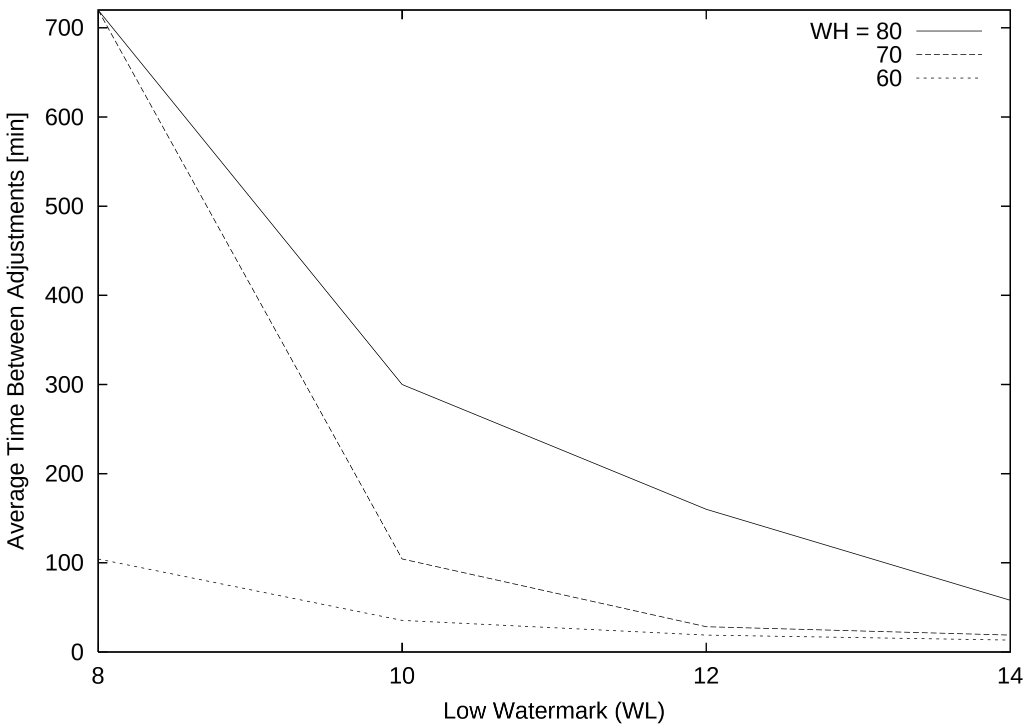

Figure 4.9: Average time between two consecutive reconfiguration steps, as a function of W_L.

Table 4.1: Average lifetime (in seconds) of the lightpaths for different watermark values

W_H	W_L			
	8	10	12	14
60	277135	167152	105369	75700
65	285980	201949	115889	75293
70	373521	262106	128661	98563
75	399535	266597	226732	119530

Table 4.2: Average lifetime (in seconds) of the lightpaths for different watermark values excluding all-time existing lightpaths

W_H	W_L			
	8	10	12	14
60	68261	72570	49727	54076
65	51309	64013	59779	43323
70	34500	63293	71365	47942
75	49260	99043	108145	51682

The gap between the watermarks affect the lifetime average, and the maximum value is at the lower left corner of the table, where $W_H = 75$ and $W_L = 8$. The lifetime average of the lightpaths that are established by the adaptation algorithm can be seen in Table 4.2. These average values exclude the lightpaths that exist during the whole simulation period, hence they show the changing lightpaths' statistics. Even in this table, one may see that the average lifetime of a lightpath established by the adaptation is quite long, changing from approximately 10 hours (for $W_H = 70$ and $W_L = 8$) to 30 hours (for $W_H = 75$ and $W_L = 12$).

4.4.4 Effect of Observation Period

A third parameter that affects the system's behavior is the length of observation period over which the lightpath loads are measured. The average traffic rate on a lightpath is calculated by measuring the total amount of traffic flowing through it and dividing that amount by the length of the period. The lightpath load is the ratio of the measured average traffic rate to the lightpath capacity. Since the traffic rate is taken as an average value during a period, the length of one period will affect the system's ability to track the changes in traffic. Momentary changes could be measured and reacted to by using enough short observation periods; but a smoothing effect on

Table 4.3: Probability of adjustment decision at the end of an observation period ($W_H = 70, W_L = 10$)

Length of Observation Period [s]	Total Number of Periods in One Day	Average Number of Adjustments	Adjustment Probability
100	864	12.33	0.014
200	432	12.17	0.028
400	216	14	0.065
800	108	17.6	0.163

traffic measurements by using longer periods would be more accurate to avoid the instabilities and frequent adjustments. On the other hand, long periods would decrease the efficiency of step-by-step adaptation on following the traffic changes closely if we do not allow more than one lightpath to be changed at the end of each period. Table 4.3 shows the effect of the observation-period length on topology-adjustment decisions. In this experiment, $W_H = 70$ and $W_L = 10$, and the number of adjustments in a day is averaged over a 5-day run.

The first column shows the length of observation period in seconds, used during the experiments. The values in the second column are the total number of observation periods in a day; therefore, they also show how many times the adaptation algorithm was activated. The average number of adjustments during one day (number of additions and number of deletions are almost equal) is shown in the third column. The probability of adjustment at the end of an observation period is the ratio of the value in the third column to the one in the second column. As the observation period gets longer, it is more likely that the algorithm will decide to make an adjustment in the virtual topology. This result confirms the claim above, namely, short observation periods keep track of the traffic more accurately.

4.4.5 Effect of Window Size

The basic algorithm in Fig. 4.1 may sometimes react very quickly because a change is made based on the measurement during the last observation period. If sudden jumps occur in traffic load, keeping a longer history information may be more useful to prevent the false alarms in the network from triggering a change. To study the effect of such history information, we store the lightpath loads measured at each period during a history window. The basic algorithm in Fig. 4.1 is the case where the window size is equal to 1. At each new period, the window is slided forward by one unit to

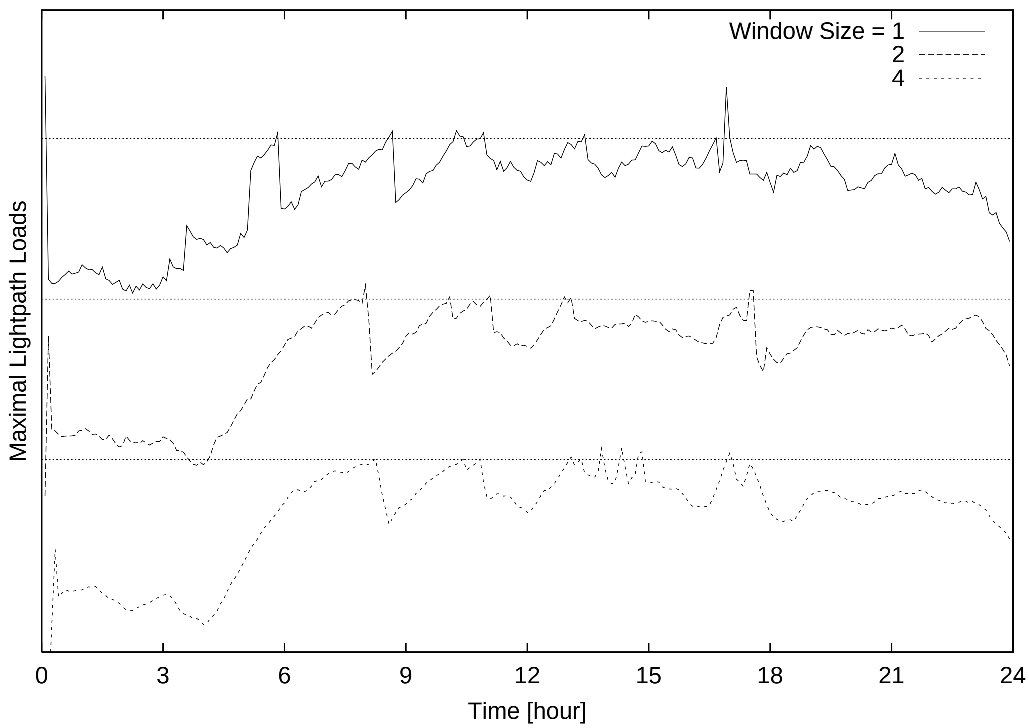

Figure 4.10: Maximal lightpath loads during a day for different history window sizes. The plots are shifted for better view.

include the new measurement and to exclude the oldest one. The decision on lightpath addition/deletion is based on the average load calculated using the entire window. Fig. 4.10 shows three plots of maximal lightpath loads in the network, interlaced for window sizes, 1, 2, and 4. In this experiment the watermark values are the same for three plots, $W_H = 70$ and $W_L = 10$. The graphics are shifted vertically for a clear view and a 70% horizontal line (high watermark value) is shown for each plot for guidance. As the window size increases, the observed average loads become smoother, and some changes are delayed or skipped. As an example to this effect, one can see that, when the window size is 1, a lightpath is added around 6 AM, but this addition does not take place for larger window sizes. However, increasing the window size would also delay some necessary changes. Specifically, when many lightpath loads increase at the same time, this delay may cause problems.

4.4.6 Load Balancing

The adaptation mechanism works to balance the lightpath loads all over the network to maintain all loads in the balance region. In this section, the load-balancing ability of the system is shown.

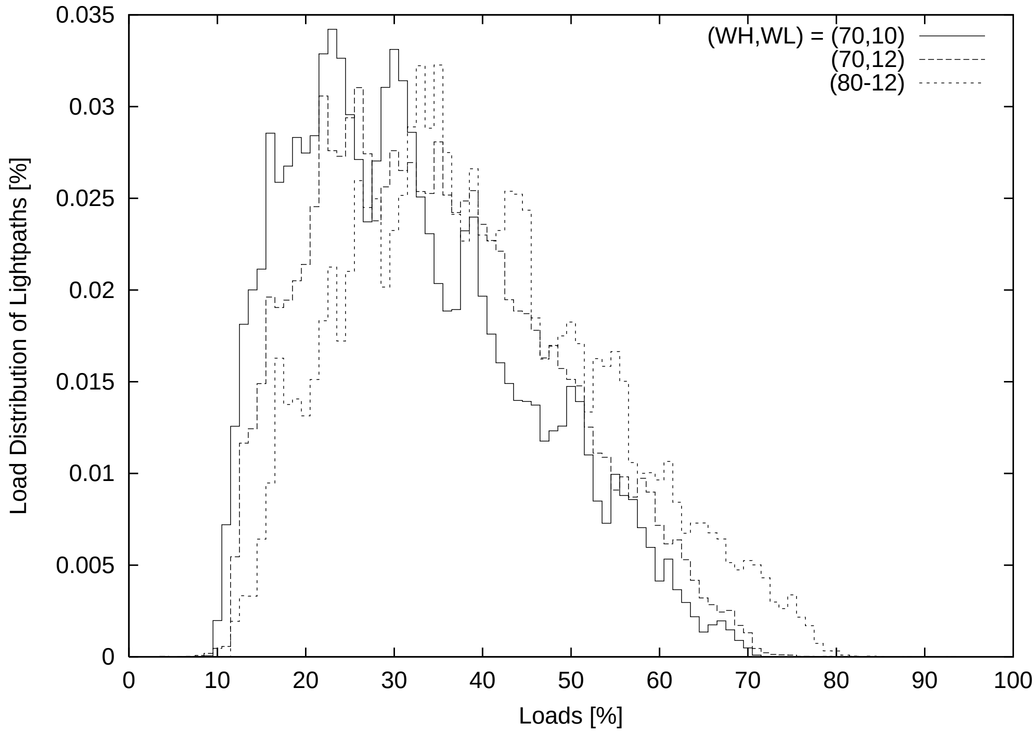

Figure 4.11: Distribution of lightpath loads for different values of watermarks over the course of the simulated five days.

Figure 4.11 plots the distribution of lightpaths according to their loads, for three different watermark-value pairs. At the end of each observation period (300 sec. for this example), the loads of all lightpaths in the network are measured and the occurrence of each load is added to the total occurrences of that load since the beginning of the simulation. This process is repeated at the end of every observation period throughout the simulation. A point in this graphic indicates the percentage of the lightpaths with that load in the course of the experiment (a 5-day run). The figure shows that a high percentage of the links are in the balance region. Another result is that most of the lightpaths are gathered in a smaller region toward the middle of the balance region, showing that few lightpaths are critically close to the watermarks, and they can trigger a topology adjustment in the following observation period. The same experiment is repeated for different values of W_H and W_L, and the percentage of the lightpaths in the balance region is tabulated in Table 4.4. As W_H and W_L values are relaxed, the percentage of lightpaths in the balance region increases. This table shows that the watermarks can be adjusted to obtain a more/less aggressive virtual topology adaptation.

Table 4.4: Percentage of the lightpath loads between the watermarks over the course of simulation.

W_H	W_L			
	8	10	12	14
60	99.9976	99.9958	99.9932	99.9882
65	99.9995	99.9986	99.9964	99.9919
70	99.9999	99.9991	99.9971	99.9949
75	99.9999	99.9991	99.9987	99.9962
80	99.9999	99.9997	99.9994	99.9984

4.4.7 Unlimited Additions and Deletions

The aim of this section is to show a variation of the basic algorithm in Figure 4.1, where more than one lightpath change is allowed at each observation period. This variation of the method provides information on how many lightpath additions or deletions would be necessary at every observation period (whose lenght is 300 seconds in this experiment) to balance the lightpath loads. The algorithm is modified to allow unlimited number of additions or deletions at each step. The only limitation on lightpath changes is that one type of change, i.e. either additions or deletions, can be performed in a specific observation period. According to the new algorithm, when a lightpath is added or deleted, the traffic is redistributed and the lightpath loads are re-calculated. If there is still an imbalance in loads, an additional change is made immediately, without waiting for the next observation period.

Figure 4.12 plots the distribution of the number of additions and deletions in one observation period for $W_H = 70$ and $W_L = 10$. The percentage is calculated over the total number of periods where at least one change is made. For example, the percentage of observation periods where only one addition was made is 87% of all observation periods where at least one addition was made. This figure shows that one change in a period is sufficient to obtain a good virtual topology in the majority of adaptation cases. Also note that, a virtual-topology change is seldom, i.e., in most of the observation periods no change is made to the virtual topology as can also be observed in Fig. 4.3. The percentage of such *silent* observation periods is 94% for the specific numerical example in Fig. 4.12. These results show that the basic adaptation algorithm, where only one lightpath change at each step is allowed, is an efficient reconfiguration method that can balance the virtual topology quickly enough.

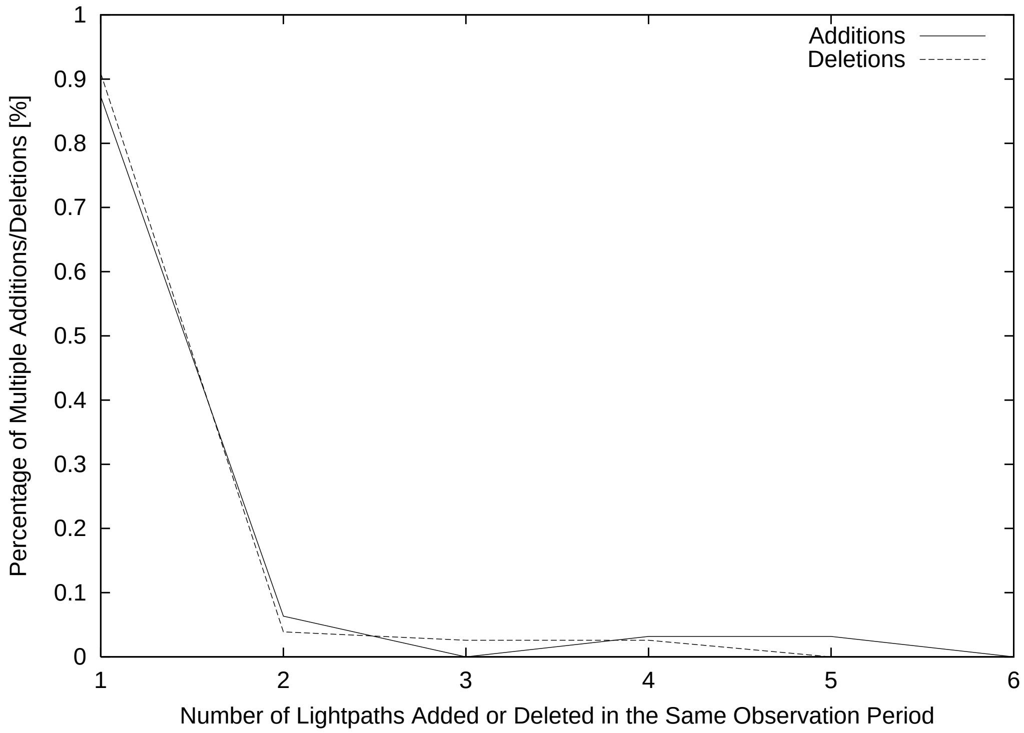

Figure 4.12: Distribution of number of lightpath additions and deletions when the number of additions and deletions are not limited.

4.5 Comparison with Local Minimum

The adaptation algorithm in Figure 4.1 selects a lightpath to be added or to be deleted at each step, by using a heuristic selection mechanism. This selection may be performed by also searching through the space of all possible lightpaths. The size of the search space is N^2 and τN for addition and deletion respectively. Therefore, a local minimisation method can be designed to choose the minimum possible maximum lightpath load in the network, by searching through the space of lightpaths. This section compares the basic adaptation algorithm with this local minimisation method, to reveal the performance of the heuristic selection decision.

Figure 4.13 shows the maximum lightpath load at the end of each observation period (300 s for this experiment), for the adaptation algorithm and the algorithm that selects the local minimum. The resulting graphic shows that the heuristic decision used by the adaptation method gives very close solutions to the local minimum, and the two methods make same changes most of the time. This observation is valid for other values of the watermarks also as can be seen in Figures 4.14 and 4.15. Hence, it is fair

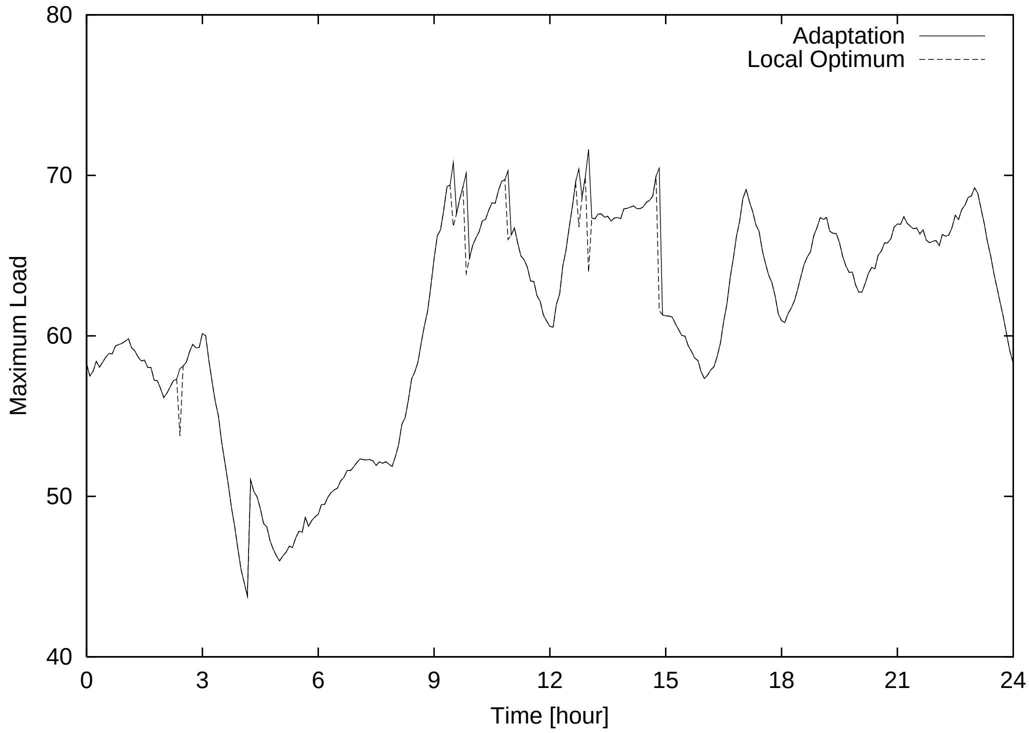

Figure 4.13: The comparison of adaptation and local minimum for maximum lightpath load during one day, for $W_H = 70$ and $W_L = 10$.

