ABSTRACT

Mahalati, Ruhiyyih Reconfiguration of Sub-wavelength Groomed Wavelength Routed Optical Networks. (Under the direction of Dr. Rudra Dutta).

Telecommunication networks recently have seen a large increase in traffic demands, especially data traffic as compared to voice traffic. With the advances in fiber optics and wavelength division multiplexing (WDM) optical networking is the key to satisfy the data-driven bandwidth demand. These technologies enable simultaneous transmission of signals on separate high-speed channels at different wavelengths. While current technologies can provide such huge bandwidth, in order to utilize efficiently the capacity of each lightpath, a number of independent lower-rate traffic streams must be multiplexed into a single lightpath. This technique is referred to as traffic grooming. Another most attractive feature of WDM and wavelength routing networks is the possibility of adaptively creating virtual topologies, or a set of lightpaths, based on network need, giving rise to the concept of reconfiguration. Till date, however, there has been little or no work on the joint consideration of the two areas of traffic grooming and reconfiguration, even though it is clear that reconfiguration is at least equally important in the realistic networks of tomorrow which will definitely need to carry sub-wavelength traffic. This is probably due to two reasons: the common wisdom has been that the two aspects can be handled separately, and also it is hard to define reasonable network design goals if the two aspects are considered jointly. In this thesis, we examine this issue of reconfiguration in groomed networks. The overall focus of this work has been the balancing of the reconfiguration cost and a good grooming solution. We define suitable goals for an integrated approach, and provide formulation of the integrated approach as an integer linear program. To allow a joint consideration of grooming and reconfiguration costs, we mathematically formulate a representation of the reconfiguration cost in terms of the OXC and DXC reconfiguration that can be related to grooming costs. We also develop a heuristic strategy for addressing the problem, which attempts to achieve minimal disturbance reconfiguration by performing local reconfiguration and delaying the need for global reconfiguration. In order to proactively avoid global reconfiguration, we introduce the concept of over provisioning at the traffic demand level. Our reconfiguration heuristic minimizes the need for solving the integer linear program, which is computationally intractable. We also present numerical results validating our claims.

Reconfiguration of Sub-wavelength Groomed Wavelength Routed Optical Networks

by

Ruhiyyih Mahalati

A thesis submitted to the Graduate Faculty of
North Carolina State University
in partial satisfaction of the
requirements for the Degree of
Master of Science

Department of Computer Science

Raleigh

2003

Approved By:

Dr. Khaled Harfoush	Dr. Douglas S. Reeves

Dr. Rudra Dutta
Chair of Advisory Committee

To my parents . . .

Biography

Ruhiyyih Mahalati was born in Indore, India, in 1979. She received her Bachelor's degree in Computer Science from Devi Ahilya University, Indore, India in 2001. Since 2001, she has been a Masters student in the Department of Computer Science at North Carolina State University.

Acknowledgements

I would like to express my gratitude to my advisor, Dr. Rudra Dutta for his guidance, patience and the countless hours he spent with me to help bring this thesis to the current form. This thesis would not have been possible without his constant support and encouragement.

I am grateful to my committee members Dr. Douglas S. Reeves and Dr. Khaled Harfoush for devoting time and providing valuable inputs.

I acknowledge the immense help I received from my colleagues Bensong Chen and Koundinya Srinivasarao. Without their help and support I would not have been able to achieve the desired level in my work.

Finally, I would also like to thank my family and all my friends for always being there for me.

Contents

List of Figures

List of Abbreviations

DXC	-	Digital Cross-connect
ILP	-	Integer Linear Program
LSP	-	Label Switched Path
LSR	-	Label Switching Router
LTE	-	Line Terminating Equipment
OAM	-	Operations, Administrations and Management
OEO	-	Opto-electro-optical
OTDM	-	Optical Time Division Multiplexing
OXC	-	Optical Cross-connect
RWA	-	Routing and Wavelength Assignment
SDH	-	Synchronous Digital Hierarchy
SONET	-	Synchronous Optical Network
WADM	-	Wavelength Add/Drop Multiplexer
WDM	-	Wavelength Division Multiplexing
WRON	-	Wavelength Routed Optical Network

Chapter 1

Introduction

Telecommunication networks recently have seen a large increase in traffic demands, especially data traffic as compared to voice traffic. The evolution of optical networking technologies has resulted in the realization of high bandwidth and high availability networks not just in the research laboratories but also in commercial environments; see [8] for a list of optical networking and switching vendors. The main breakthrough was in implementing optical *wavelength division multiplexing* (WDM). WDM is the process of transmitting data simultaneously at multiple carrier wavelengths over an optical fiber cable keeping the wavelengths sufficiently far apart so that they do not interfere with each other. Thus, a single strand of fiber can be thought of as a collection of high capacity *virtual fibers*. A single strand of optical fiber has the potential bandwidth of 50 THz. This bandwidth can be split across multiple WDM channels with each channel being capable of carrying several Gbps of data (e.g. Lucent's offering of 160 wavelengths of 10Gbps capacity [42]). Over the last few years we have witnessed a wide deployment of point-to-point WDM transmission technology in the Internet infrastructure. Today, WDM systems are widely deployed in long-haul (wide-area) networks. They have a major presence in the metro-area networks as well. The corresponding massive increase in bandwidth due to WDM has heightened the need for faster switching at the core of the network. This coincides with a growing desire to support different levels of Quality of Service (QoS), which makes it necessary to not only efficiently utilize the greatly increased available bandwidth, but to manage it with high predictability. Label Switching Routers (LSRs) running Multi-Protocol Label Switching (MPLS) [14] are being deployed to address the dual issues of faster switching and QoS support. On one hand, LSRs simplify the forwarding function, thereby making it possible to operate at higher data rates. On the other hand, MPLS enables the Internet architecture, built upon the connectionless Internet Protocol, to behave in a connection-oriented fashion that is more conducive to QoS and traffic engineering [3]. With the further advancement of wavelength selective routing, it

has become possible to use wavelengths as generalized labels, and create *lightpaths* or λ-circuits as the primary transport mechanism in such networks; lightpaths are clear optical channels spanning several physical fiber links, switched entirely optically at intermediate nodes.

While current technologies can provide such huge bandwidth, a network would typically need to support sub-wavelength traffic connections as well. In order to utilize efficiently the capacity of each lightpath, a number of independent lower-rate traffic streams must be multiplexed into a single lightpath. These observations give rise to the concept of *traffic grooming*, which refers to the techniques used to combine lower speed traffic flows onto available wavelengths in order to meet network design goals such as cost minimization.

With the deployment of commercial WDM systems, it has become apparent that the cost of network components, especially line terminating equipment (LTE), is a dominant cost in building optical networks, and is a more meaningful metric to optimize than, say, the number of wavelengths. In the existing literature, LTE cost minimization is addressed by formulating the problem so that its objective reflects either the number of LTEs directly, or in a more indirect way by minimizing the amount of electronic (as opposed to optical) routing performed. Most existing studies concentrate on some aggregate representation of these i.e., the total number of LTEs (or the total amount of electronic routing performed) over all of the network nodes is minimized.