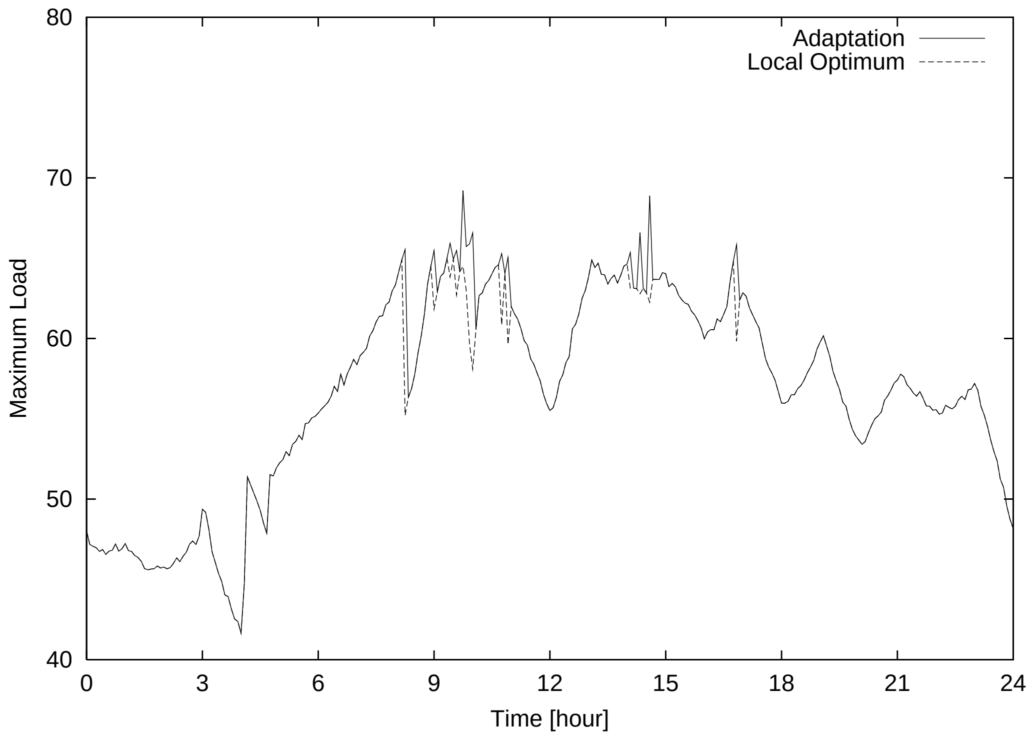

Figure 4.14: The comparison of adaptation and local minimum for maximum lightpath load during one day, for $W_H = 65$ and $W_L = 10$.

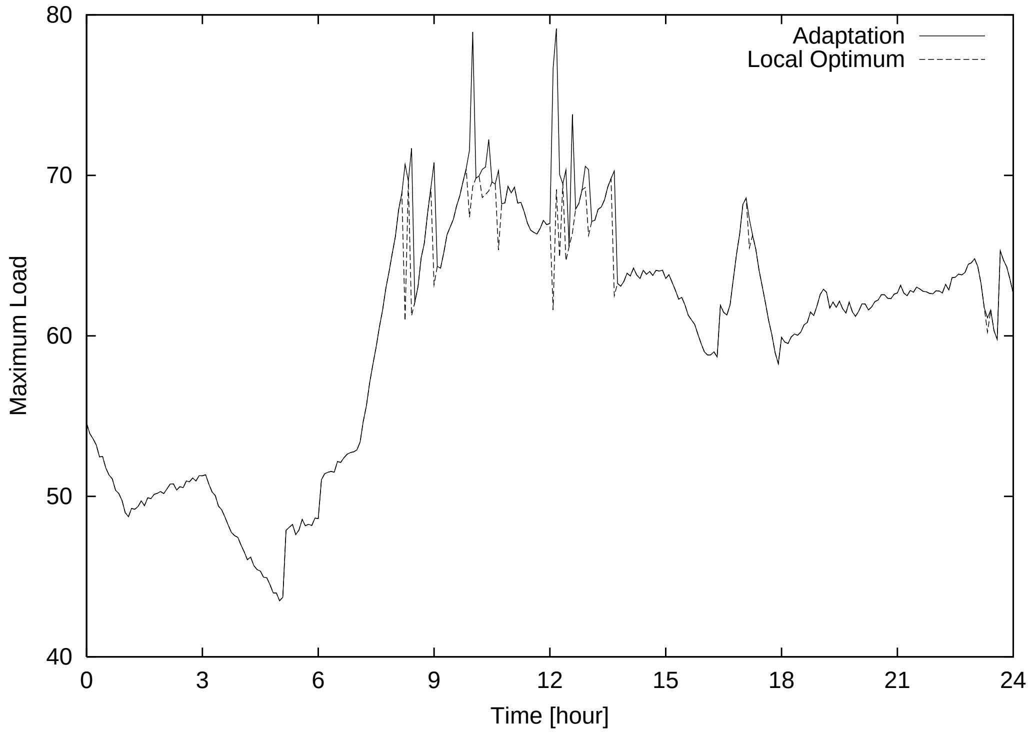

Figure 4.15: The comparison of adaptation and local minimum for maximum lightpath load during one day, for $W_H = 70$ and $W_L = 12$.

to conclude that the heuristic lightpath selection decision proves to be very close to the local minimum, while its complexity is τN times lower.

4.6 Comparison with Virtual Topology Design Methods

In this section, two virtual topology design algorithms are used to evaluate the performance of the adaptation algorithm: SMLTDA and GLTDA. These methods are chosen because of their relative computational simplicity among the virtual topology design methods, and their average performance compared to other design heuristics [92].

The comparison is made as follows. For a simulation time of 24 hours, the adaptation method is allowed to adjust the virtual topology every five minutes (observation period). In accordance with this, the selected design method is run each five-minute interval with the same traffic demand at that moment. Hence, the virtual topology design method is expected to create a new virtual topology each five-minutes without considering the number of changes to be made for transition from the previous topology, but aiming at evenly distributing the traffic loads on lightpaths.

The comparison metric used is the maximum and minimum lightpath loads in the network at each observation period. The total number of lightpaths in the network is a parameter that affects the performance of the algorithms.

4.6.1 Single-hop Maximization Logical Topology Design Algorithm (SMLTDA)

Define, for each lightpath that can be included in the virtual topology, the binary variables $b_{ij} \epsilon \{0, 1\}$, such that $b_{ij} = 1$ if the lightpath from node i to node j belongs to the virtual topology, and $b_{ij} = 0$ otherwise. Assuming that the traffic intensities between nodes is given by values t^{ij}, find the set of b_{ij} that maximises

$$f = \sum_{i} \sum_{j} b_{ij} \cdot t^{ij} \tag{4.1}$$

where t^{ij} is the traffic load from node i to node j. The variables b_{ij} are subject to

$$\forall i \sum_{j} b_{ij} \leq \delta_O^i \tag{4.2}$$

$$\forall i \sum_{i} b_{ij} \leq \delta_I^i \tag{4.3}$$

where δ_O^i is the numberof receivers and δ_I^i is the number of transmitters at node i.

This problem can be solved by picking up the largest t^{ij} in the traffic matrix and establishing a lightpath if possible. The algorithm will stop when the network resources are used up, i.e., no more lightpaths can be established. The time complexity of this method is $O(\delta N^3 \log N)$, where δ is assumed to be the number of transceiver pairs at a node.

4.6.2 Greedy Logical Topology Design Algorithm (GLTDA)

This algorithm is a greedy matching method that allows the establishment of multiple parallel lightpaths attempting to introduce connections between nodes that exchange large amounts of traffic. The pseudocode of GLTDA is given in Figure 4.16. The notation used in this figure is similar to the one used for SMLTDA.

Algorithm:

```
S1. Find  $t^{sd}_{max} = \max_{ij}(t^{ij})$  and  $t^{s'd'}_{max} = \max_{i \neq s \vee j \neq d}(t^{ij})$
S2. If  $t^{sd}_{max} = 0$ then
        Go to S4
    Endif
S3. If  $\delta^s_O \neq 0$  and  $\delta^d_I \neq 0$  then
        Create a lightpath between  $s$  and  $d$
        Decrement  $\delta^s_O$  and  $\delta^d_I$
        Update $t^{sd} = t^{sd} - t^{s'd'}_{max}$
        Go to S1
    Else
        Let  $t^{sd} = 0$
        Go to S1
    Endif
S4. Place required number of lightpaths at random, without
    violating the degree and wavelength constraints
```

Figure 4.16: Pseudocode of GLTDA.

4.6.3 Results

The results of the experiments explained above are shown in Table 4.5. First three lines of the table shows the maximum and minimum lightpath loads over 24-hour period, by the adaptation method. The results are given for three different values of high and low watermarks. The last column of the table shows the maximum number of lightpaths established at a specific point in time, during one day.

Over the cours of 24-hour period, i.e., 288 different runs of five minutes apart, the virtual topology design methods SMLTDA and GLTDA design the virtual topologies that include as many lightpaths as possible, since they stop only when all the network resources are used up. The results obtained when these algorithms are run with this termination condition, are marked in the table as *no limit* (Last two lines).

Both SMLTDA and GLTDA may result in disconnected virtual topologies if the number of lightpaths established are small. On the other hand, for a clearer understanding, the number of lightpaths in the virtual topologies designed by these methods are limited in a second experiment. Here, the aim is to obtain connected topologies for all 288 intervals of the day, with a limited percent increase over the number of lightpaths in the topologies created by the adaptation method. As a results, SMLTDA and GLTDA could create connected virtual topologies for all intervals, when they are allowed to establish 80% more lightpaths compared to the adaptation method

Table 4.5: Comparison of lightpath loads for different algorithms.

Extra lightpaths	Algorithm	Maximum Load	Minimum Load	Number of Lightpaths
0%	Adaptation (70,10)	71.61	8.91	66
	Adaptation (60,10)	64.45	8.20	68
	Adaptation (60,8)	61.92	7.71	72
80%	SMLTDA	46.59	2.82	119
	GLTDA	65.78	2.72	119
No Limit	SMLTDA	37.29	0.26	147
	GLTDA	60.20	0.55	150

(with $W_H = 70$ and $W_L = 10$). The topologies they provided for smaller number of lightpaths were in parts either disconnected or unable to accomodate the given traffic. Therefore, the first feasible solution with smallest number of lightpaths are for the case 80% extra lightpaths, and the results are shown in the fourth and fifth lines of the table.

The maximum loads given in the third column of the table is the maximum lightpath load among all virtual topologies designed during the day. The minimum load represents similarly the minimum lightpath load found during a day. According to the results in the table, the adaptation method provides virtual topologies with good load balancing, and with very small number of lightpaths. Both design methods are unable to provide solutions with so few lightpath as in the adaptation method. Their feasible solutions start at virtual topologies of 80% more lightpaths. Hence, the maximum loads obtained by using SMLTDA are lower than the adaptation method. GMLTDA is outperformed by SMLTDA in these comparisons, and adaptation with $W_H = 60$ and $W_L = 8$ has better performance compared to GLTDA 80% and very close performance compared to GLTDA without limit.

These experiments show that the adaptation method has advantages over even the virtual topology design methods. One would expect that the topologies designed anew, each time a reconfiguration is made, to have a better performance than any reconfigured topology. However, simple heuristics for topology design establish to many lightpaths with very low utilisation. Since there is no correlation between consecutive topologies, the performance of these topologies may be very different from each other. On the other hand, even though the adaptation makes only one change each time, it provides a virtual topology that is evolving in long period of time, and whose performance is more stable as can be seen in Figure 4.17. This figure plots the maximum and

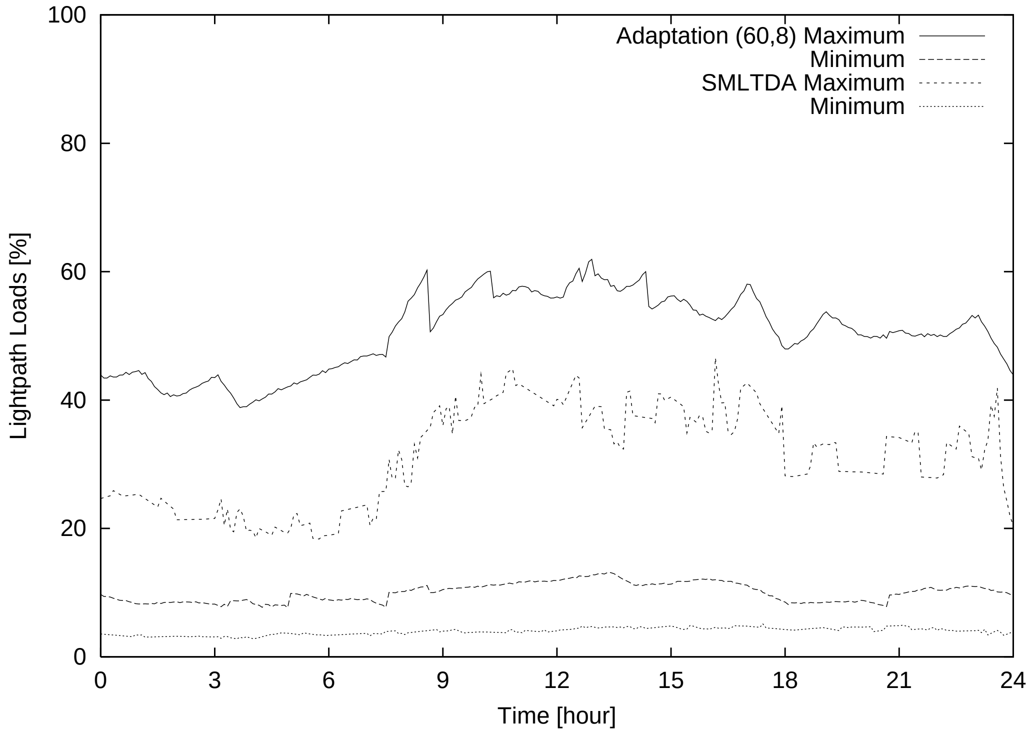

Figure 4.17: The comparison of maximum and minimum lightpath loads obtained by the adaptation and SMLTDA.

minimum lightpath loads on virtual topologies designed by adaptation and SMLTDA (whose results are better than GLTDA's) over the course of one day.

One last remark is that the virtual topology methods are difficult to use for reconfiguration, since they do not consider the number of changes necessary to reach the new topology starting from the previous one, and eventually they require high number of changes each time the topology is designed. Such comparison is not used as a performance metric here, since the aim is to show the quality of results obtained by the adaptation method in comparison with virtual topology design heuristics.

4.7 Conclusion

Based on a new view of the virtual-topology reconfiguration problem, this chapter presented an adaptation scheme for WDM mesh networks under dynamic traffic. The problem was defined as tracking the long-term traffic fluctuations by adapting the topology in a measurement-adaptation cycle with the following constraints: No assumption should be made on future traffic rates and the ongoing traffic should not be interrupted by a transition phase. This work is motivated by the traffic characteristics of backbone networks, which change slowly and according to the time-of-the-day.

In this chapter, a heuristic algorithm was developed, which adjusts the virtual topology by adding or deleting one lightpath at the end of a measurement cycle, if it is necessary. Two system parameters were introduced, called high watermark (W_H) and low watermark (W_L), on lightpath loads to detect any over- or under-utilized lightpath. This chapter showed through simulations that these two parameters can be effectively used to regulate the operation of the proposed topology-adaptation method. The effect of the length of observation period was also shown, and the load balancing property of the algorithm was illustrated. Two variations of the adaptation method were implemented. In the first variation, the traffic loads were measured according a history window whose size could be multiple observation periods. In the second variation, an unlimited number of changes were allowed at every step, and the simulation results showed that one change in one observation period was enough to maintain the loads to be balanced most of the time.

5. COST-EFFICIENT BANDWIDTH ALLOCATION AND RECONFIGURATION IN A WORLD-WIDE WDM NETWORK

This chapter investigates cost-efficient methods for bandwidth allocation and bandwidth reconfiguration for global IP-over-WDM backbone networks. The organisation of the chapter is as follows: Background information, emerging business models and bandwidth service types are given in Section 5.1. The motivation of bandwidth reconfiguration is explained in Section 5.2. In Sections 5.3 and 5.4 the network model and the traffic model used in this work are introduced. New bandwidth assignment and reconfiguration schemes are presented in Section 5.5, and the numerical examples and their results are discussed in Section 5.6.

5.1 Introduction

Today, bandwidth is rapidly becoming a commodity that can be sold, purchased, or leased on a global scale. Bandwidth markets have started to appear [93, 94], and they are expected to mature quickly into markets similar to today's electricity markets, where even ten-minute blocks are priced, traded, and settled [95]. In this perspective bandwidth markets will allow service providers (or large institutional users of bandwidth) to drastically reduce their capital and operational expenditure, by employing on-demand bandwidth provisioning.

In today's networks, topologies and capacities have to be planned well in advance, because bandwidth provisioning can take weeks, or even months. With on-line bandwidth provisioning, service providers will be able to purchase or lease bandwidth in a matter of seconds or less, and also to end a lease quickly, instead of setting up their own fibre infrastructure.

Bandwidth markets provide the opportunity to choose among different prices available from several bandwidth providers. On-demand provisioning provides flexibility of upgrading or changing connections and their bandwidths. A third opportunity has already started to be provided by many global carriers: Different service types can be chosen for a communication channel. Among several service types, *flat-rate*, *tiered*, and *burstable* are of interest for this study. A flat-rate channel is priced with a constant value based on the channel's bandwidth. (In optical networks channel bandwidths are scaled based on OC (optical channel) units, where OC-1 is approximately 51.84 Mb/s.) A tiered channel's bandwidth is capped in increments to limit bandwidth usage. The customer can upgrade to another bandwidth tier at any time as long as it is below the major bandwidth level (e.g. OC-12, OC-48, etc.) chosen at the beginning of the lease. In this model the price of the channel changes when the bandwidth tier changes. A burstable channel is the most flexible among the three models, and the price is based on the average usage of the customer. The channel's bandwidth usage is monitored during a time interval (e.g. a month), the highest 5 percent of the traffic rate samples are discarded to prevent traffic burst affect the average, and the price is based on the average usage of the 95 percent of the samples.

Together with the advances of technologies, markets and pricing models, new business approaches are evolving. These approaches can be grouped into four business models [96]. The following three are of interest in this study.

National/Global Carriers: Their objective is to provide high bandwidth circuits through their fibre infrastructure, and to offer services to Internet Service Providers (ISPs).

Internet Service Providers: They provide services to end users (whether at a scale of a person or an institution), by leasing wavelength or sub-wavelength granularity channels from Global Carriers, and flexibly managing their backbone to resell the bandwidth with added values (QoS, security, applications for end users, etc.).

Carrier-Neutral Inter-Exchange (IX) Points or Carrier Hotels: They provide neutral collocation space to Global Carriers and ISPs to lease, trade, buy, and sell bandwidth on demand.

The opportunities cited above pose new challenges in operating the network cost-efficiently, since there are several decisions to be made that can affect the operational expenses. With the growth of data traffic, customer's traffic requirements are becoming increasingly dynamic in time. Therefore, the challenge is not only designing the network, but also continuously following the bandwidth markets, updating and reconfiguring the network connections, to maximise the profit.

In this study, we investigate the reconfiguration problem in an IP-over-WDM network from the point of view of an ISP under the light of these new challenges. In most of previous work on virtual topology design and bandwidth assignment, the problem was considered as embedding an appropriate virtual topology on an existing physical fibre layout. In such a view, the business model is a combination of a Global Carrier and an ISP, where it owns fibre infrastructure and optical cross-connects, and provides network services to the end users. Such a model does not consider the emerging bandwidth markets and carrier hotels, and the business is limited by its physical layer topology as well as physical layer constraints (number of wavelengths, wavelength continuity, etc.).

Previous studies on virtual topology design concentrate on routing the IP traffic appropriately, so that the lightpath connectivity can carry a large amount of traffic. These studies consider a constant capacity for any lightpath and do not take into account the sub-wavelength bandwidth granularities. They also contravene the objective of the ISP business model defined above, to minimise the network cost for high profit, because they try to set up as many lightpaths as possible, instead of an appropriate and cost-effective working set of lightpaths. Network reconfiguration studies are similar to design studies in both aspects, and their application to future real networks is difficult.

This chapter investigates cost-efficient methods for bandwidth allocation and bandwidth reconfiguration for global IP-over-WDM backbone networks. IP traffic is very dynamic, but in a backbone network, due to traffic aggregation, this dynamism becomes smoother, following a well-shaped daily profile (Figure 3.1). This property can be exploited by WDM technology, since wavelength-routed WDM networks provide a flexible infrastructure that can adjust the network topology (connectivity and bandwidth) by reconfiguring the optical cross-connects. On the other hand, when a world-wide network traffic structure is examined, another dimension of dynamism can be seen: traffic shifts based on the time zones. A typical traffic profile between two nodes follow a sinusoidal shape, where the peak is around midday and the valley around early hours of the morning, e.g. 03.00. Since different locations in the world are in different time zones, each node will create a traffic whose intensity changes according to that node's time zone. This effect is clearly observed on long span connections, such as London - New York connection, where the traffic intensities in two directions are asymmetric, because of five hours time difference [97]. When this new type of dynamism is considered, the bandwidth cost of the network can be reduced, even in the design step of the virtual topology. In this chapter, we show that a large reduction on bandwidth cost can be achieved by reconfiguring the lightpath capacities.

5.2 Motivation

To date, topology-design algorithms have considered static traffic demands, which are the maximum traffic intensities observed or expected between node pairs. When virtual topology is designed using these value, the allocated lightpaths and the capacities are expected to accommodate the maximum amount of traffic. Since the traffic fluctuates during the day, and the maximum traffic demand occurs a limited amount of time, such as few hours, the connection capacities are over-allocated.

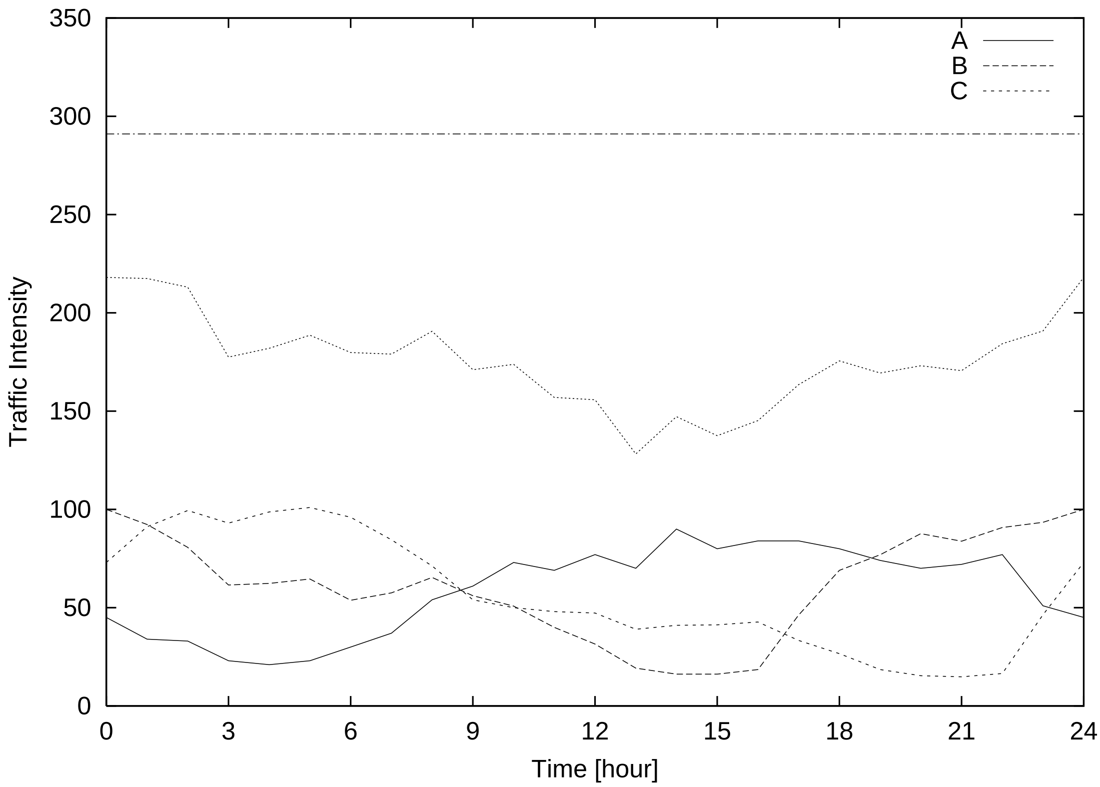

Figure 5.1: Illustration of the intensities of three traffic flows, and their total intensity during a day.

In a global network, this over-allocation is more severe. The following example illustrates this problem. Figure 5.1 shows three traffic flows with destinations A (in time zone +2), B (in time zone -5) and C (in time zone +11). Assuming that in the virtual topology, these flows are routed through the same lightpath, the total traffic flowing through this lightpath is also shown in the figure. At the design step of the virtual topology, if the daily traffic fluctuations with the time shifts are considered, the total bandwidth assigned to this connection would be the maximum value of this total curve which is 218 unit in this example. On the other hand, if the maximum values of the three profiles are considered separately, without taking into account the time zone based shifts, the bandwidth assigned will be 291 unit, the total of the peak traffic intensities of the three flows. This numerical illustration shows that traffic dynamism can be exploited to assign the connection bandwidths cost-effectively.

When on-line bandwidth provisioning becomes commonplace, it is possible for an ISP to operate its network more cost-effectively, using bandwidth reconfiguration. This is a different approach when it is compared to virtual topology reconfiguration, because the virtual connectivity between nodes remains the same, while the allocated bandwidths of the connections change. In this approach the traffic flowing through the connections is observed, and the connection's capacity is increased or decreased according to the observed intensities.

5.3 Network Model

In this study, a world-wide fibre backbone topology with 20 nodes and 47 fibre-links is used (Figure 5.2). This topology has been created by combining and simplifying several real backbone topologies of global carriers. This topology represents the physical infrastructure that is available to the ISP for creating its virtual topology. In that sense, it may be a combination of available physical paths through a bandwidth market. Length of a fibre-link is assigned as the air distance of the two cities connected by that link. Each network node is assumed to have 8 transceiver pairs. The number of wavelengths per fibre is not assumed to be a constraint since the physical topology does not belong to the ISP. For illustration purposes, the cost of a connection of unit capacity per km per hour is assumed to be the same throughout the network, although this may not be the case in bandwidth markets. Each wavelength channel has a maximum capacity of OC-192, while sub-capacities can also be assigned at the granularity of OC-3, OC-12, or OC-48.

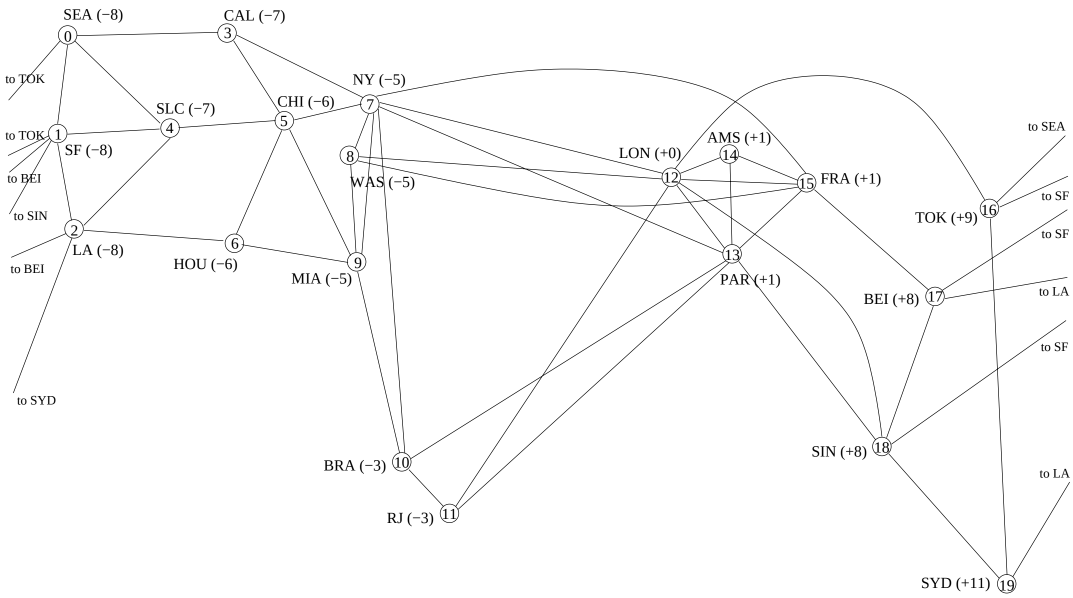

Figure 5.2: Global topology used in this study. The numbers beside the city names show the time zone that the city belongs.

Each node in the network is equipped with electronic routers, and optical cross-connects that can route the multi-hop optical connections. The bandwidth-reconfiguration concept is based on the property that next-generation switch architectures will be capable of on-line bandwidth provisioning.

5.4 Traffic Model

The traffic model used in this study reflects both time-of-the-day dynamism where traffic intensity between nodes increase at busy hours of the day, and time-zone dynamism where traffic intensities increase or decrease at different locations of the world according to the local time. For each communicating node pair, a daily traffic profile is chosen randomly among several sampled backbone traffic profiles. The traffic intensities for the node pair is then scaled by a factor that is a multiplication of the two populations and the country's development factor. Each city represents a population of the geographic area that the city is situated. For example, four European cities (Paris, London, Amsterdam and Frankfurt) represent the population of the whole European continent.

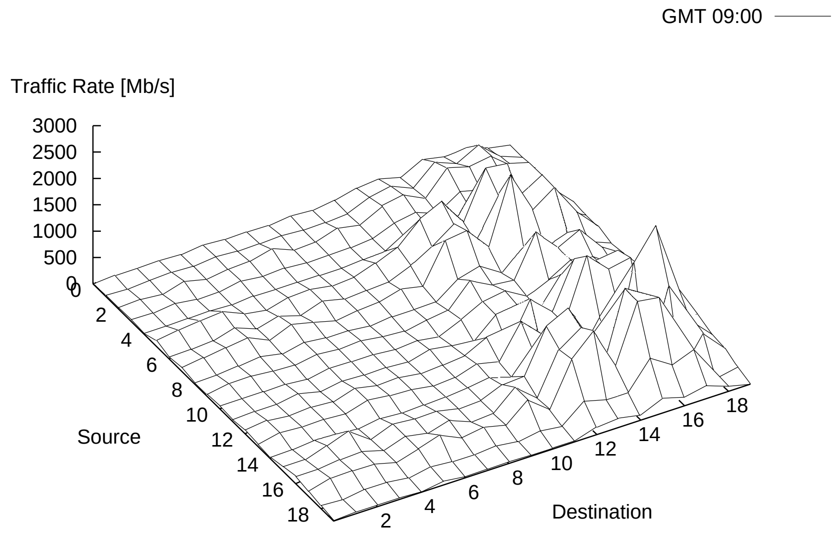

Figure 5.3: Traffic matrix at GMT 09.00.

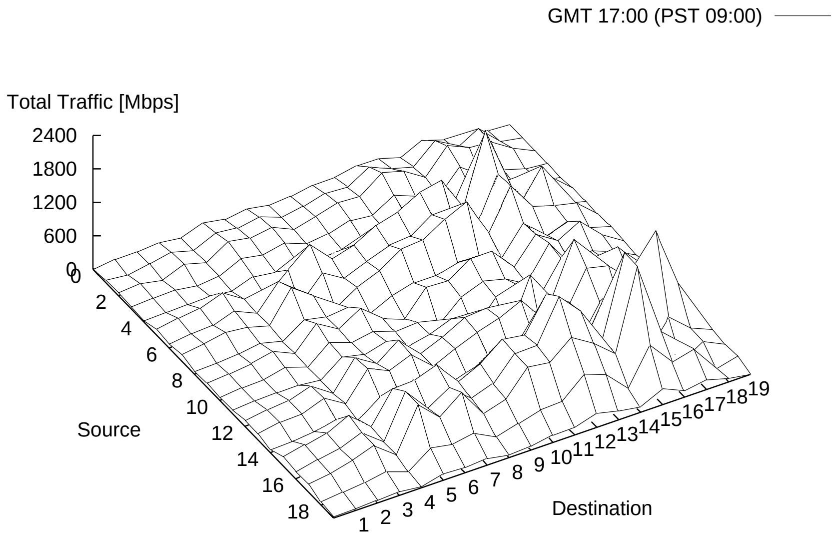

Figure 5.4: Traffic matrix at GMT 17.00.

To simulate the time zone effect, the traffic profiles are shifted by the time zone of the destination city. An example of this effect is shown in Figure 5.3, where the traffic matrix at 09:00 GMT is shown as a surface plot. According to this plot, traffic intensities in Europe (nodes 12, 13, 14, and 15) start to increase while in North America (nodes 0 through 9), traffic intensities remain low. At 17:00 GMT (09:00 PST), business hours start in all over America, which can be seen in Figure 5.4, increasing the traffic at North American nodes.

Figure 5.5 shows a traffic matrix at different hours of the day, namely, each three hours starting from midnight for one day. The changes in traffic intensities according to the location of the nodes and the time of the day can be seen clearly in this figure.

For the numerical examples, three different traffic matrices are created randomly using the method above. Figure 5.5 illustrates one of these matrices. The other two matrices are illustrated in Appendix C.

5.5 Bandwidth Assignment Schemes

The new bandwidth reconfiguration method called *CATZ (Capacity Assignment Using Time-Zones)* is presented in this section. For illustrating the performance of the dynamic method, two static bandwidth allocation methods are also introduced.

5.5.1 Static Bandwidth Assignment

Two static bandwidth assignment methods used in this study are *static-max* and *static-daily*, which are explained below.

Static-Max: Traditional topology-design and bandwidth-allocation methods do not consider the hourly changes in traffic intensities. Instead, they allocate bandwidths based on peak expected traffic intensity between node pairs. The virtual topology is designed, and the bandwidths of the channels are assigned based on a traffic matrix where the traffic between each node pair is represented by one value, which is the maximum intensity expected. This method is a static bandwidth allocation scheme, since the bandwidths of the channels are not expected to change for a long period of time, until the traffic intensities necessitate re-designing the virtual topology. This method may lead to over-allocation of bandwidths in a real backbone network, since the traffic fluctuates according to the time-of-the-day and time zones.

Static-Daily: A straightforward improvement would allocate the bandwidths of the connections by routing the actual traffic between nodes instead of the peak traffic values, during a 24-hour period. Periodically, e.g. every 1 hour, the bandwidth of a channel can be calculated by observing the total traffic to be routed on that channel.

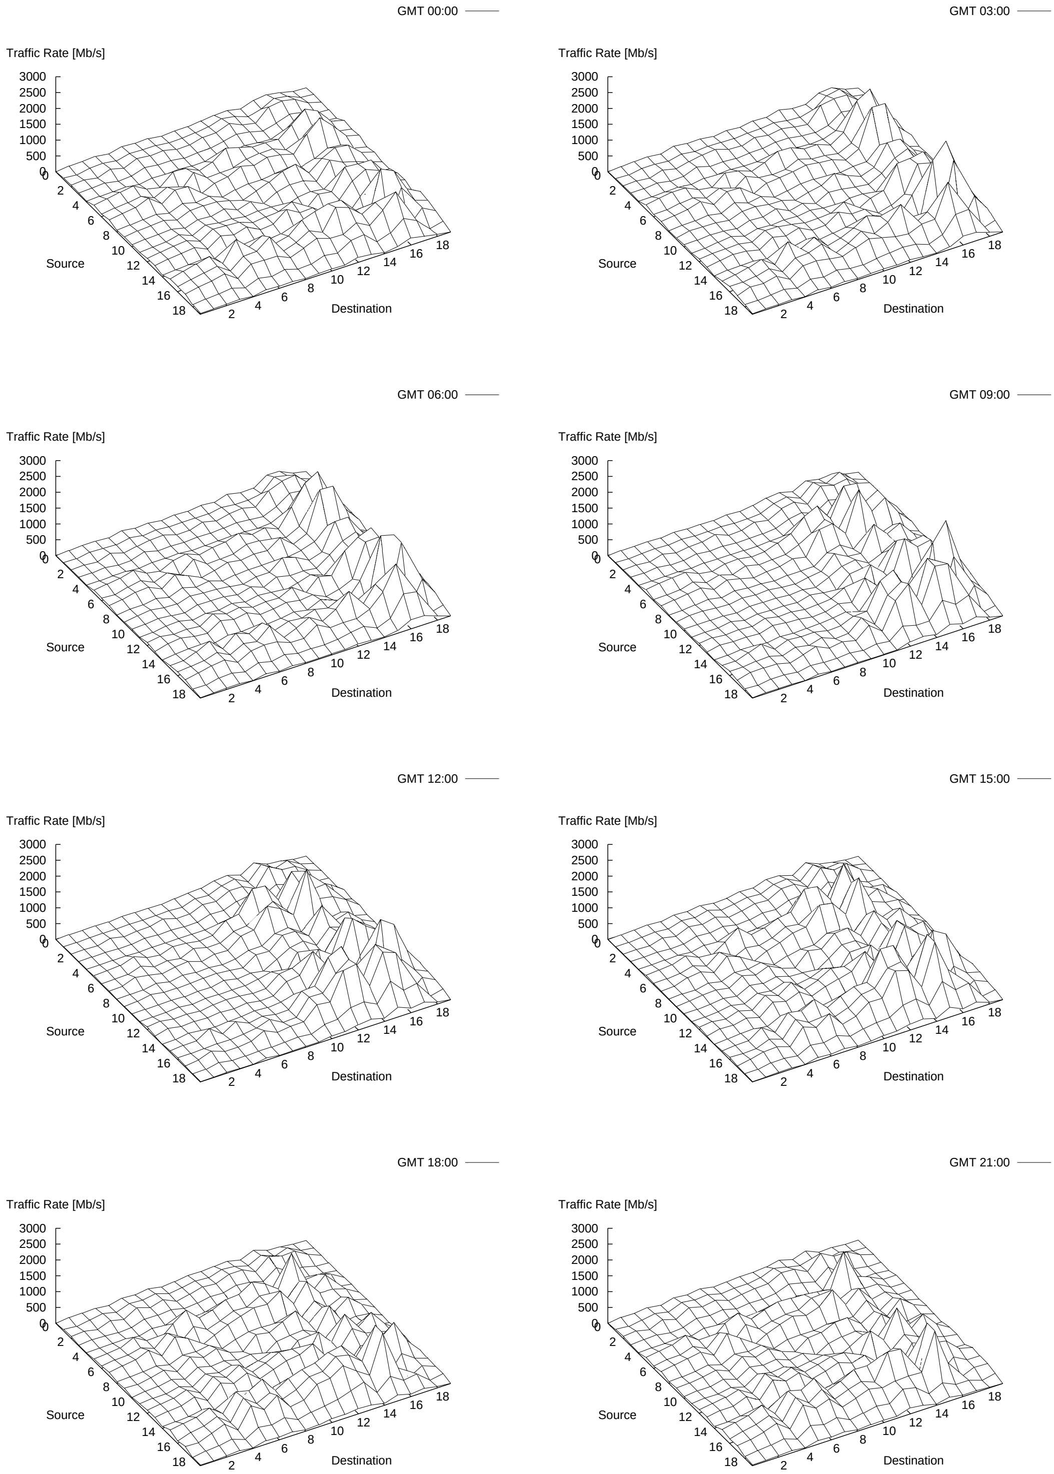

Figure 5.5: Traffic matrices at different hours of the day.