One of the most attractive features of WDM and wavelength routing networks is the possibility of adaptively creating *virtual topologies*, or a set of lightpaths, based on network need, as is possible when the optical switches can be reconfigured on-line. To higher client layers, the lightpaths appear to be physical links due to the very small and highly invariant propagation delays along them when compared to connections mediated by switching/routing by electronic devices at intermediate nodes: a consequence of the circuit-switched nature of lightpaths. Changing the physical topology of the network is usually a very long and costly operation, and cannot usually be undertaken in response to short- or middle-term changes in network traffic demand. The virtual topology, on the other hand, can be changed over comparatively far shorter time scales, by simply reconfiguring the optical switches. Not surprisingly, the topic of efficient and effective reconfiguration strategies for optical networks has received significant attention in literature.

Till date, however, there has been little or no work on the joint consideration of the two areas of traffic grooming and reconfiguration, even though it is clear that reconfiguration is at least equally important in the realistic networks of tomorrow which will definitely need to carry sub-wavelength traffic. This is probably due to two reasons: the common wisdom has been that the two aspects can be handled separately, and also it is hard to define reasonable network design goals if the two aspects are considered jointly.

In this thesis, we examine the issue of reconfiguration in groomed networks. We develop a framework for considering the question, and propose methods to develop joint design goals. Our numerical results indicate that a disjoint treatment of the two issues would incur undesirable effects, thus an OAM strategy based on common wisdom would not be the best practical one.

The rest of this thesis is arranged as follows. In the next chapter we briefly describe the context of this thesis and other related work, describing the relevant network architecture mentioned above in more detail. In Chapter 3 we present a formal definition of the new problem we identify. The formulation of the integrated approach as an ILP is presented in the Chapter 4. In Chapter 5 we present the over-provisioning algorithm and the reconfiguration heuristic which is practically more useful since the computation done by the ILP is prohibitively expensive. Numerical experimentation validating the performance of the proposed algorithm are presented in Chapter 6. Chapter 7 discusses possible directions for future research and concludes this thesis.

Chapter 2

Context

2.1 Architecture of Optical Networks

Computer networks have evolved from transmission links made of copper (Ethernet) to the copper being replaced by point-to-point fiber (SONET, Gigabit Ethernet) to today's optical networks with wavelength routing. The architecture of current day optical networks has been designed to exploit the unique features of optics and provides functionality beyond point-to-point optical transmission.

One of the most significant features of current optical networks is *wavelength division multiplexing* (WDM). A single fiber has enormous usable bandwidth (nearly 50 THz). Exploiting all of this bandwidth using a single very high capacity channel is impossible and impractical. It is essential to divide this bandwidth into channels of lower capacity and WDM essentially achieves this. WDM involves transmitting several signals simultaneously on different wavelengths. Each wavelength runs at the peak rate that the electronic devices at the endpoints can handle. While conceptually optical time division multiplexing (OTDM) is also an option it is yet to mature as a technology [43], especially when compared to WDM systems which have already had presence in the marketplace for several years.

The earliest architectures that were proposed were *opaque* in that the optical path of transmission between two endpoints was interrupted at all intermediate nodes, converted to the electronic form, processed, converted back to the optical form and then transmitted on the outgoing link to the next intermediate node or the destination node. This conversion process is referred to as optical-electrical-optical (OEO) conversion. In such opaque networks the buffering and routing (more precisely, switching) is handled by the electronic equipment and consequently, the OEO conversion

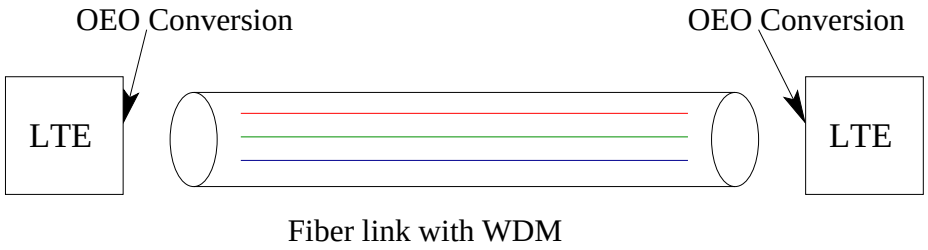

Figure 2.1: Wavelength Division Multiplexing in a Fiber Link

points become the bottleneck in such networks. The other undesirable side effects of OEO conversion are that it introduces delay, variability of delay, loss of throughput, increased probability of errors and buffer overflow. Figure 2.1 shows WDM in an optical fiber link. Each wavelength transmitted requires line terminating equipment (LTE) to convert the signal from the optical form in which it is transmitted, to the electronic form suitable for use by today's computers. A WADM passes traffic on certain wavelengths through without interruption or optoelectronic conversions. Figure 2.2 shows a WADM that supports W wavelengths. In the figure, the wavelengths λ_1 and λ_2 are dropped (terminated) and added whereas the other wavelengths pass through optically. A node can be made *transparent* with respect to some of the wavelengths by making such nodes optically bypass any signal transmitted on these wavelengths using *wavelength add/drop multiplexers* (WADM).

Shifting away from *opaque* networks, the other extreme is a *transparent* optical network architecture. There exists no electronic routing in such networks; all the signals are switched in the optical form at the intermediate nodes. Some of the advantages of transparency is that the throughput of the fiber can be exploited to the fullest without any restrictions on data rates and data formats. Transparency also makes architecture of the optical nodes and network control simple, but they fail to perform 3R (regeneration, reshaping, retiming), which is a benefit of OEO conversion. Such transparent optical network architectures may however not be feasible in realistic situations due to constraints on the number of wavelengths as well as hardware constraints at the nodes. From a network design point of view, both opaque and transparent networks can be just as costly, as illustrated in an example in [19]. Consequently, the *hybrid* or "almost all-optical" network which combines the features of both the worlds has been accepted to be the candidate for next generation wide area networks.