At the end of the day, the bandwidth of a channel can be selected as the maximum allocated bandwidth to that channel during the day. This method is also static and based on an expected traffic matrix, however, it considers the dynamic changes in traffic intensities. The bandwidths are calculated off-line and they cannot be changed during the network operation.

5.5.2 Dynamic Bandwidth Assignment - CATZ

The objective of the dynamic bandwidth assignment is to exploit the dynamism of the traffic to reconfigure the bandwidths of connections when necessary. The proposed method increases or decreases the bandwidth of a connection when total traffic flowing through it increases or decreases.

The main proposal for cost-effective bandwidth assignment (CATZ) exploits the on-line bandwidth provisioning capability of the network. The channels are observed periodically to measure the traffic intensities flowing through them, and the bandwidth of the channel for the next period is set according to these intensities. In this method, the capacities of the channels are allowed to change, since the underlying transport system, optical switches, and the global carriers are assumed to provide on-line bandwidth services. The amount of step increase or decrease can be a multiple of OC-1, since the bandwidth markets provide optical channels at capacities of OC-3, OC-12, OC-48, or OC-192. Parallel channels are also possible between nodes, as long as there are enough transceivers. On the other hand, step increase/decrease in channel capacities are also possible through burstable service type, as explained in Section 5.1. The pseudocode of CATZ is given in Figure 5.6.

5.6 Illustrative Numerical Examples

In this section, the performance of the two static bandwidth allocation methods and CATZ are compared in term of the network's total bandwidth cost. The time interval for periodical measurements is 1 hour in methods Static-Daily and CATZ. The network cost is calculated for one day using the three methods. The resulting cost graphic is shown in Figure 5.7. The horizontal axis shows the time of the day in hours, and the vertical axis represents the total bandwidth cost of the whole network, in terms of \$1000. Note that this amount is a representative value, rather than the real prices in the bandwidth markets. The total bandwidth cost for the network is calculated as the total

Input:

- Capacity of connections, C_{sd}
- Maximum capacity of connections, C_{sd}^{max}
- Current traffic rates, λ_{sd}
- Allowed granularity of step increase/decrease, Δ
- The minimum allowed bandwidth gap between C_{sd} and λ_{sd}, ε

Algorithm:
At the end of every observation period:

```
For  each node pair  <s,d>  do
     if  λ_sd > C_sd − ε  then
         if  C_sd < C_sd^max  then
             C_sd ← C_sd + Δ
         endif
     else
         if λ_sd < C_sd − Δ − ε  then
             C_sd ← C_sd − Δ
         endif
     endif
```

Figure 5.6: Pseudocode of CATZ.

cost of every channel where the channel cost is the product of bandwidth and distance unit price (unit price = 0.0075/km/hour/OC3 in the examples).

In Figure 5.7, one can observe the cost improvement by using Static-Daily method instead of Static-Max. Although the improvement is 8%, it shows the advantage of considering the traffic dynamism during the design step of the network. If the result curve of CATZ examined, this figure shows that when the bandwidths are changed according to the daily traffic fluctuations, the cost can be reduced dramatically. The total bandwidth cost for one day in this illustrative example is reduced by approximately 36% compared to maximum-traffic-based assignment of the Static-Max method. The hourly cost of the network changes dynamically due to the bandwidth increase or decrease on connections, achieving a peak at around 16:00 GMT where majority of the world population is creating traffic load simultaneously in the network. Note that the cost changes are smoother than the changes in traffic, due to the aggregation with time-shifting of several traffic flows on a typical connection (This aggregation effect has been illustrated in Figure 5.1).

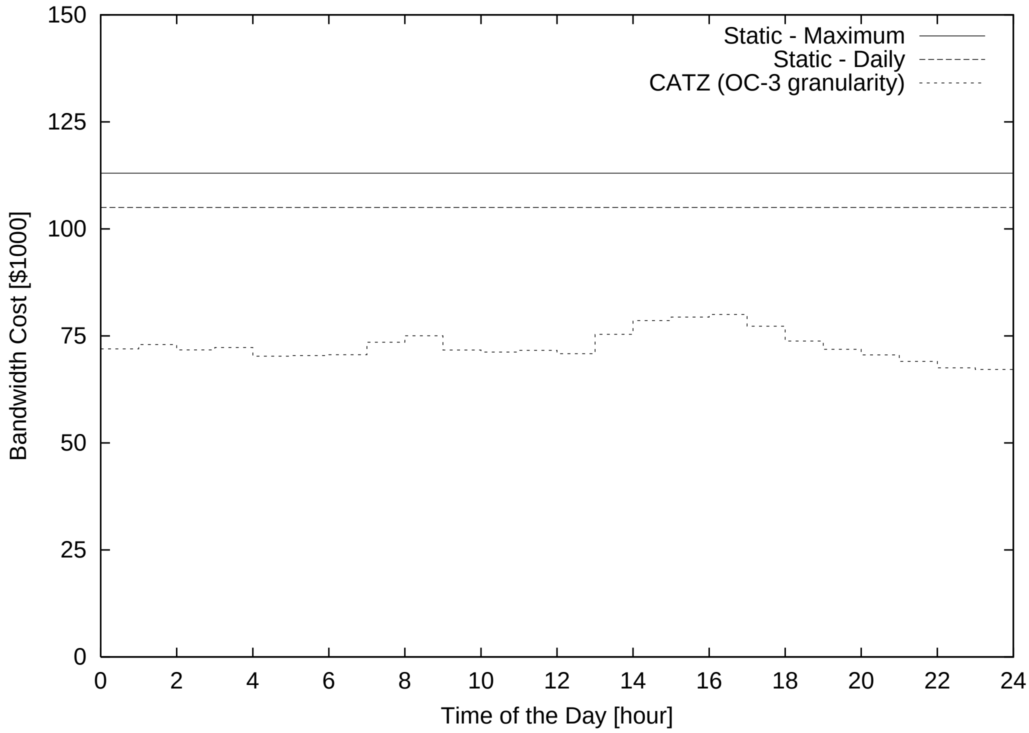

Figure 5.7: Bandwidth costs of different methods.

Figure 5.8 plots the bandwidth costs, for different channel bandwidth increase/decrease granularities. The increase or decrease in channel bandwidths are allowed to be in steps, and three different step values are used in this experiment: OC-3, OC-12, and OC-48. As expected, when fine-granularity bandwidth adjustments are possible, the cost is lower, compared to coarse-granularity case (OC-48). (Comparing Figures 5.7 and 5.8, one can observe that the network cost of CATZ with OC-48 bandwidth steps, is higher than Static-Max. The reason is that in both static methods OC-3 bandwidth steps are used in the design step.) This result motivates the use of different bandwidth granularities in a network, as well as, the use of traffic-sensitive services, such as burstable service.

The experiment is repeated for other two traffic matrices, and the results are plotted in Figures 5.9 and 5.10. (The bandwidth steps are OC-3.) Similar results are obtained in these experiments as can be observed in the figures. The exact cost values for the three different traffic matrices are presented in Table 5.1. The values shown for CATZ in the fourth column are the average value over 24 hours. The fifth column of this table shows the percentage improvement in cost obtained by using CATZ, instead of Static-Max.

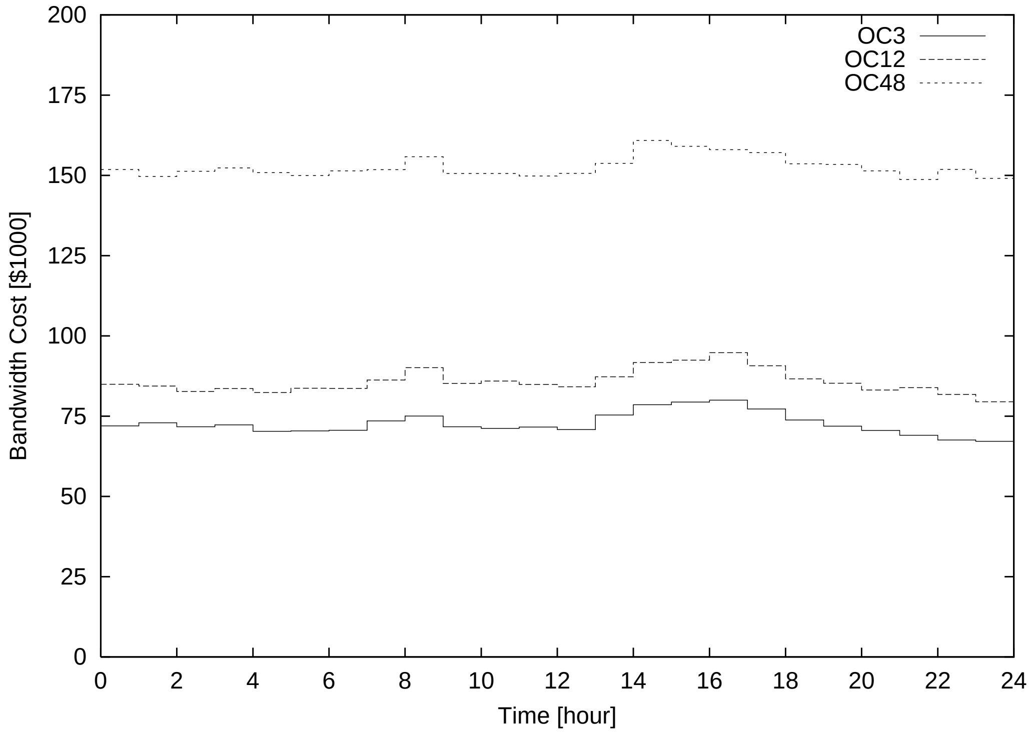

Figure 5.8: Bandwidth costs for different channel granularities with CATZ.

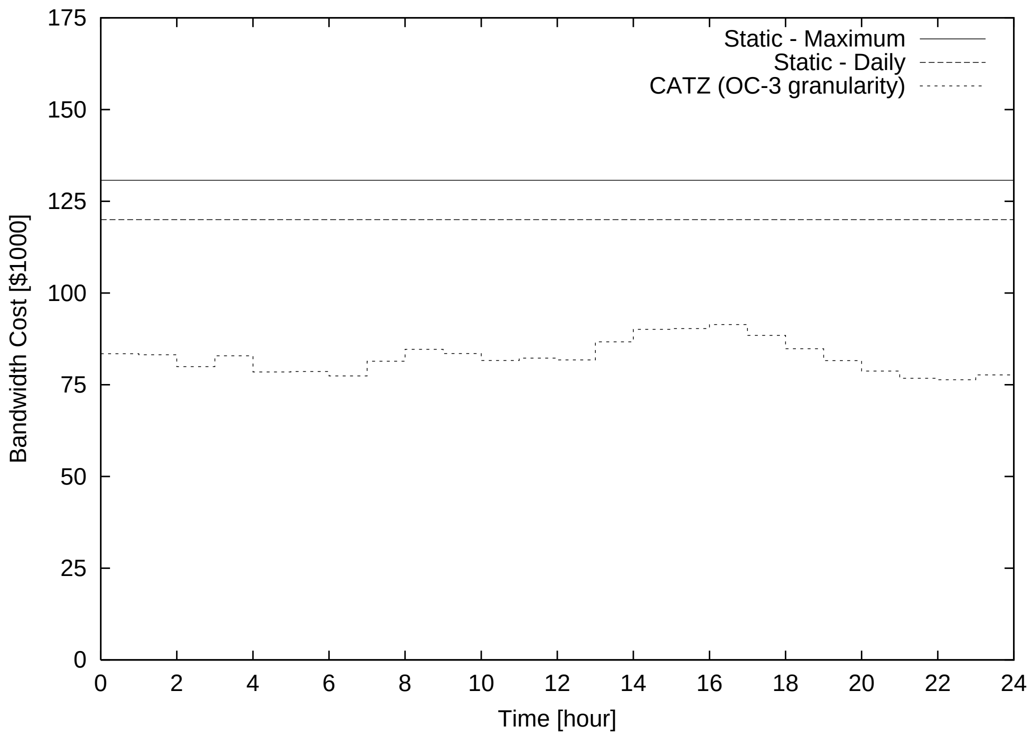

Figure 5.9: Bandwidth costs of different methods for the second traffic matrix given in Fig. C.1.

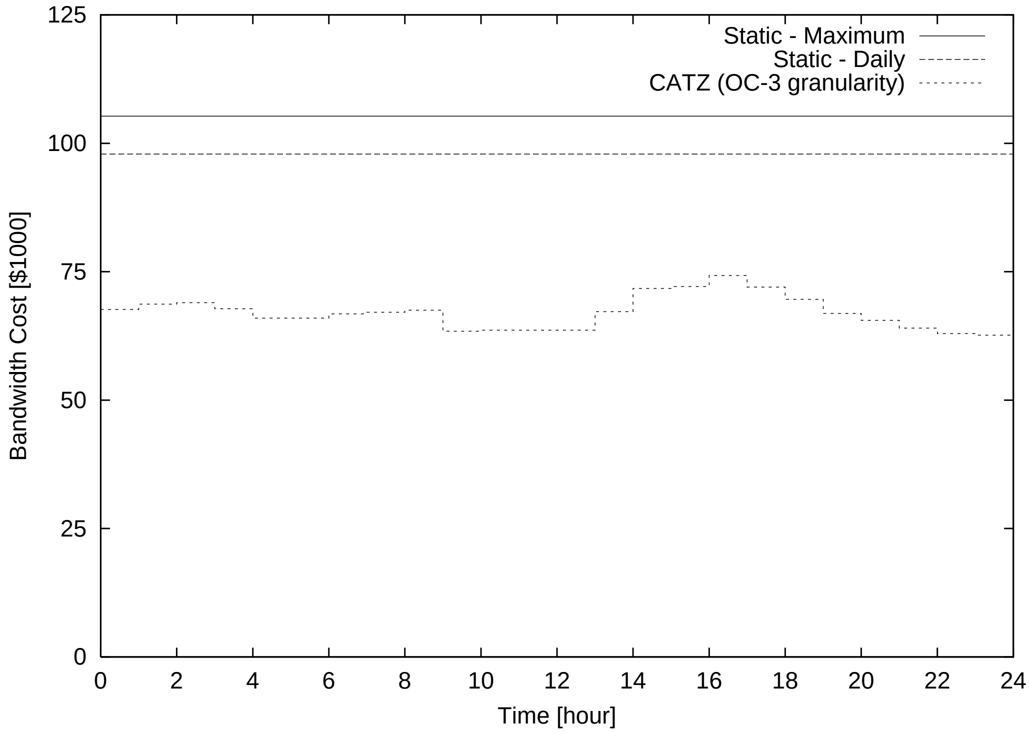

Figure 5.10: Bandwidth costs of different methods for the third traffic matrix given in Fig. C.2.

Table 5.1: Hourly bandwidth costs (in $1000) of the network for different traffic matrices

Traffic Matrix	Static-Max	Static-Daily	CATZ	% Improvement over Static-Max
1	113.534	105.073	72.700	36
2	130.729	120.002	82.580	37
3	105.297	97.909	67.072	36

5.7 Conclusion

This chapter introduced a bandwidth reconfiguration scheme for backbone networks, where dynamic channel provisioning is possible, and the traffic is dynamic. The daily fluctuations of the traffic and the time-zone effect are considered in this study. The new approach is compared with two static bandwidth allocation schemes using network bandwidth cost as the performance metric. The experiments show that an important cost improvement can be obtained by using a bandwidth reconfiguration scheme that takes into consideration the dynamism of the traffic.

6. CONCLUSION

Reconfiguration of wavelength-routed WDM networks under dynamic traffic has been an active area of research. With the improvement of optical technology, dynamic switching products are becoming available, and there is a need for intelligent methods to operate these increasingly capable network elements. On the other hand, it is becoming clear that the traffic demand will increase with the diversity of the data applications, and Internet will be the dominant data network which has a dynamic backbone traffic structure. Therefore, the challenge for the researchers is to design and operate the future's optical network such that, the ever increasing traffic demand can be accommodated, and the optical network resources can be utilised efficiently in a dynamic traffic environment.

As WDM networks offer reconfigurability at the transport layer by its nature, the problem maps to the design of intelligent algorithms for managing the network. What this study has aimed, was to provide methods for reconfiguring the optical layer connectivity, i.e. the virtual topology, under realistic assumptions on traffic.

This thesis developed a new perspective for the reconfiguration of the virtual topology, in a WDM backbone network. This new view was motivated by several judgements. First, a reconfiguration method should not rely on assumptions on the future traffic intensities to decide how the topology should be changed. Second, the algorithm should run on-line, since reconfiguration is an on-line process by its nature, i.e., the decision process should be quick enough to react promptly to the changes in traffic. Next, the process should not be broken into two independent fragments; deciding to trigger a reconfiguration and transition from the old topology to the new one. Instead, the reconfiguration should be a continuous process of observation-adjustments since the changes in backbone traffic are also observed to be smooth and continuous.

Given these bases, two new methods of topology reconfiguration were proposed in this thesis.

The first method was an adaptation scheme for virtual topology of a WDM network. Here, the problem was approached to as tracking the long-term traffic fluctuations by adapting the virtual topology in a continuous measurement-adaptation cycle. Two parameters were introduced, namely high watermark and low watermark, to control the loads of each lightpath in the virtual topology. These values worked as the upper and the lower bounds for the lightpath loads. The objective of the adaptation process was to maintain all of the lightpath loads between the two watermarks, by adding or deleting lightpaths whenever it is necessary. This method was different from the previous topology reconfiguration methods proposed, in terms of the properties of the new perspective explained above. No assumption was made on future traffic rates, the algorithm was simple, and the ongoing traffic was not interrupted by a transition phase where many lightpaths were changed in the same period. These advantages allowed the proposed method, to adapt any change in the traffic pattern, and to be able to run for large networks.

The performance of the new adaptation method was measured by using two methodologies. Using mixed-integer linear programming technique, the adaptation method was formulated as an optimisation problem. This formulation was solved for an examplary network by using a standard solver, and the results were compared to a previous reconfiguration method that optimises the number of lightpaths and the number of changes during the reconfiguration. The results of this comparison showed that the new method had a better adaptation property, since it required far less changes to adapt itself to the changing traffic. The adaptation method was simulated on real backbone network topologies to reveal several properties of the method. These simulation experiments showed that the adaptation method can follow the traffic changes, by making necessary lightpath additions and deletions.

The second method was a bandwidth reconfiguration scheme for world-wide WDM networks, where on-line bandwidth provisioning is commonplace. This method exploited the daily traffic fluctuations and time-zone based traffic shifts in the network, to operate it cost-effectively. The method provided considerable reduction in terms of total bandwidth cost. This result showed that the reconfigurability of the bandwidths of connections in a network might provide an efficient network operation. It also showed that a cost reduction could be achieved by considering the traffic dynamism in a world-wide network, even when the lightpath connectivity of the network remained unchanged.

There are several open issues in extending the virtual topology adaptation method. One such problem is determining the watermark values automatically, depending on the performance of the virtual topology. As the network performance metrics, such as total number of lightpaths, average packet hop distance, and virtual topology change frequency are affected by the watermarks, these metrics can be used as feedback information for an automated system. Such an automated scheme would allow changing the watermarks dynamically, in response to the dynamic traffic and the changing requirements of the service provider.

The adaptation algorithm is assumed to be centralised, however, a distributed algorithm can be developed. The signaling issues should be also considered, since the algorithm is assumed to be used by backbone networks which may include fibres spanning thousands of kilometers.

The bandwidth reconfiguration method, CATZ, can be extended to include different service types that carriers provide to the ISPs. The cost-related issues need to be modeled more accurately to have a better understanding of the real world networks. Currently this information is very difficult to obtain, however, developing bandwidth markets are expected to reveal the pricing schemes.

The results obtained in this thesis are encouraging for building the backbone reconfigurable WDM networks, since they show that simple and efficient algorithms that could properly manage the network resources can be designed. The new view of the problem introduced in this work, namely *adaptation* instead of *reconfiguration*, may propel new methods that are practical for using in large networks. Although the economics of the network are not a primary issue for researchers, it is vital for the realisation of the next generation WDM networks. The network operators should be given simple and cost-effective network management solutions, before they are convinced to invest for a completely new and expensive technology.

REFERENCES

[1] Abilene network traffic statistics, http://www.abilene.iu.edu/.

[2] Ryan, Hankin, and Kent (RHK), http://www.rhk.com/.

[3] **Mukherjee, B.**, 2000. WDM optical communication networks: Progress and challenges, *IEEE Journal on Selected Areas in Communications*, **18**, 1810-1824.

[4] **Ghani, N., Dixit, S. and Wang, T. S.**, 2000. On IP-over-WDM integration, *IEEE Communications Magazine*, **38**, 72-84.

[5] **Banerjee, A., Drake, J., Lang, J. P., Brad, T., Kompella, K. and Rekhter, Y.**, 2001. Generalized multiprotocol label switching: An overview of routing and management enhancements, *IEEE Communications Magazine*, **39**, 144-150.

[6] **Awduche, D. and Rekhter, Y.**, 2001. Multiprotocol lambda swithching: Combining MPLS traffic engineering control with optical crossconnects, *IEEE Communications Magazine*, **39**, 111-116.

[7] Internet Engineering Task Force, http://www.ietf.org/.

[8] **Wei, J. Y.**, 2002. Advances in the management and control of optical internet, *IEEE Journal on Selected Areas in Communications*, **20**, 768-785.

[9] **Yao, S., Mukherjee, B. and Dixit, S.**, 2000. Advances in photonic packet switching: An overview, *IEEE Communications Magazine*, **38**, 84-94.

[10] **Yao, S., Mukherjee, B., Yoo, S. J. B. and Dixit, S.**, 2001. All-optical packet switching: Challenges and opportunities, *IEEE Communications Magazine*, **39**, 142-148.

[11] **El-Bawab, T. S. and Shin, J.-D.**, 2002. Optical packet switching in core networks: Between vision and reality, *IEEE Communications Magazine*, **40**, 60-65.

[12] **Mukherjee, B.**, 1997. Optical Communication Networks, McGraw-Hill, New York.

[13] **Chlamtac, I., Ganz, A. and Karmi, G.**, 1992. Lightpath communications: An approach to high bandwidth optical WAN's, *IEEE Transactions on Communications*, **40**, 1171-1182.

[14] **Labourdette, J.-F. P. and Acampora, A. S.**, 1991. Logically rearrangeable multihop lightwave networks, *IEEE Transactions on Communications*, **39**, 1223-1230.

[15] **Labourdette, J.-F. P., Hart, G. W. and Acampora, A. S.**, 1994. Branch-exchange sequences for reconfiguration of lightwave networks, *IEEE Transactions on Communications*, **42**, 2822-2832.

[16] **Bala, K., Ellinas, G., Post, M., Shen, C.-C., Wei, J. and Antoniades, N.**, 1996. Toward hitless reconfiguration in WDM optical networks for ATM transport. *IEEE GLOBECOM*, London, England, November 18-22, pp. 316-320.

[17] **Banerjee, D. and Mukherjee, B.**, 2000. Wavelength-routed optical networks: Linear formulation, resource budgeting tradeoffs, and a reconfiguration study, *IEEE/ACM Transactions on Networking*, **8**, 598-607.

[18] **Ramamurthy, B. and Ramakrishnan, A.**, 2000. Virtual topology reconfiguration of wavelength routed optical WDM networks, *IEEE GLOBECOM*, San Francisco, USA, November 27-December 1, pp. 1269-1275.

[19] **Sreenath, N., Siva Ram-Murthy, C., Gurucharan, B. H. and Mohan, G.**, 2001. A two-stage approach for virtual topology reconfiguration of WDM optical networks, *Optical Networks Magazine*, **2**, 58-71.

[20] **Baldine, I. and Rouskas, G. N.**, 2001. Traffic adaptive WDM networks: A study of reconfiguration issues. *IEEE Journal of Lightwave Technology*, **19**, 433-455.

[21] **Ramaswami, R. and Sivarajan, K. N.**, 2002. Optical Networks A Practical Perspective, Morgan Kaufmann, San Francisco.

[22] **Stern, T. E. and Bala, K.**, 1999. Multiwavelength Optical Networks: A Layered Approach, Addison Wesley, Boston.

[23] **Kramer, G. and Pesavento, G.**, 2002. Ethernet passive optical networks (EPON): Building a next-generation optical access network. *IEEE Communications Magazine*, **40**, 66-73.

[24] **Gencata, A., Singhal, N. and Mukherjee, B.**, 2003. Overview of Optical Communication Networks: Current and Future Trends, in *Handbook of Optical Communication Networks*, pp. 2-21, Eds. Ilyas, M. and Mouftah, H. T., CRC Press, Boca Raton.

[25] **Bishop, D. J., Giles, C. R. and Austin, G. P.**, 2002. The lucent lambdarouter: MEMS technology of the future here today, *IEEE Communications Magazine*, **40**, 75-79.

[26] **Chu, P. B., Lee, S. and Park, S.**, 2002. MEMS: The path to large optical crossconnects, *IEEE Communications Magazine*, **40**, 80-87.

[27] OIF–Optical Internetworking Forum, http://www.oiforum.com.

[28] **Payne, D.**, 2002. The evolution and impact of photonics components. *ECOC–28th European Conference on Optical Communications*, Copenhagen, Denmark, September 8-12.

[29] **Green, P.**, 2001. Progress in optical networking, *IEEE Communicaions Magazine*, **39**, 54-61.

[30] **Mukherjee, B.**, 1992. WDM-based local lightwave networks–Part I: Single-hop systems, *IEEE Network*, **6**, 12-27.

[31] **Mukherjee, B.**, 1992. WDM-based local lightwave networks–Part II: Multihop systems, *IEEE Network*, **6**, 20-32.

[32] **Kazovsky, L. G. and Poggiolini, P. T.**, 1993. STARNET: a multi-gigabit-per-second optical LAN utilizing a passive WDM star, *IEEE/OSA Journal of Lightwave Technology*, **11**, 1009-1027.

[33] **Jia, F., Mukherjee, B. and Iness, J.**, 1995. Scheduling variable-length messages in a single-hop multichannel local lightwave network, *IEEE/ACM Transactions on Networking*, **3**, 477-488.

[34] **Duanyang, G. and Acampora, A. S.**, 1996. Scalable multihop wdm passive ring with optimal wavelength assignment and adaptive wavelength routing, *IEEE/OSA Journal of Lightwave Technology*, **14**, 1264-1277.

[35] **Marsan, M. A., Bianco, A., Leonardi, E., Meo, M. and Neri, F.**, 1996. MAC protocols and fairness control in WDM multirings with tunable transmitters and fixed receivers, *IEEE/OSA Journal of Lightwave Technology*, **14**, 1230-1244.

[36] **Hamad, A. M. and Kamal, A. E.**, 2002. A survey of multicasting protocols for broadcast-and-select single-hop networks. *IEEE Network*, **16**, 36-48.

[37] **Tridandapani, S. B., Mukherjee, B. and Hallingstad, G.**, 1997. Channel sharing in multi-hop WDM lightwave networks: do we need more channels?, *IEEE/ACM Transactions on Networking*, **5**, 719-727.

[38] **Kramer, G., Mukherjee, B. and Pesavento, G**, 2002. IPACT: A dynamic protocol for an Ethernet PON (EPON), *IEEE Communications Magazine* **40**, 74-80.

[39] **Green, P. E.**, 2002. Paving the last mile with glass, *IEEE Spectrum* **39**, 13-14.

[40] **Davis, C. C., Smolyaninov, I. I. and Milner, S. D.**, 2003. Flexible optical wireless links and networks, *IEEE Communications Magazine*, **41**, 51-57.

[41] **Chiu, A. L. and Modiano, E. H.**, 2000. Traffic grooming algorithms for reducing electronic multiplexing costs in WDM ring networks, *IEEE Journal of Lightwave Technology*, **18**, 2-12.

[42] **Gerstel, O., Ramaswami, R. and Sasaki, G. H.**, 2000. Cost-effective traffic grooming in WDM rings, *IEEE/ACM Transactions on Networking*, **8**, 618-630.

[43] **Zhang, X. and Qiao, C.**, 2000. An effective and comprehensive approach for traffic grooming and wavelength assignment in SONET/WDM rings, *IEEE/ACM Transactions on Networking*, **8**608-617.

[44] **Berry, R. and Modiano, E. H.**, 2000. Reducing electronic multiplexing costs in SONET/WDM rings with dynamically changing traffic, *IEEE Journal on Selected Areas in Communications*, **18**, 1961-1971.

[45] **Wang, J., Cho, W., Vemuri, V. R. and Mukherjee, B.**, 2001. Improved approaches for cost-effective traffic grooming in WDM ring networks: ILP formulations and single-hop and multihop connections, *IEEE/OSA Journal of Lightwave Technology*, **19**, 1645-1653.

[46] **Dutta, R. and Rouskas, G. N.**, 2002. On optimal traffic grooming in WDM rings, *IEEE Journal on Selected Areas in Communications*, **20**,110-121.

[47] **Grover, W. D. and Doucette, J.**, 2001. A novel heuristic for topology planning and evolution of optical mesh networks, *IEEE GLOBECOM*, San Antonio, USA, November 25-29, pp. 2169-2173.

[48] **Ali, M.**, 2002. Optimization of splitting node placement in wavelength-routed optical networks, *IEEE Journal on Selected Areas in Communications*, **20**, 1571-1579.

[49] **Ramamurthy, B., Datta, D., Feng, H., Heritage, J. P. and Mukherjee, B.**, 1999. Impact of transmission impairments on the teletraffic performance of wavelength routed optical networks, *IEEE/OSA Journal of Lightwave Technology*, **17**,1713-1723.

[50] **Zang, H., Ou, C. and Mukherjee, B.**, 2003. Path-protection routing and wavelength assignment (RWA) in WDM mesh networks under duct-layer constraints, *IEEE/ACM Transactions on Networking*, **11**, 248-258.

[51] **Zang, H. and Mukherjee, B.**, 2001. Connection management for survivable wavelength routed WDM mesh networks, *Optical Networks Magazine*, **2**, 17-28.

[52] **Ramamurthy, S. and Mukherjee, B.**, 1999. Survivable WDM mesh networks, Part I–protection, *IEEE INFOCOM–21st Conference on Computer Communications*, New York, USA, March 23-25, pp. 744-751.

[53] **Ramamurthy, S and Mukherjee, B.**, 1999. Survivable WDM mesh networks, Part II–restoration, *19th IEEE International Conference on Communications*, Vancouver, Canada, June 6-10, pp. 2023-2030.

[54] **Gerstel, O. and Ramaswami, R.**, 2000. Optical layer survivability-an implementation perspective, *IEEE Journal on Selected Areas in Communications*, **18**, 1885-1899.

[55] **Gerstel, O. and Ramaswami, R.**, 2000. Optical layer survivability-a service perspective, *IEEE Communications Magazine*, **38**,104-113.

[56] **Sahasrabuddhe, L., Ramamurthy, S. and Mukherjee, B.**, 2002. Fault management in IP-over-WDM networks: WDM protection versus IP restoration, *IEEE Journal on Selected Areas in Communications*, **20**, 21-33.

[57] **Modiano, E. and Narula-Tam, E.**, 2002. Survivable lightpath routing: a new approach to the design of WDM-based networks, *IEEE Journal on Selected Areas in Communications*, **20**, 800-809.

[58] **Wang, J., Sahasrabuddhe, L. and Mukherjee, B.**, 2002. Path vs. subpath vs. link restoration for fault management in IP-over-WDM networks: performance comparisons using GMPLS, *IEEE Communications Magazine*, **40**, 80-87.

[59] **Sasaki, G. H. and Ching-Fong, S.**, 2003. The interface between IP and WDM and its effect on the cost of survivability, *IEEE Communications Magazine*, **41**, 74-79.

[60] **Cox, L. A. and Sanchez, J. R.**, 2001. Cost savings from optimized packing and grooming of optical circuits: Mesh versus ring comparisons, *Optical Networks Magazine*, **2**, 72-90.

[61] **Modiano, E. and Lin, P. J.**, 2001. Traffic grooming in WDM networks, *IEEE Communications Magazine*, **39**, 124-129.

[62] **Lardies, A., Gupta, R. and Patterson, R. A.**, 2001. Traffic grooming in a multi-layer network, *Optical Networks Magazine*, **2**, 91-99.

[63] **Zhu, K. and Mukherjee, B.**, 2002. Traffic grooming in an optical WDM mesh network, *IEEE Journal on Selected Areas in Communications*, **20**, 122-133.

[64] **Dutta, R. and Rouskas, G. N.**, 2002. Traffic grooming in WDM networks: past and future, *IEEE Network*, **16**, 46-56.

[65] **Cinkler, T.**, 2003. Traffic and λ grooming, *IEEE Network*, **17**, 16-21.

[66] **Zhu, K., Zhu, H and Mukherjee, B.**, 2003. Traffic engineering in multigranularity heterogeneous optical WDM mesh networks through dynamic traffic grooming, *IEEE Network*, **17**, 8-15.

[67] **Sahasrabuddhe, L. and Mukherjee, B.**, 1999. Light-trees: Optical multicasting for improved performance in wavelength-routed networks, *IEEE Communications Magazine*, **37**, 67-73.

[68] **Sun, Y., Gu, J. and Tsang, D. H. K.**, 2001. Multicast routing in all-optical wavelength routed networks, *Optical Networks Magazine*, **2**, 101-109.

[69] **Chen, B. and Wang, J.**, 2002. Efficient routing and wavelength assignment for multicast in WDM networks, *IEEE Journal on Selected Areas in Communications*, **20**, 97-109.

[70] **O'Mahony, M. J., Simeonidou, D., Hunter, D. K. and Tzanakaki, A.**, 2001. The application of optical packet switching in future communication networks, *IEEE Communications Magazine*, **39**, 128-135.

[71] **Yuan, X. C., Li, V. O. K., Li, C. Y. and Wai, P. K. A.**, 2003. A novel self-routing address scheme for all-optical packet-switched networks with arbitrary topologies, *IEEE/OSA Journal of Lightwave Technology*, **21**, 329-339.

[72] **Qiao, C. and Yoo, M.**, 1999. Optical burst switching (OBS) - a new paradigm for an optical internet, *Journal on High Speed Networks*, **8**, 69-84.

[73] **Baldine, I., Rouskas, G. N., Perros, H. G. and Stevenson, D.**, 2002. Jumpstart: a just-in-time signaling architecture for WDM burst-switched networks, *IEEE Communications Magazine*, **40**, 82-89.

[74] **Banerjee, D. and Mukherjee, B.**, 1996. A practical approach for routing and wavelength assignment in large wavelength-routed optical networks, *IEEE Journal on Selected Areas in Communications*, **14**, 903-908.

[75] **Wauters, N. and Demeester, P.**, 1996. Design of the optical path layer in multiwavelength cross-connected networks, *IEEE Journal on Selected Areas in Communications*, **14**, 881-892.

[76] **Cinkler, T., Marx, D., Larsen, C. P. and Fogaras, D.**, 2000. Heuristic algorithms for joint configuration of the optical and electrical layer in multi-hop wavelength routing networks, *IEEE INFOCOM-22nd Conference on Computer Communications*, Tel-Aviv, Israel, March 26-30, pp. 1000-1009.

[77] **Karasan, E. and Ayanoglu, E.**, 1998. Effects of wavelength routing and selection algorithms on wavelength conversion gain in WDM optical networks, *IEEE/ACM Transactions on Networking*, **6**, 186-196.

[78] **Xu, S., Li, L., and Wang, S.**, 2000. Dynamic routing and assignment of wavelength algorithms in multifiber wavelength division multiplexing networks, *IEEE Journal on Selected Areas in Communications*, **18**, 2130-2137.

[79] **Kodialam, M. and Lakshman, T. V.**, 2001. Integrated dynamic IP and wavelength routing in IP over WDM network, *IEEE INFOCOM–23rd Conference on Computer Communications*, Anchorage, USA, April 22-26, pp. 358-366.

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

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

[82] **Mukherjee, B., Banerjee, D., Ramamurthy, S. and Mukherjee, A.**, 1996. Some principles for designing a wide-area WDM optical network, *IEEE/ACM Transactions on Networking*, **4**, 684-696.

[83] **Krishnaswamy, R. M. and Sivarajan, K. N.**, 2001. Design of logical topologies: A linear formulation for wavelength-routed optical networks with no wavelength changers, *IEEE/ACM Transactions on Networking*, **9**, 186-198.

[84] **Dutta, R. and Rouskas, G. N.**, 2000. A survey of virtual topology design algorithms for wavelength routed optical networks, *Optical Networks Magazine*, **1**, 73-89.