A *hybrid* network may comprise of a few *opaque* nodes and *transparent* nodes. However opaque nodes are known to have several problems associated with them. The large amount of LTE required at such nodes makes them too expensive or impractical to deploy due to high interconnection costs, power consumption and space requirements. As illustrated in [19], having none of the nodes perform any electronic routing is also not a viable solution. Thus, the networks need to be equipped with *translucent* nodes, in that such nodes are transparent with respect to some of the wavelengths

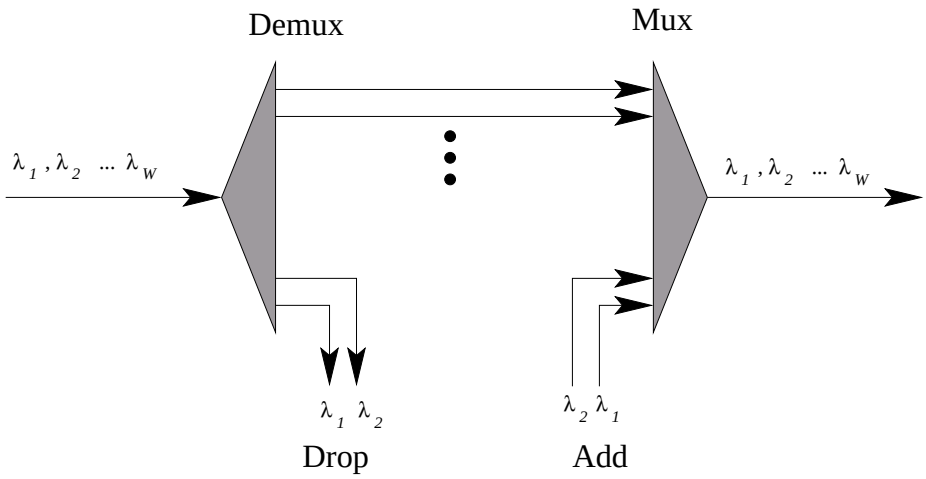

Figure 2.2: Wavelength Add/Drop Multiplexer

and opaque with respect to others.

2.1.1 Optical Cross Connects

A more powerful component in optical networks is the *optical cross-connect* (OXC) which can optically switch a signal from a particular input port (and hence, a fiber link) to a particular output port, independent of the other wavelengths. OXC allows optical network to be reconfigurable by adapting to the changed traffic conditions without having the need to change the physical topology. The design of OXC must adhere to the characteristics of the traffic flowing, in order to implement a cost-effective and flexible architecture as illustrated in [62]. There are different OXC architectures which are distinguished primarily on the basis of the switching fabric [62].

The OXCs provide the switching and routing functions for supporting logical data connections. An all-optical OXC is one whose switching fabric is optical and does not require OEO conversion of the incoming signals for switching purposes. All-optical OXCs operate completely in the optical domain and have higher capacity and lower cost per port than other OXC architectures. Figure 2.3 shows the architecture of a typical all-optical OXC. The signals coming in over different fibers are first demultiplexed. All the signals at a particular wavelength are sent to an optical switch dedicated for that wavelength. The signals from the outputs of the optical switches are multiplexed back together onto the outgoing fibers. If W is the number of wavelengths each fiber carries, N_s the number of incoming/outgoing fibers together with any inputs/outputs from the DXC, W mux-demux pairs and W $N_s \times N_s$ switches are required to implement the OXC shown in the figure. The $N_s \times N_s$ optical switches can be configured such that any $N_s - N_f$ wavelengths, N_f the number of incoming/outgoing fibers, can be added or dropped locally from and to any fiber. For the OXC

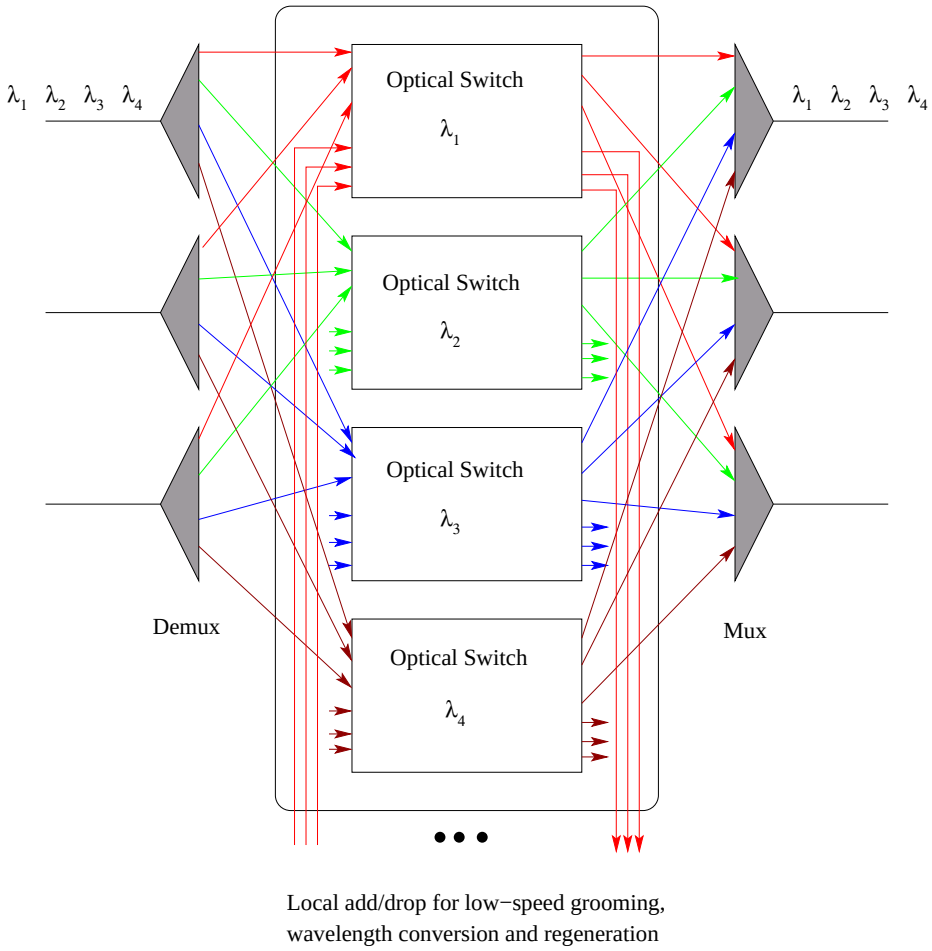

Figure 2.3: An all-optical OXC with local add/drop

shown in Figure 2.3 N_f is 3 and N_s is 6. With this configuration most of the switching is in the optical domain, minimizing the cost and maximizing the capacity of the network, while allowing us to route the signals up to the electronic layer for the purposes of low-speed grooming, wavelength conversion and signal regeneration. It should be noted that what has been presented is a functional description of the working of an OXC. An optical cross-connect may be implemented using various technologies such as liquid crystal and MEMS (micro electromechanical systems), but the essential function remains the same.