[85] **Dutta, R. and Rouskas, G. N.**, 2000. Design of Logical Topologies for Wavelength Routed Networks, in *Optical WDM Networks: Principles and Practice*, pp. 79-102, Eds. Sivalingam, K. M. and Subramaniam, S., Kluwer Academic Publishers, Boston.

[86] **Jeong, H. Y., Li, S. S. and Seo, S.-W.**, 2000. An adaptive distributed wavelength routing algorithm in WDM networks, *IEEE GLOBECOM*, San Francisco, USA, November 27-December 1, pp. 1259-1263.

[87] **Nakamura, H., Kato, T., Yokoyama, H. and Yamamoto, S.**, 2000. Proposal of statistical lambda multiplexing network (SLAMNet), *IEEE GLOBECOM*, San Francisco, USA, November 27-December 1, pp. 1319-1323.

[88] **Narula-Tam, A. and Modiano, E.**, 2000. Dynamic load balancing for WDM based packet networks, *IEEE INFOCOM-22nd Conference on Computer Communications*, Tel-Aviv, Israel, March 26-30, pp. 1010-1019.

[89] CAIDA–Cooperative Association for Internet Data Analysis, http://www.caida.org.

[90] **Rouskas, G. N. and Ammar, M. H.**, 1995. Dynamic reconfiguration in multihop WDM networks, *Journal of High Speed Networks*, **4**, 221-238.

[91] ILOG Inc., http://www.ilog.com.

[92] **Leonardi, E., Mellia, M. and Marsan, M. A.**, 2000. Algorithms for the logical topology design in WDM all-optical networks, *Optical Networks Magazine*, **1**, 35-46.

[93] RateXchange, http://www.ratexchange.com.

[94] Fiberloops. http://www.fiberloops.com.

[95] **Kenyon, C.**, 2001. Stochastic models for telecom commodity prices, *Computer Networks*, **36**, 533-555.

[96] **Mukherjee, B.**, 2002. R&D priorities and challenges for optical networking: Migrating theory to practice, preprint.

[97] **Thompson, K., Miller, G. J. and Wilder, R.**, 1997. Wide-area internet traffic patterns and characteristics, *IEEE Network*, **11**, 10-23.

[98] **Subramaniam, S. and Barry, R. A.**, 1997. Wavelength-assignment in fixed routing WDM networks, *IEEE ICC–17th International Conference on Communications*, Montreal, Canada, June 8-12, pp 406-410.

A. ROUTING AND WAVELENGTH ASSIGNMENT ALGORITHMS

This section describes some basic heuristic methods used in routing and wavelength assignment for connection establishment in optical networks. A thorough study of these methods and their performance comparison can be found in [80].

A.1 Basic Routing Algorithms

Fixed Routing

This is the most simple method, which chooses the same fixed route for a given source-destination pair. When the shortest path is required for a connection, the shortest paths for each node pair is calculated off-line, and these paths are used to route the connection when a connection request arrives. Since this method does not consider the current state of the network resources and chooses the same path whether the network is congested or not, it may lead to high blocking performance, or may result in a large number of wavelengths being used.

Fixed Alternate Routing

In this approach multiple routes are considered between node pairs. Each node in the network is required to maintain a routing table which contains an ordered list of a number of fixed routes to each destination node, such as, the shortest path, the second shortest path, the third shortest path. When a connection request arrives, the routes in the routing table are tried in sequence until the connection is established succesfully. This method improves the blocking performance significantly compared to fixed routing.

Adaptive Routing

In adaptive routing, the route from a source node to a destination node is chosed dynamically, depending on the state of network resources. In shortest-path adaptive routing, when a connection request arrives, the shortest path between the source node and destination node is determined. The routing algorithm considers only available network resources (transceivers and wavelengths). A connection is blocked only when there is no route from the source node to the destination node. Adaptive routing method requires more support from the control and management entity to continuously update the routing tables, but it results in lower connection blocking than fixed and fixed-alternate routing methods.

A.2 Basic Wavelength Assignment Algorithms

Many heuristics have been proposed in literature, and they can be combined with different routing schemes to provide solutions for dynamic routing and wavelength assignment problem. Here, we define three basic wavelength assignment methods. Other methods are, min-product, least loaded, MAX-SUM, relative capacity loss, wavelength reservation, and protecting threshold.

First-fit Wavelength Assignment

In this scheme all wavelengths are numbered. When the algorithm searches for an available wavelength it considers first the lowest numbered wavelength. The wavelengths are searched in ascending order of their number until an available one is found. The idea of this scheme is to pack all of the in-use wavelengths towards the lower end of the wavelength space, so that continuous longer paths towards the higher end of the wavelength space will have a higher probability of being available. This scheme performs well in terms of blocking probability, and it is preferred in practise because of its simplicity.

Least-used Wavelength Assignment

This method selects the least used wavelength in the network, attempting to balance the usage of all wavelengths. This schemes causes the long wavelength paths break quickly, therefore only connection requests which traverse a small number of links can be services. The performance of this scheme may be worse than random wavelength assignment.

Most-used Wavelength Assignment

This scheme is the opposite of least-used, since it selects the most used wavelength in the network when trying to set up a new connection. It outperforms least-used moethod [98]. It also packs connections into fewer wavelengths, as first-fit.

B. SIMULATION RESULTS FOR VIRTUAL TOPOLOGY ADAPTATION

This section includes additional simulation results, for the heuristic adaptation algorithm introduced in Chapter 4. First, the operation of the algorithm is shown for different watermark values, for the 19-node physical topology. Next, the simulation results for another telco network with 24 nodes are illustrated.

B.1 Typical Operation

A typical operation of the adaptation algorithm has been illustrated in Figure 4.3 of Chapter 3, with watermark values $W_H = 70$ and $W_L = 10$. Following figures plot the operation of the algorithm with different watermark values. Here, the 19-node physical topology shown in Figure 4.2 is used, and other network parameters are the same as in Section 4.4.

Figure B.1 shows the highest and lowest lightpath loads in the network and the times of lightpath addition and deletion, for $W_H = 65$ and $W_L = 10$. This figure can be compared to Figure 4.3 and Figure B.2 ($W_H = 75$ and $W_L = 10$) to observe the effect of the high watermark value on the operation of the adaptation. As the high watermark value increases, the number of changes to the virtual topology decreases as expected.

The cases where W_L is changed and W_H is held constant are shown in Figures B.3 and B.4. The value of W_L is 8 in Figure B.3, and 12 in Figure B.4. These two figures also should be observed by comparing them to Figure 4.3. Note that the specific values used for the watermarks, are related to the properties of the physical topology and the randomly generated traffic matrix.

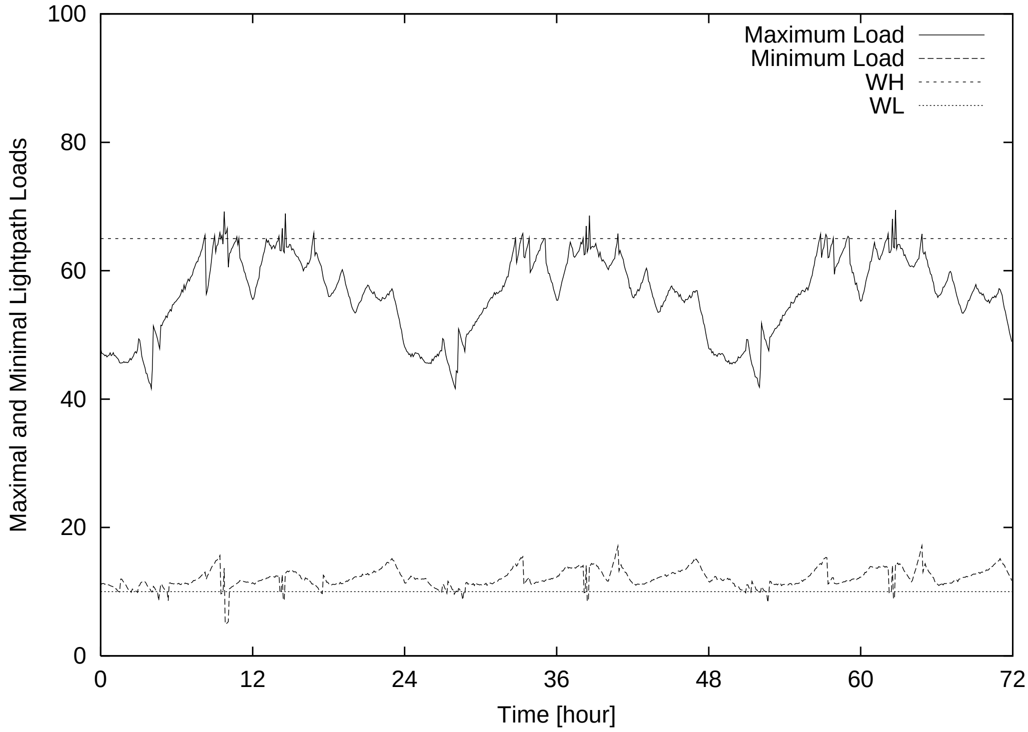

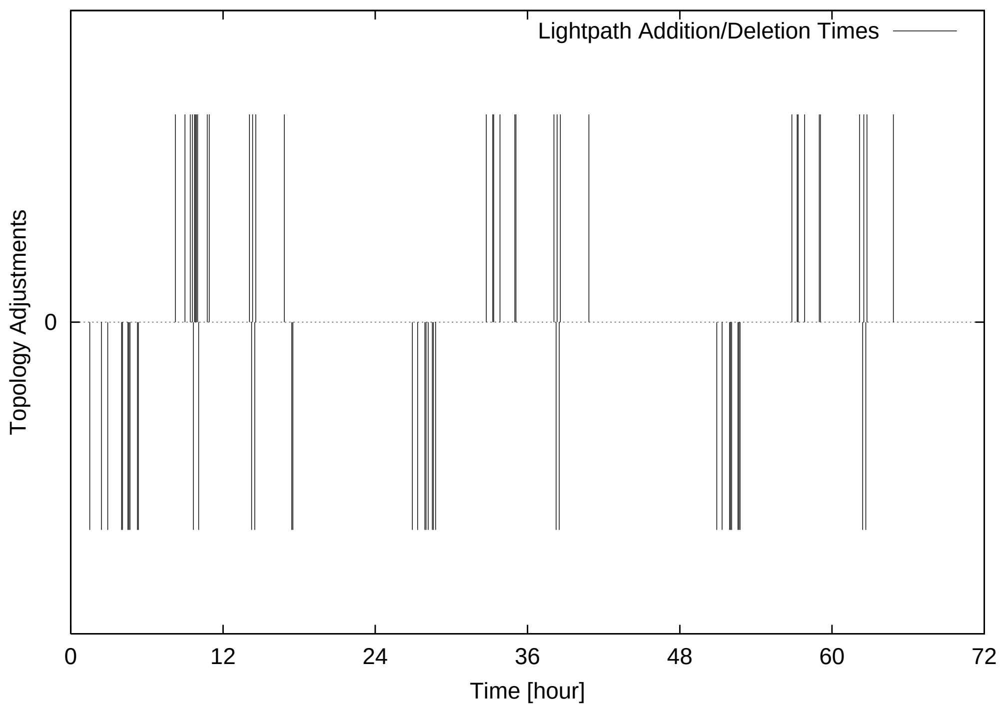

Figure B.1: a. Maximal and minimal lightpath loads in the network during a 3-day run ($W_H = 65, W_L = 10$). b. Impulse graphic indicating times of lightpath addition or deletion.

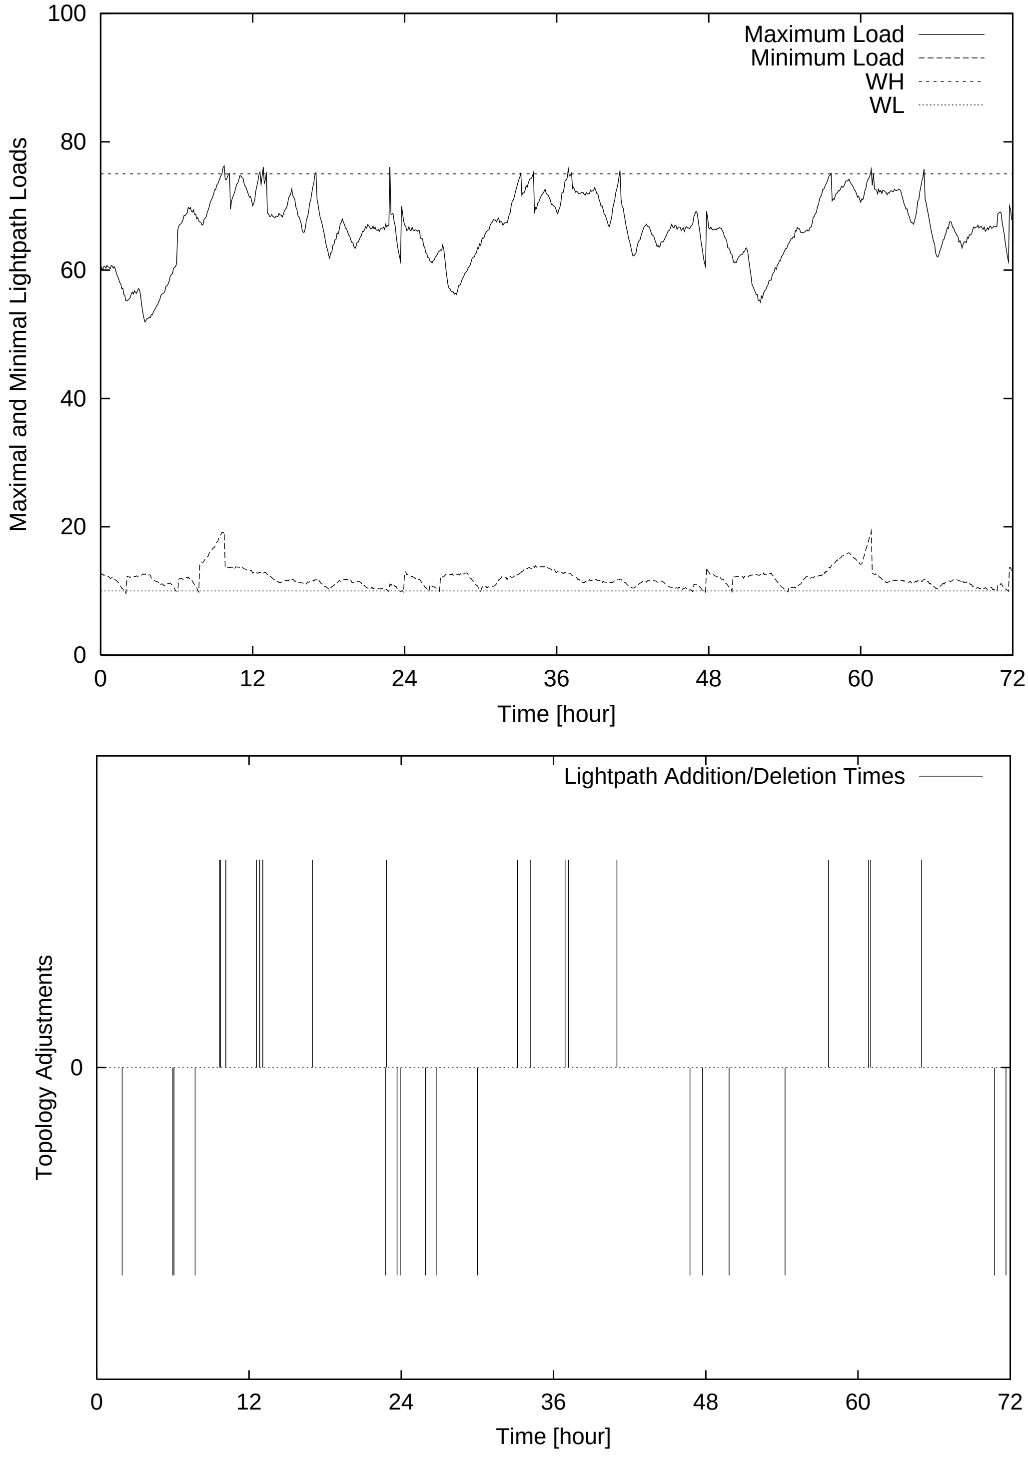

Figure B.2: a. Maximal and minimal lightpath loads in the network during a 3-day run ($W_H = 75, W_L = 10$). b. Impulse graphic indicating times of lightpath addition or deletion.

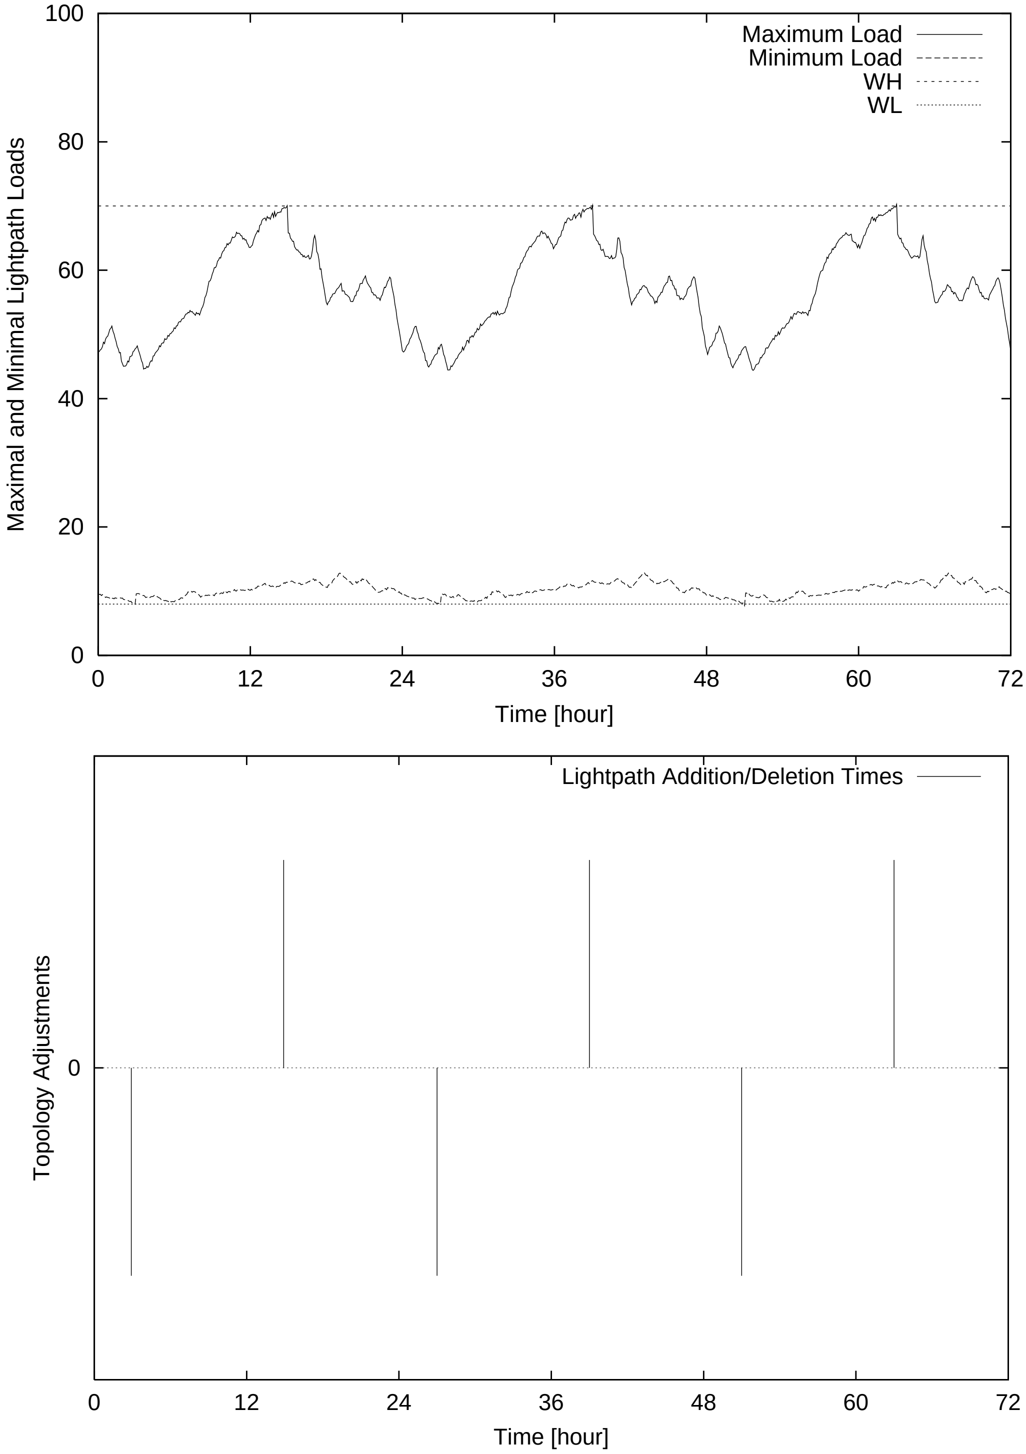

Figure B.3: a. Maximal and minimal lightpath loads in the network during a 3-day run ($W_H = 70, W_L = 8$). b. Impulse graphic indicating times of lightpath addition or deletion.