2.1.2 Digital Cross Connects

The structure of an optical node is shown in Figure 2.4. *Digital Cross Connect* (DXC) is deployed along with each OXC, to offer a limited degree of OEO capability, for regeneration,

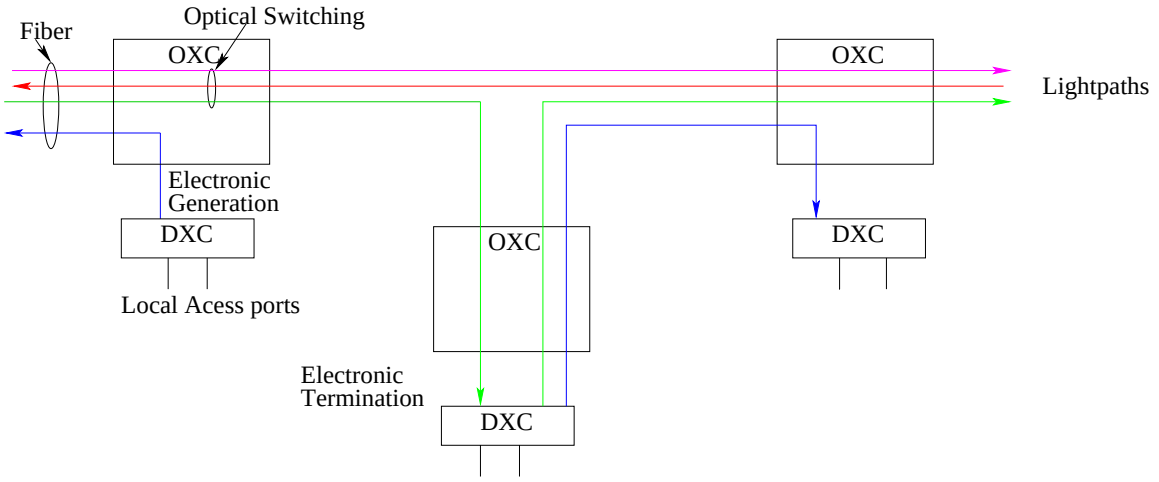

Figure 2.4: An optical node

protection and grooming. Some of the wavelengths on some of the input ports of the OXC may be carrying traffic destined for the access node directly connected to this OXC. These wavelengths are terminated at the DXC, where the signal is extracted from the optical domain and converted into electronic form, which can then be delivered to the higher layers. Certain signals on these wavelengths may be destined for other nodes, these are converted back to optical form at the DXC and injected to an output port of the OXC. Thus, the DXCs perform grooming by extracting and switching a large number of lower-speed streams at the nodes. The DXC considered here is a monolithic DXC with limited capacity. A monolithic DXC has a single switching fabric that is capable of switching TDM slots on any of the wavelengths from the input ports to any other wavelength on any output port. A full capacity DXC can terminate all the signals from the OXC from all the incoming fiber links and generate traffic for all the outgoing fiber links of the OXC, however, a limited capacity DXC operates at capacity less than this.

2.1.3 Optical Mesh Topology

The rapid advancement and evolution of optical technologies makes it possible to move beyond point-to-point WDM transmission systems to an all-optical backbone network that can take full advantage of the available bandwidth by relaxing the need for per-hop packet forwarding. Typically, the backbone networks have comprised of several high-speed routers interconnected as SONET ring networks and such rings themselves interconnected to form rings of rings. However, with the advent of the optical cross-connect slowly the trend is shifting towards a generalized arbitrary mesh topology as shown in Figure 2.5. Such a network consists of a number of optical cross-connects

(OXCs) arranged in some arbitrary topology. Each OXC can switch the optical signal coming in on a wavelength of an input fiber link to the same wavelength in an output fiber link. The OXC may also be equipped with converters that permit it to switch the optical signal on an incoming wavelength of an input fiber to some other wavelength on an output fiber link without OEO. The main mechanism of transport in such a network is the lightpath, which is a communication channel established between two LSRs (or other edge devices) over the network of OXCs, and which may span a number of fiber links. A lightpath is a generalization of the MPLS concept of a label-switched path (LSP) with the wavelength color corresponding to the label and the cross-connect matrix corresponding to the label forwarding table [14]. Lightpaths are setup between different LSRs by appropriately configuring the LSRs and the OXCs in the network. The lightpaths usually carry several lower rate circuits so that its bandwidth is used efficiently. An LSR, as shown in the figure, is connected to an access network. The access networks sink and source traffic to and from the backbone network. To the backbone network it appears as if the LSR is generating all the traffic. Thus, in the current context a 'node' refers to the OXC and the LSRs it is connected to. Since the abstraction does not alter the problem, such an assumption is made for the purpose of succinctness.

2.2 Virtual Topology Design

Often computer networks are viewed from a graph theoretic viewpoint as a graph. The vertices of such a graph typically represent routers, switches etc., and the graph edges the physical links. The graph equivalent of a computer network denotes the physical topology of the network. In optical WDM networks, the graph denoting the physical topology of the network has an edge between nodes if there exists an optical fiber connecting the nodes. Selectively optically bypassing certain wavelengths at certain nodes in optical networks, introduces the concept of *logical* or *virtual* connectivity between the endpoints of a lightpath; while the two nodes are not physically connected directly, the absence of any interruption or electronic conversion in the channel between the nodes *logically* connects them. A graph connecting nodes that are logically or virtually connected can be obtained for a configured optical network. A graph thus obtained is said to denote the *logical* or *virtual topology* of the network. The logical topology hides the details of the physical topology. A given physical topology can be used to realize several logical topologies and similarly, a given logical topology can realized using several different physical networks.

One of the major concerns for service providers and an active research problem has been to design a logical topology (or stated differently to determine what lightpaths need to be set up) that can satisfy the requirements of the upper layers. The virtual topology design problem has similarities with several topology design problems [17]. An elaborate survey of research in the area of virtual topology design can be found in [17]. In general, the virtual topology design problem can and has been formulated as an optimization problem optimizing some metric of interest such as throughput,