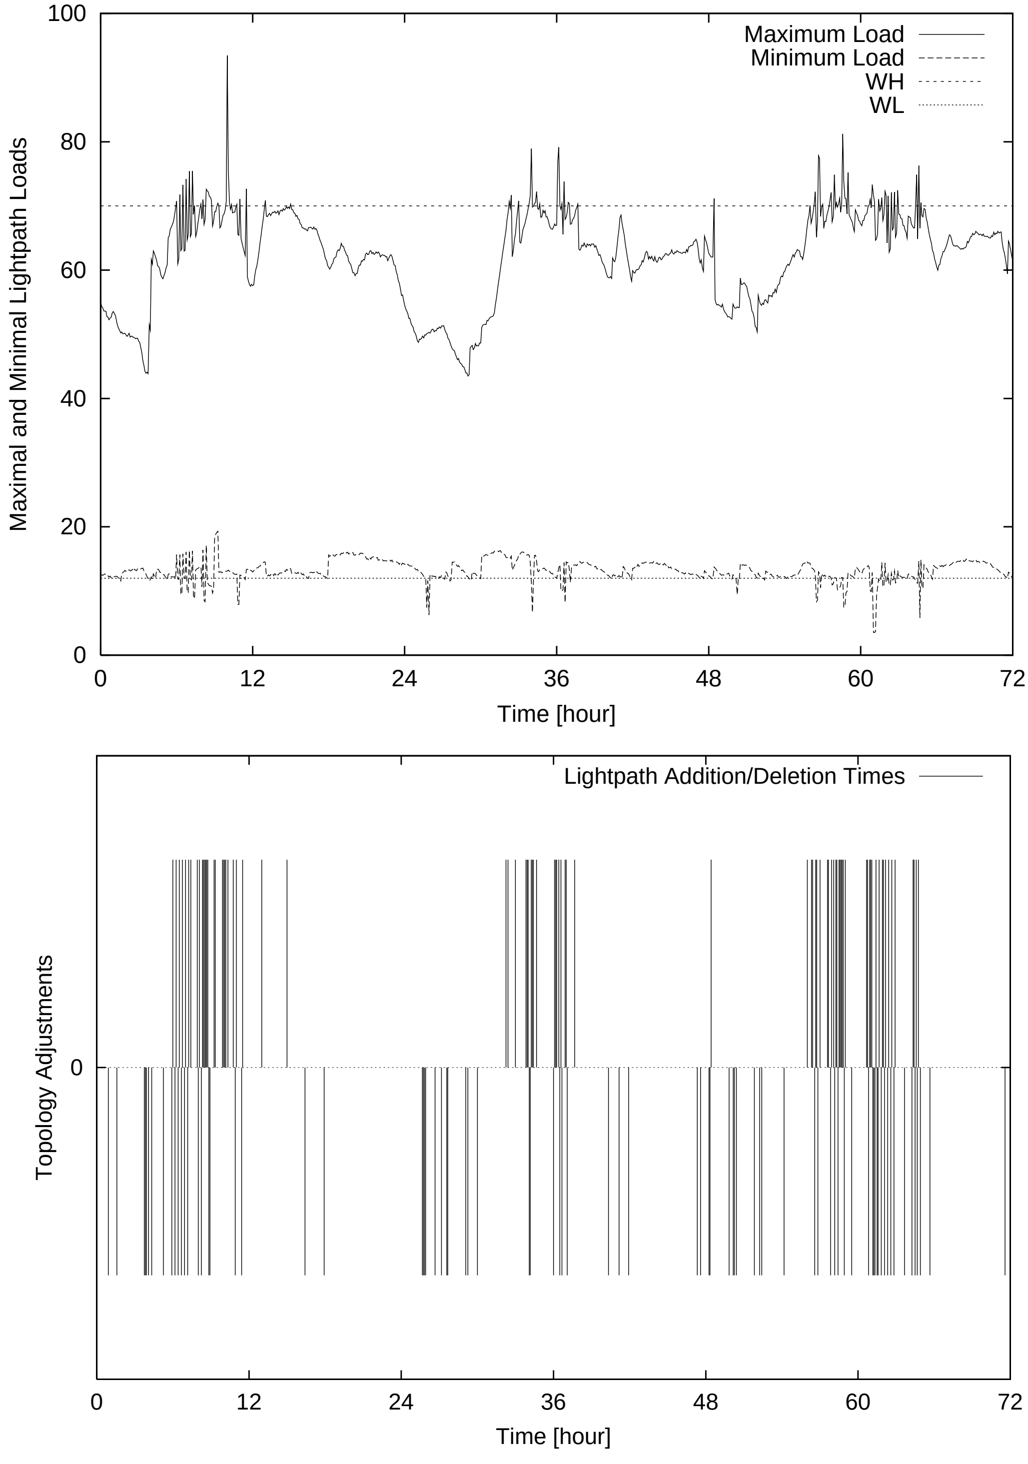

Figure B.4: a. Maximal and minimal lightpath loads in the network during a 3-day run ($W_H = 70, W_L = 12$). b. Impulse graphic indicating times of lightpath addition or deletion.

B.2 Results of Adaptation for a 24-node Topology

The adaptation method is applied to a second physical topology shown in Figure B.5. This section includes the results obtained for this topology. The traffic is randomly generated using the technique introduced in Section 4.4.1. 60% of node pairs have non-zero traffic rates at any time of the day. The simulations were run for 1-day period.

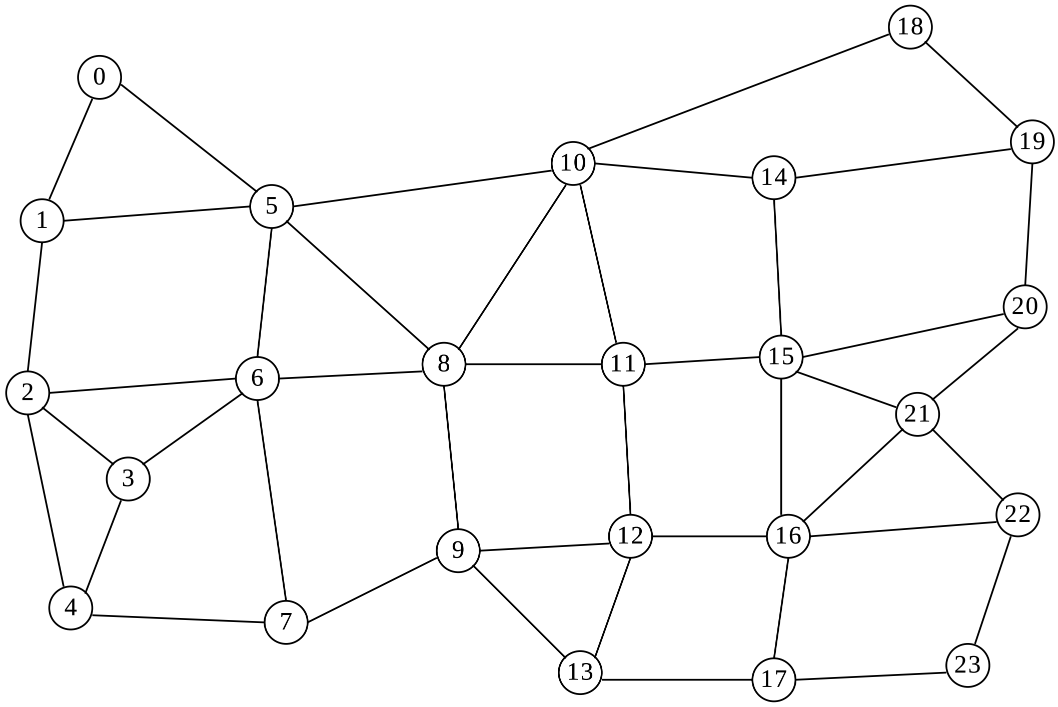

Figure B.5: A 24-node telco network used as the physical topology in the simulations.

Figures B.6 and B.7 show typical runs for 1 day, with $W_H = 70$, $W_L = 10$, and $W_H = 70$, $W_L = 12$. Figures B.8 and B.9 plot the similar graphics for $W_H = 65$. These figures show that the algorithm performs as expected, in a different physical topology and traffic configuration. Also, the values to be chosen for the watermarks can be different for different network configurations. These values are related to the traffic between nodes, the average number of changes that the operator would allow, and the total number of lightpaths that the operator would decide to have in its virtual topology.

Table B.1 shows the number of changes (total number of lightpath additions and deletions) during a day, for different watermark values. The number of changes increases when the gap between the two watermarks decreases.

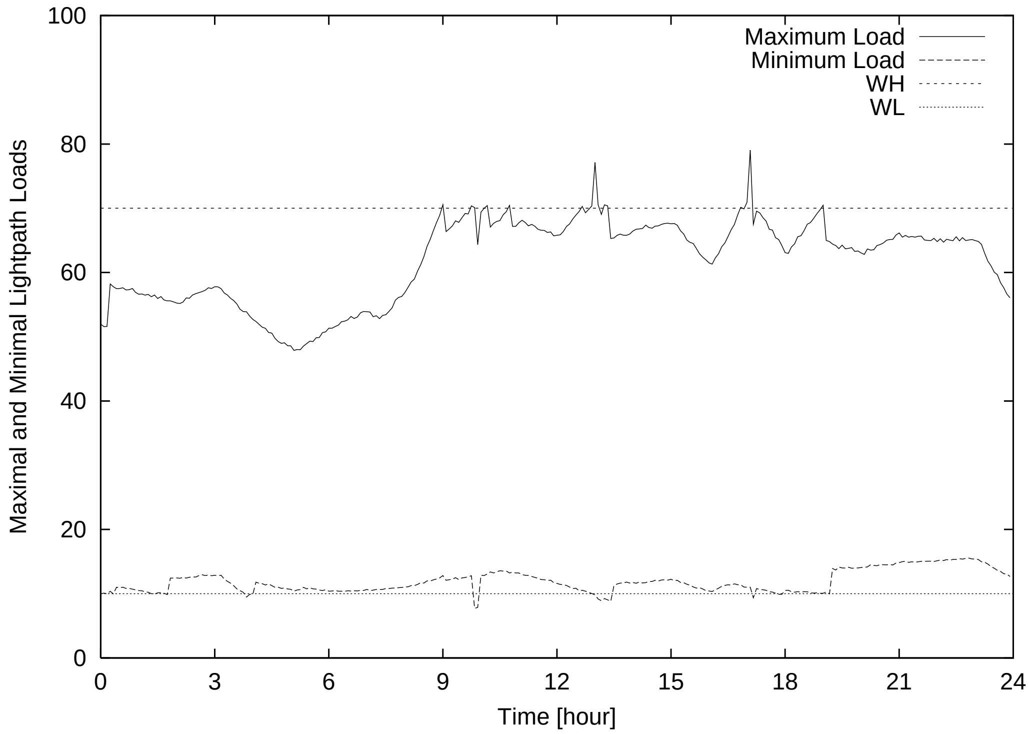

Figure B.6: Maximal and minimal lightpath loads in the network during a 1-day run ($W_H = 70, W_L = 10$)

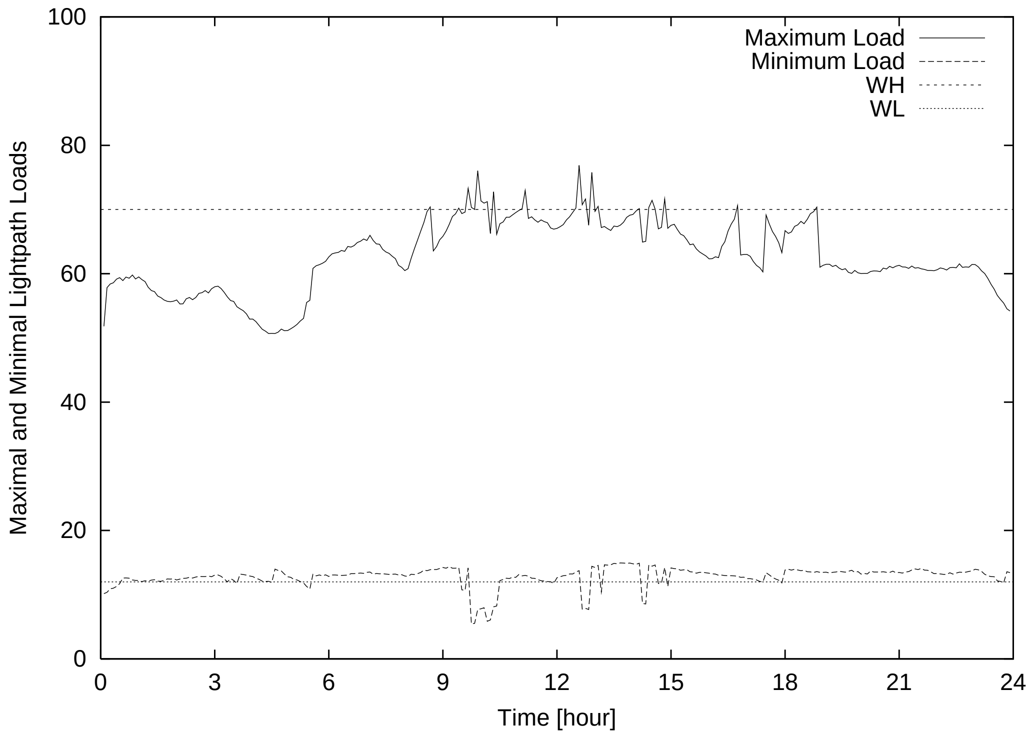

Figure B.7: Maximal and minimal lightpath loads in the network during a 1-day run ($W_H = 70, W_L = 12$)

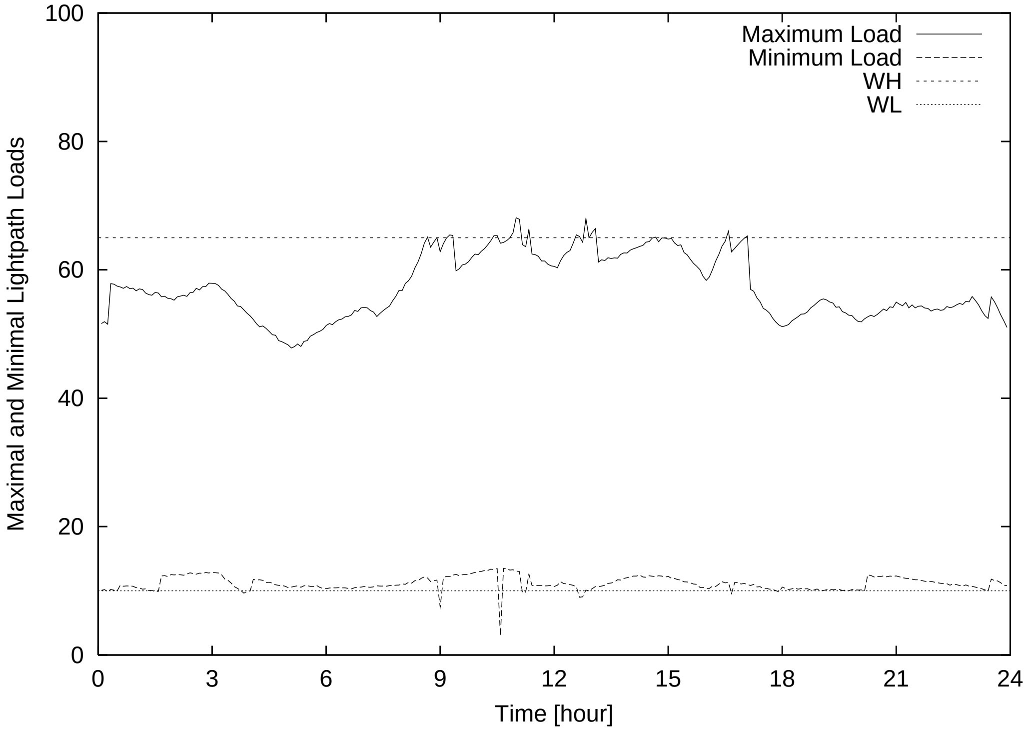

Figure B.8: Maximal and minimal lightpath loads in the network during a 1-day run ($W_H = 65, W_L = 10$)

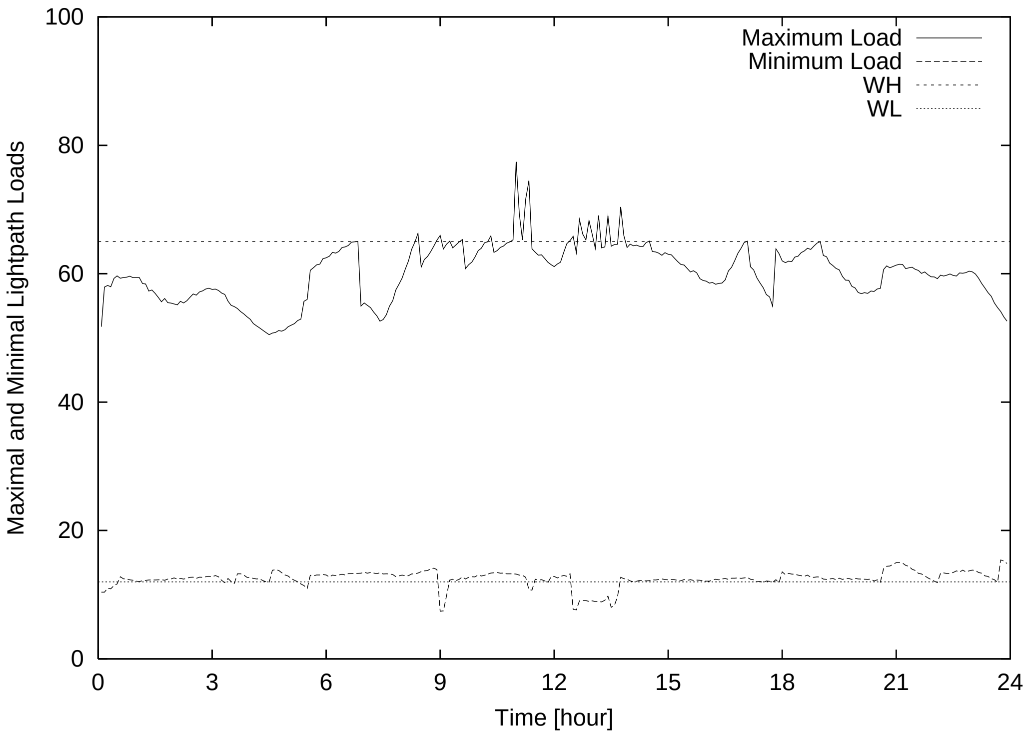

Figure B.9: Maximal and minimal lightpath loads in the network during a 1-day run ($W_H = 65, W_L = 12$)

Table B.1: Number of lightpath additions and deletions during a day for different watermarks.

W_H	W_L	
	10	12
65	34	56
70	29	49

C. TRAFFIC MATRICES USED IN THE SIMULATION OF CATZ

Two traffic matrix sets are used in the simulation experiments in Chapter 5. Figures C.1 and C.2 show 8 snapshot matrices for each set.

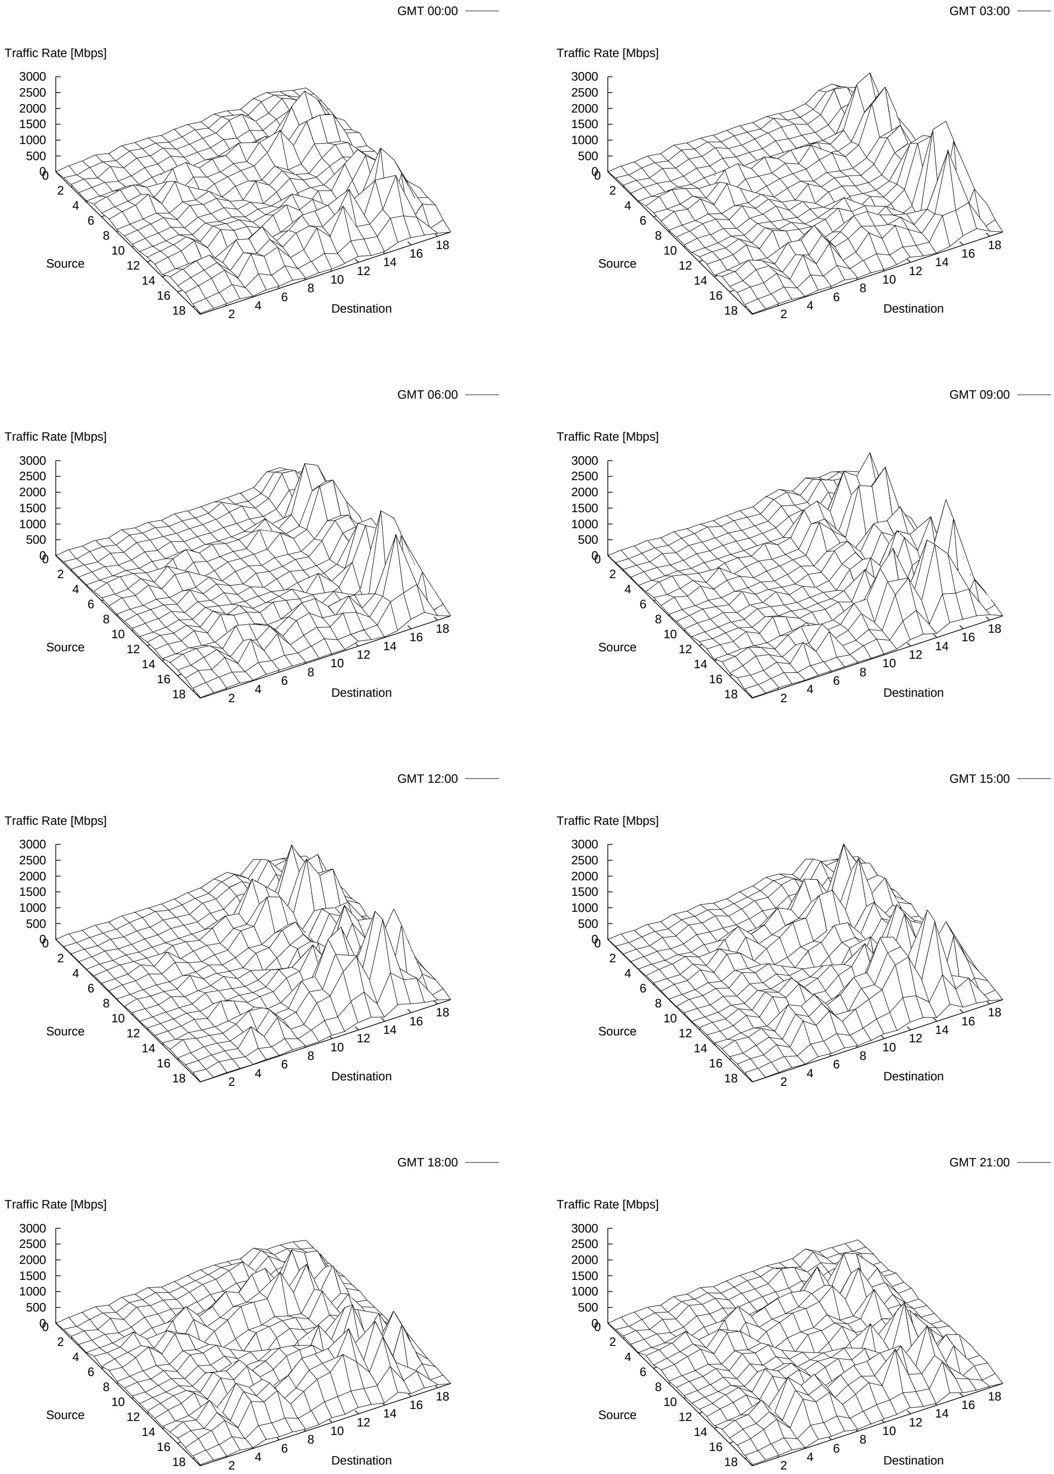

Figure C.1: Traffic Matrix Set 2 at different hours of the day.

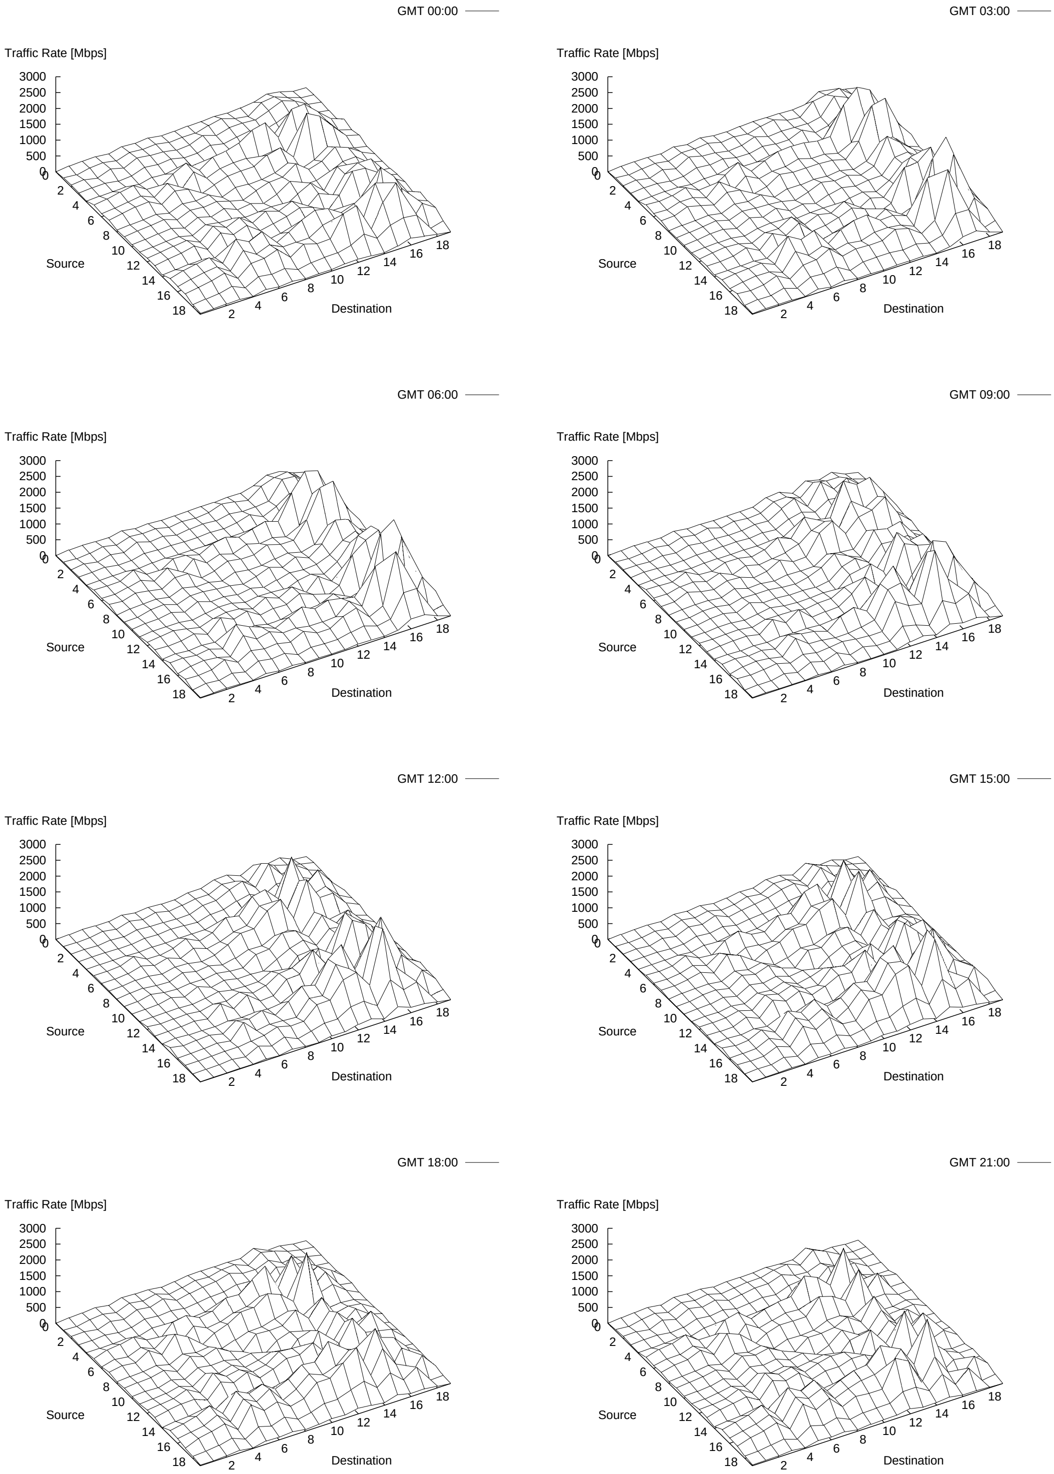

Figure C.2: Traffic Matrix Set 3 at different hours of the day.

D. THE STRUCTURE OF THE SIMULATION

This section describes the structure of the simulation designed for the adaptation algorithm in Figure 4.1. The whole simulation is written in C++ using event-driven methodology. The simulation time is given as an input parameter, and the experiment runs until the given simulation time is achieved.

Every event during the runtime is inserted into the *Event Queue* as soon as it is created. The events are processed according to their time stamps, hence the *Current Time* is a discrete variable that progress from one event to the next. The main loop of the simulation is as follows:

```
while (Event Queue is not empty){
     Get an event from the front of the Event Queue;
     Call event handler;
}
```

The simulation starts with creating one data packet and inserting it into the queue. The end of the run is also an event that is put into the queue at the setup stage. Different types of events may be created during the run, most important ones are: PacketArrival, UpdateRates, Lightpath, IPreroute, RTchange, VTchange, IPreport, and EndSimulation. The general structure of the Event Queue is shown in Figure D.1.

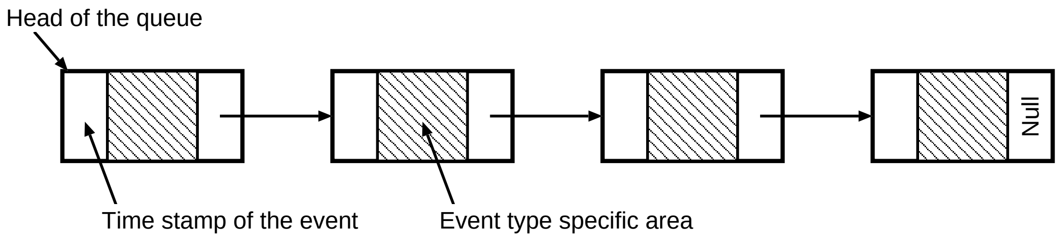

Figure D.1: Event Queue model.

Two functions manage the operation of the queue. *Schedule()* inserts a newly created event into the queue, in increasing order of its time stamp. *Get()* removes an event from the front of the queue, where exists the event with the smallest timestamp.

Traffic Generator

The average traffic rates for each hour during a day are taken from an input stream, and installed into a data structure. Every minute an event UpdateRates is activated to update the mean traffic rates between each source-destination pairs at that time, by interpolating the given values.

Every time a packet arrival occurs, i.e. a packet is dequeued from the Event Queue, a new packet is created and inserted into the queue. Next packets to be generated for each communicating node pair, are stored in the table *NextPackets*. When a new packet creation is required from the traffic generator, the NextPackets table is searched for the packet with the smallest timestamp. This packet is removed from the table, and inserted into the Event Queue, while its place in the table is filled up with the next packet for the same source-destination pair (Figure D.2).

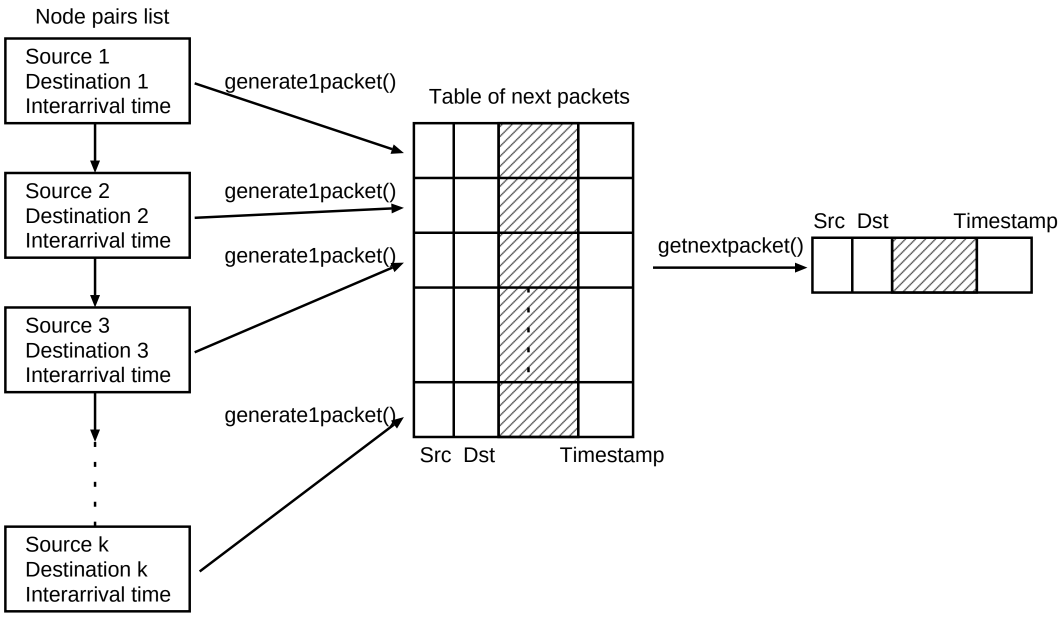

Figure D.2: Traffic generator model.

Reconfiguration

An IPreport event is scheduled at the setup stage for the end of the first observation period. The event handler calls the function *TE()* where the lightpath loads are calculated, and the reconfiguration algorithm in Figure 4.1 is run. The next IPreport event is scheduled also by this function.

If a change needs to be done in the virtual topology, either a lightpath addition or a ligthpath deletion, a Lightpath event is scheduled. The Ligthpath handler calls the *request()* function where the virtual topology change request is processed. For a lightpath add request, routing and wavelength assignment function is called to establish the new lightpath. Than, for both types of requests, VTchange and IPreroute events are scheduled. These two events take care of the maintenance of the virtual topology data structure and IP routing tables.

Traffic Observation

PacketArrival handler passes the packet to *routepacket()* function, where the load created by this packet is added to the route it is following. Each packet is assumed to arrive its destination node.

BIOGRAPHY

Ayşegül Gençata graduated from Galatasaray Lisesi in 1989. She received her B.Sc. in 1993 from the Control and Computer Engineering Department of Istanbul Technical University, and her M.Sc. in 1995 from the Institute of Science and Technology of the same university. She enrolled in the Ph.D. program of the same institute in 1995, and is a research and teaching assistant at the Computer Engineering Department of İ.T.Ü. since 1993. She was a visiting research scholar at the University of California at Davis, Computer Networks Laboratory between December 2000 and June 2002. She is a student member of IEEE.

During her work on Ph.D. thesis, she published the following related papers:

Book chapter:

- A. Gençata, N. Singhal, and B. Mukherjee, Overview of Optical Communication Networks: Current and Future Trends, in *Handbook of Optical Communication*, CRC Press, to appear in April 2003.

Journal paper:

- A. Gençata, B. Mukherjee, "Virtual Topology Adaptation for WDM Mesh Networks Under Dynamic Traffic," *IEEE/ACM Transactions on Networking*, to appear in April 2003 issue.

Conference papers:

- A. Gencata, L. Sahasrabuddhe, and B. Mukherjee, "Virtual-Topology Adaptation with Minimal Lightpath Change for Dynamic Traffic in WDM Mesh Networks," *Proc. of Optical Fiber Communications Conf. - OFC*, pp. ThGG119, Anaheim, CA, Mar. 2002.
- A. Gençata, B. Mukherjee, "Virtual Topology Adaptation for WDM Mesh Networks Under Dynamic Traffic," *Proc. of IEEE INFOCOM*, New York, Jun. 2002.
- A. Gençata, B. Mukherjee, "Time-zone Based Traffic Engineering and Capacity Optimization in a World-wide Backbone Mesh Network," *High-Speed Networking Workshop - HSN 2002*, Invited talk, New York, Jun. 2002.
- A. Gençata, L. Sahasrabuddhe, B. Mukherjee, "Dynamic Network Planning and Design using Optical WDM Infrastructure," *Proc. of National Fiber Optic Engineers Conference - NFOEC*, Dallas, TX, Sept. 2002.
- A. Gençata, B. Mukherjee, "CATZ (Capacity Allocation with Time Zones): A Methodology for Cost-efficient Bandwidth Allocation and Reconfiguration in a World-wide WDM Network," *Proc. European Conference on Optical Communications - ECOC 2002*, Copenhagen, Sept. 2002.