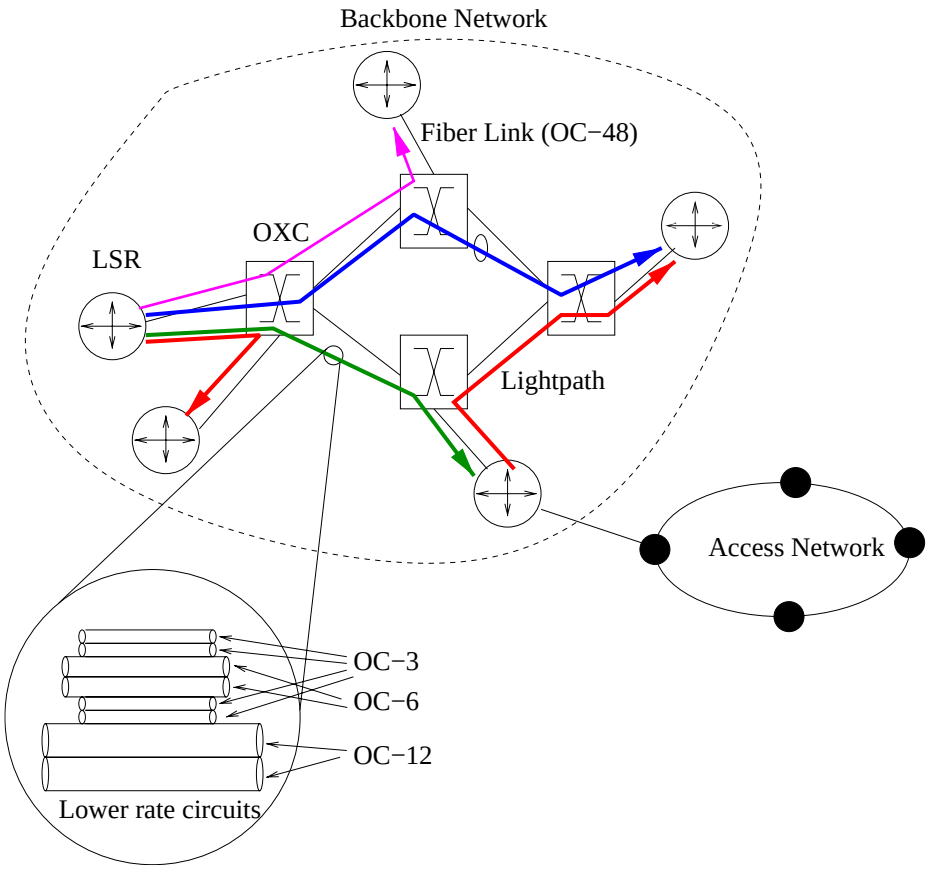

Figure 2.5: A typical optical network

amount of electronic routing or the amount of LTE used.

The fact that a given physical topology can be used to realize several logical topologies, implies that optical networks lend themselves very well to reconfigurability. Indeed, reconfigurability is seen as one of the salient strengths of optical networks. A virtual topology is designed on the basis of observed patterns in traffic for a given physical topology. Being able to realize a different virtual topology on the same physical topology provides adaptability (when traffic patterns change), self-healing capability (when physical topology changes due to failure of components) and upgradability. The virtual topology design problem is known to be NP-complete because it is in essence an m-commodity flow problem. Hence, several heuristic solutions have been proposed in literature.

2.3 Traffic Grooming

With the deployment of commercial WDM systems, it has become apparent that the cost of network components, especially line terminating equipment (LTE), is a dominant cost in building optical networks, and is a more meaningful metric to optimize than, say, the number of wavelengths. Furthermore, with currently available optical technology, the data rate of each wavelength is of the order of 2.5-10 Gbps, while channels operating at 40 Gbps will be commercially available in the near future. In order to efficiently utilize the capacity of each lightpath, a number of independent lower-rate traffic streams must be multiplexed into a single lightpath. These observations give rise to the concept of **traffic grooming**, which refers to the techniques used to combine lower speed components onto available wavelengths in order to meet network design goals such as cost minimization. Traffic grooming has received considerable attention recently. For reviews of the traffic grooming problem area, see [19]; it also provides a detailed example and a precise formulation as an ILP.

An instance of the traffic grooming problem is specified by, (i), the traffic demand matrix specifying the amount of traffic that flows between all nodes in the network, (ii), the number of wavelengths supported by each fiber (usually assumed to be the same), (iii), the grooming factor and (iv), a relevant objective (typically cost) function. The elements of the traffic demand matrix represent aggregate traffic between the various source-destination pairs in multiples of some base rate and the grooming factor represents the bandwidth of a single wavelength in multiples of the same unit. An instance of the traffic grooming problem can conceptually be considered to comprise of the following three subproblems (SPs):

1. **Virtual topology SP:** The SP comprises of determining a set of lightpaths that need to be set up to satisfy the given traffic requirements. The set of lightpaths determine the virtual topology to be realized and can be represented as a directed graph, $G_v(V, E_v)$ with E_v comprising of unweighted edges between nodes that are connected by a lightpath and V comprising of vertices each of which corresponds to a node in the network.

2. **Lightpath routing and wavelength assignment SP:** The RWA SP concerns assigning a path of links and a wavelength on each of the links for each of a multiset of lightpaths. The virtual connectivity graph along with lightpath routing and wavelength assignment is denoted by $\mathcal{T}$. One version of the RWA SP assigns the same wavelengths in all the links of a lightpath. Such an assignment is said to obey the *wavelength continuity constraint*. The version of RWA SP that does not enforce the constraints requires the nodes to be equipped with *wavelength converters* to convert signal on wavelength to another. The RWA SP is NP-hard in general network topologies [17]. Polynomial time algorithms exist to solve the RWA SP for simpler elemental topologies such as the path, star [63] and spider networks [64].

3. **Traffic routing SP:** The traffic routing SP concerns routing the sub-wavelength traffic demands of a source-destination pair to one or more lightpaths set up. This mapping of traffic

on the virtual topology $\mathcal{T}$ is denoted by $\mathcal{G}$. This traffic routing SP too is inherently hard. Even in the simplest physical topology, the path network, in which the RWA SP can be solved in polynomial time, the traffic grooming problem is NP-hard even when a candidate virtual topology is provided [15]. The fact that even when two of the SPs of traffic grooming are eliminated (virtual topology design and RWA) traffic grooming is hard, implies the traffic routing problem is inherently hard, even in its simplest form.

While the traffic grooming problem comprises of these SPs, an optimal solution can be obtained only by considering all the three SPs in an integrated manner. However, due to the inherent hardness of the problem the practical approach taken is to solve the SPs separately and then combine the sub-solutions to obtain a feasible solution. Note that arbitrarily solving the three SPs separately and then trying to combine may not in general provide feasible solutions.

2.3.1 Traffic Grooming Literature Survey

In this section we briefly review and survey recent literature in the area of traffic grooming in general.

The static version of the problem is a network design problem—a problem of how best to configure the network to support a static traffic matrix. The static traffic matrix may be obtained from estimates of traffic in the network. The goal is to minimize some measure of network cost, such as the number of LTEs used for the design. The dynamic version of the problem, on the other hand is mostly one of how best to set up connections (lightpaths), usually again to minimize network equipment used. However, minimizing blocking probability may also be a concern. Most of the research considers the network design aspect of grooming rather than blocking characteristics for dynamic traffic. Most studies have considered the static version of the problem [11, 12, 18, 24, 26, 39, 41, 54, 55, 60, 61, 65]. There have been a few studies that have considered the dynamic version of the problem [9, 26].

The grooming problem being NP-complete, all solutions proposed in the literature are sub-optimal, in general but for specific scenarios (however unrealistic) optimal solutions are obtainable and have been proposed [10]. The approaches in literature are mainly one of three kinds: heuristics specific to traffic grooming, decomposition based approaches and solutions based on transforming the grooming problem to a different well known problem and then employing heuristic solutions proposed in those contexts.

Several heuristic solutions that specifically address grooming (as opposed to those that employ meta-heuristics like genetic programming etc.) [9, 10, 26, 61, 66] have been proposed in literature. Typically, the performance of the heuristics is studied for different traffic patterns and compared with theoretical bounds. A few of the studies have accorded a rigorous theoretical analysis of the heuristic solutions proposed [60] and provided approximation ratios for the heuristic algorithms.

As stated earlier the grooming problem comprises of three other hard problems which need to be considered in an integrated manner to obtain the optimal solution. The decomposition based approaches, break down the grooming problem by ignoring some of the constraints of the problem and considering the three problems in isolation. This significantly reduces the complexity of the problem but at the cost of sub-optimality. Some of the constraints become trivial or computationally tractable for certain topologies e.g. the routing constraints in the RWA SP for the unidirectional ring or path network. Wavelength assignment is solvable in polynomial time in path networks [29]. Thus, the RWA SP is completely eliminated in path networks. One of the approaches in literature has been to consider complex topologies as consisting of simpler topologies (solving which is relatively easier) and to combine solutions to several instances of the simpler topology to obtain an approximate solution for the more complex topology. Finally, some of the studies employ meta-heuristics such as simulated-annealing and genetic algorithms to solve the grooming problem [61].

2.4 Reconfiguration

One of the key features of multi-wavelength optical networks is rearrangeability, i.e. the ability to dynamically optimize the network and this ability is a consequence of the independence between the virtual and the physical topology as discussed in Section 2.2. A virtual topology would not be more useful than a physical topology if it was as incapable of change as the latter. The primary advantage of a virtual topology is that it can be changed on demand much more easily than the physical topology. Reconfiguration of a virtual topology may be carried out for two reasons: (i) To reoptimize the virtual topology under a changed traffic pattern, or even a changed cost metric, (ii) To create a new topology capable of serving the current traffic pattern, such that the new topology does not use some network components which have failed or are being withdrawn from service. Reconfiguration to recover from failure is usually called restoration.

In this thesis we focus on reconfiguration for reoptimization, when the traffic pattern changes. The objective function value determines how best the virtual topology is suited for the given traffic demand. With the change in the traffic pattern the objective function may not remain optimal. The virtual topology needs to be changed to optimally reflect the new objective. This change requires reconfiguration of the network components (OXCs and DXCs) to establish the lightpaths present in the new virtual topology but absent in the current topology. Similarly the lightpaths absent in the new virtual topology must be deleted. The metric of reconfiguring the OXCs and the DXCs appropriately reflects the amount of delay experienced by the traffic being switched. The reconfiguration solution is a tradeoff between the objective function value related to the general objective function value and the number of changes to the virtual topology. Some of the other issues involved in reconfiguration are discussed below.

1. *Disruption to the traffic:* The number of lightpath changes in the process of transition to

the new topology determines the extent of disruption to the traffic in the network. Since the traffic carried by the lightpaths is of the order of gigabits per second, the disruption in the traffic needs to be minimized. A goal of reconfiguration process is to minimize the impact of reconfiguration on the carried traffic, making the transition from the old virtual topology to the new one transparent to the users.

2. *Network availability during reconfiguration:* The transition to the new virtual topology can either be done by reconfiguring all the network components concurrently or by applying step-by-step changes until the new topology is implemented. For the first approach a significant amount of the network may be unavailable during the transition period. But if the network must continue to operate during the process of reconfiguration (which is more practical), shutting down the links and reconfiguring each OXC etc. is not permitted. Then the transition from the old virtual topology to the new topology must be carried out in a sequence of reconfiguration steps, where network component reconfigurations in any given step can be performed concurrently. Only a few nodes will be shut down at a time, but the overall transition to the new virtual topology would take longer to complete in this case.

3. *Time spent on reconfiguration:* During the time period when changes in the optical layer occur, the lightpaths involved in transition cannot be used. Also, the network is very vulnerable during the transition period to the failure of the protocol controlling the reconfiguration process. Thus, the aim of reconfiguration process is to reduce the total time spent on reconfiguring to the new topology. However, there is a tradeoff between the time required to reconfigure and the network availability during the process of transition.

The virtual topology design problem is computationally intractable [17,57]. The reconfiguration problem not only optimizes the objective function for the virtual topology design problem but also tries to minimize the number of changes. Hence, reconfiguration problem is also computationally intractable.

2.4.1 Reconfiguration Literature Survey

In this section we briefly review the literature on reconfiguration studies.

The possible approaches to this problem as found in literature are discussed below.

1. The first approach is to solve the virtual topology design problem for the new traffic, without any consideration of the current virtual topology. The second step in this approach is to implement the new virtual topology with the aim of minimizing the number of changes from the current virtual topology. This approach always results in an optimal value with respect to the objective function at the cost of the number of changes to be made.

2. The second approach is to solve the the virtual topology design problem for the new traffic at the same time minimize the number of changes to be made. This approach is NP-complete in nature [5]. Several heuristics have been proposed to solve this problem.

The literature on reconfiguration can be viewed as falling in two broad categorizations, namely, reconfiguration of broadcast optical networks and reconfiguration of wavelength routed optical networks (WRONs). A broadcast optical network has a single broadcast medium which is accessed by all the nodes in the network. On the other hand, a WRON is a WDM network where signals at intermediate nodes are switched on the basis of their wavelengths. The virtual topology design problem is different for broadcast and wavelength routed optical networks because with a broadcast medium, the physical topology does not constraint the virtual topologies that can be implemented. Further, there is no possibility of wavelength reuse in broadcast networks since each lightpath needs a unique wavelength. Accordingly, the reconfiguration issues are different, and we discuss the literature below under these different heads.

As mentioned above, the two conflicting objectives of the reconfiguration process are the optimal objective function and the number of changes required to achieve the new virtual topology. The various objective functions considered are average packet delay [20,34,36,40,50], total network throughput [2,5,6,33,35,45–47,59] the average weighted hop distance [7,21,48,49,51,53]. The frequency with which reconfiguration is undertaken also varies; it may be at regular intervals [31,45–47], triggered by a given amount of change in the traffic matrix or on-demand [7,20,33–35,48,51], upon network events such as packet loss or even packet arrival, etc.

2.4.1.1 Broadcast Optical Networks

Initial interest in reconfiguration was concentrated on broadcast optical medium LANs [5,6,21,35–37] where every station accessing the shared medium has a given number of optical transmitters and receivers, and reconfiguration consists of retuning these. Broadcast networks where traffic is sent from one node to another using a sequence of lightpaths are also known as multihop networks. The issues involved in reconfiguring multihop networks by retuning a set of slowly tunable transmitters or receivers have been studied in [36,37]. In [36], the problem of obtaining a virtual topology that minimizes the maximum link flow, given a set of traffic demands, was studied. While in [37], algorithms were developed for minimizing the number of branch-exchange operations required to achieve a target virtual topology, when traffic changes. This was an attempt at logically reconfiguring the network in a way that is minimally disruptive to the traffic, since only two links were disrupted at any given time. A branch-exchange operation can be defined as the swapping of destination nodes between two source-destination pairs. The problem of reconfiguring a single-hop network by retuning a subset of slowly tunable receivers in response to changing traffic and network conditions was studied in [5]. The reconfiguration problem was formulated as a Markov decision process in [6],

with the degree of load balancing and the number of retunings as the two conflicting objectives. A survey of reconfiguration issues arising in broadcast optical networks has been done in [6].

2.4.1.2 Wavelength Routed Optical Networks

More recently, attention has shifted to wavelength routed wide area networks and the reconfiguration issues involved in wavelength routed wide area networks [2, 4, 7, 20, 23, 28, 33, 34, 40, 44–48, 50, 51, 53, 56, 57, 59], where the reconfiguration steps consist of reconfiguring the optical switching state of the OXCs in the network. The different approaches for solving the reconfiguration problem as found in literature are either exact solutions or heuristics. We briefly describe these approaches here.

Exact solutions employing Integer Linear Programming (ILP) formulations have been proposed in [7,33,36,51]. The approach in [7] is to obtain the new virtual topology without considering the current virtual topology. Followed by the effort to minimize the number of changes required in the implementation of the new virtual topology. This method guarantees an optimal virtual topology for the new traffic, but with no provision for the tradeoff between the number of changes and the objective function value. In [51], a modification to the formulation given in [7] was proposed to include an upper bound for the number of reconfiguration steps in the formulation. However, these exact solutions become intractable for large networks.

Several heuristics have been proposed that try to minimize the disruption to the network by making small elemental changes to the virtual topology at each step [23, 45–47, 59]. In [46], an iterative reconfiguration algorithm has been proposed that reduces the maximum link load of the network for rapid changes in the traffic pattern. The "dynamic single step optimization" algorithm proposed here, implemented a single 3-branch exchange or a 2-branch exchange at regular intervals, to minimize the disruption caused to the network. The adaptation mechanism proposed in [23], assumed that the future traffic pattern was not known and the fluctuations in the traffic are adopted by adding or deleting a single lightpath at a time. This algorithm tried to delete a lightpath only after the traffic on that virtual link was rerouted. This was an attempt to reconfigure in a "hitless" (minimal disruptive) manner. The algorithm proposed in [59] also made only a single lightpath change at every step and aimed at maximum network availability.

In [56] reconfiguration has been proposed as an on-line process based on a two stage approach. The first stage, reconfiguration stage, determined the changes to be made in the virtual topology with the aim of minimizing the objective function value and limiting the number of changes within a pre-specified value. Due to the continual application of heuristics to reconfiguration, the objective function value might deviate from the optimal value. Thus, the second stage in the proposed algorithm, optimization stage, was used to bring the objective function value of the virtual topology closer to the optimal value. A reconfiguration algorithm for optical ring networks has been proposed in [20]. This algorithm was based on the concept of splitting and merging existing lightpaths, and

cost-benefit analysis was used to reduce the network reconfiguration cost. The objective of the algorithm was to reduce the number of lightpaths to be reconfigured, while the network congestion was kept low. In [48], a heuristic tool known as the Virtual Permanent Reconfiguring Mechanism has been proposed. This tool was used to dynamically reconfigure wavelength connections in a WDM transport network with a given physical topology. A standard-deviation threshold detection scheme was used to detect load distribution changes and a new optimum logical connectivity was computed whenever the threshold limit was reached.

While different authors take different approaches to the reconfiguration problem, we note two commonalities: (i) Usually routing of sub-wavelength traffic on the virtual topology is not integrated into the reconfiguration problem, and reconfiguration is not allowed to change the routing of traffic on lightpaths, (ii) as a consequence, if the optimality of the new virtual topology is considered inside the scope of the problem, it must be possible to express such optimality in terms of lightpath traffic, without referring to sub-wavelength traffic.

The reconfiguration heuristic proposed by us considers the reconfiguration of sub-wavelength traffic on lightpaths as opposed to the reconfiguration of the entire lightpath. To the best of our knowledge, there is not a single study in literature on reconfiguration that considers sub-wavelength traffic. Our reconfiguration heuristic for WRONs, discussed in Chapter 5, tries to achieve minimal disturbance reconfiguration by performing frequent local changes accompanied with infrequent global optimal changes.

2.5 Our Contribution

Our primary contribution is in recognizing the need for investigating the usefulness of an integrated approach for reconfiguration of optical networks with sub-wavelength traffic. The overall focus of this work has been the balancing of the reconfiguration cost (disturbance to the network) and a good grooming solution. We define suitable goals for an integrated approach, and provide formulation of the integrated approach as an integer linear program. To allow a joint consideration of grooming and reconfiguration costs, we mathematically formulate a representation of the reconfiguration cost in terms of the OXC and DXC reconfiguration that can be related to grooming costs.

We also develop a heuristic strategy for addressing the problem, which attempts to achieve minimal disturbance reconfiguration by performing local reconfiguration and delaying the need for global reconfiguration. In order to proactively avoid global reconfiguration, we introduce the concept of over provisioning at the traffic demand level. Our reconfiguration heuristic minimizes the need for solving the integer linear program, which is computationally intractable.

Chapter 3

Problem Definition

In this chapter, we discuss the need for an integrated approach for reconfiguration of a network where the traffic is of sub-wavelength nature. We also define the integrated problem and the various assumptions and notations used in this thesis.

3.1 Integrated Approach

Reconfiguration strategies that do not account for sub-wavelength traffic are of limited or little use in groomed networks. Such strategies are sufficient when the bandwidth of the typical wavelength channel provided by the optical technology of the time is small enough that it is reasonable to think of source-destination traffic on the backbone in units of lightpaths. With the coming age of Tbps (and beyond) this may no longer be true. Thus, reconfiguration problem needs to be attended for sub-wavelength traffic networks. Independently operating grooming and reconfiguration strategies, however good in themselves, are not likely to work well together. Indeed, they may end up working against each other. A reconfiguration strategy that seeks to minimize the number of reconfiguration steps may dictate a final virtual topology that results in very high OEO switching, thus destroying the good grooming characteristics obtained by the grooming strategy.

From this point of view, there are two shortcomings of past reconfiguration strategies: (i) They cannot rearrange traffic routing over lightpaths, (ii) Temporary disruption to traffic is considered in terms of lightpath disruption, not source-to-destination traffic. Several approaches to reconfiguration in literature are based on multiple phases [46, 50, 56, 57] (the phases may be identical to the steps we mentioned in Section 2.4). The objective of such phasing is often to minimize the disruption experienced by existing traffic. Spare store-and-forward (i.e. OEO switching) capability

in the network at the time has been assumed sufficient to protect the traffic from disruption for broadcast LANs [6, 37]. With the much higher speed and bandwidth of the optical networks of tomorrow, this may be reasonably assumed for each phase of the reconfiguration, but not the overall reconfiguration process (because variability of delay and probability of loss would be correspondingly higher for the much longer time scale of all the phases combined). Moreover, OEO switching is exactly what we want to minimize from a grooming standpoint. Finally, for a groomed network, we need to specify not only the final topology (with RWA) but also the grooming solution (with traffic routing), and disruptions to traffic streams must take into account all disruptions to source-to-destination traffic (some of which may not involve disruptions to lightpaths because of OEO switching). The above considerations force us to recognize a new problem which cannot be addressed by existing strategies: reconfiguration of a topology as well as traffic assignment in a groomed network, to minimize or limit the disruption caused to groomed traffic components in each phase, and no loss of traffic between phases.

3.2 Cost Calculation

In this section we discuss the various cost functions considered for traffic grooming and reconfiguration in literature. We also formulate a cost function for our integrated approach.

3.2.1 Grooming and Reconfiguration Cost Calculation

Grooming Cost Functions: In the existing literature, grooming cost minimization is addressed by formulating the problem so that its objective reflects the network cost incurred in electronically forwarding traffic at intermediate nodes. One measure is to directly incorporate the number of LTEs into the cost function. A more fine-grained measure of electronic routing is to focus on the actual amount of traffic (in bytes per unit time) that is electronically routed in any grooming solution. This does not directly capture equipment cost, but it more accurately captures indirect costs such as delay or probability of loss in electronic equipment, and avoids penalizing traffic injection/termination in the network. Either of these measures may be used on the basis of total over all the network nodes, or the maximum which occurs at any node.

In our work, we adopt the total amount of electronic routing in the network as the grooming cost measure. We discuss the usefulness of this measure later in this section.

Reconfiguration Cost Functions: The reconfiguration cost functions used in literature have been of many types. They have been generally expressed in terms of the network equipment that are involved in a reconfiguration, such as the number of wavelength routers that need to have their optical switching reconfigured, or the total number of optical switchings that need to be changed to implement the new lightpaths and eliminate old ones. Other metrics have focused on the number of

elementary reconfiguration operations, such as lightpath pairwise exchanges, that must be performed to migrate from one topology to another.

In general terms, we focus on the number of individual reconfigurations of network equipment as well. However, for reasons we explain below, we investigate some detailed options of defining reconfiguration costs such that they will be suitable for comparison with grooming costs.

3.2.2 Integrated Cost Calculation

In this thesis, we formulate an integrated approach for implementing reconfiguration of wavelength routed optical networks, where the traffic carried by the network is sub-wavelength in nature. To the best of our knowledge, such formulations do not currently exist in literature. We have formulated the integrated approach as an integer linear program (ILP) with the objective of striking a balance between the grooming gain obtained by reconfiguring to the new virtual topology and the reconfiguration cost. The grooming gain is obtained in terms of the reduced amount of electronic switching required by reconfiguring to the new virtual topology under the changed traffic conditions. As discussed in Section 2.1 one of the undesirable effects of OEO conversion is that it introduces delay. This is because the lightpath is terminated at the node where OEO conversion needs to be done. The signal is then extracted from the optical domain to the electronic domain, processed and then converted back to the optical domain. Thus, the grooming cost calculated in terms of total amount of electronic switching models the total traffic weighted delay in the network. On the other hand, we calculate the reconfiguration cost in terms of the number of OXCs and DXCs that need to be reconfigured. Every time an OXC or a DXC needs to be reconfigured, the traffic being switched through them is delayed. Thus, the reconfiguration cost as calculated by us models the total delay experienced by the traffic at the nodes that need to be reconfigured. This provides the common basis on which to combine and compare grooming and reconfiguration costs.

The grooming cost is a constant operating cost to the network while the reconfiguration cost is a transient cost incurred only at the time of reconfiguration. A badly groomed network will continue to incur a high grooming cost until either the network conditions change again or reconfiguration is undertaken. On the other hand, a bad reconfiguration strategy would result in a high reconfiguration cost only at the time of reconfiguring the network. Thus, to be able to reflect the grooming gain effectively in comparison to the reconfiguration cost we multiply it by a factor of γ. γ is therefore related to the average delay between reconfigurations.

In order to avoid frequent reconfigurations for a very little gain in grooming we introduce another parameter δ. We allow a transition from the old virtual topology to a new virtual topology only when the grooming gain is at least δ times more than the reconfiguration cost, thus, avoiding "flapping". Both these parameters γ and δ are highly dependent on the physical topology, traffic pattern being supported by the network, operating policy etc. It is upto the network administrator to determine the values of these parameters based on the characteristics of the network. The deter-

mination of the exact values of these factors for any particular set of conditions is beyond the scope of this thesis work.

3.3 Assumptions

In this thesis we have assumed the following:

- Each node is equipped with WADMs to selectively optically bypass signals of certain wavelength(s).
- Each node is equipped with digital cross-connects so that a traffic component arriving at a node on one wavelength may be switched electronically to a different one as it is routed to its destination.
- The physical links and lightpaths are directed and each physical link is represented by a single fiber.
- *Wavelength continuity constraint* is assumed to hold i.e. a lightpath should be assigned the same wavelength on all the links between its endpoints. From a practical perspective, this means we assume wavelength converters are absent in the network.

The above are standard assumptions made in the literature, which is our reason for making them. However, we make two additional specific assumptions regarding bifurcation of traffic and multiple lightpaths that are more closely related to our work.

- *Bifurcated routing* of traffic [19] is not allowed and the total traffic for any source-destination (s-d) pair is at the most equal to the capacity C of a wavelength (i.e. $t^{(sd)} \leq C$). Disallowing bifurcated routing implies that the entire traffic ($t^{(sd)}$) for a particular s-d pair follows the same lightpath or a sequence of lightpaths from source to destination.
- No more than a single lightpath is allowed between any two nodes.

The need for the last two assumptions is generally in reducing the complexity of the problem to manageable levels. In particular, if multiple lightpaths from a source to a destination is allowed, deciding which of them a particular sub-wavelength traffic component traverses is often not included as part of a realistic grooming solution, and picking an option which is most useful from our point of view is non-trivial. Though this is outside the scope of our work, we demonstrate this briefly in Section 5.1.5. The other assumption is prompted by similar considerations.

3.4 Notations

The notations used for the problem formulation in this thesis are mostly consistent with [19]. We consider an arbitrary mesh topology of N nodes, with the nodes numbered from 0 to $N-1$. Each directed link in the topology has a unique link id. All the nodes are assumed to be equipped with wavelength add/drop multiplexers (WADM) (similar to the one shown in Figure 2.2), which can pass each wavelength through optically, or add/drop it from/to an electronic cross-connect.

The network under study supports dynamic (i.e. time varying) traffic streams of lower rates that are multiplexed onto higher capacity lightpaths. We assume that the traffic demand placed on the network can be reasonably represented in an aggregate static manner, as is likely to be the case for some backbone networks, then the virtual topology is also static. Let us denote this initial topology (with RWA) by $\mathcal{T}$ and the initial groomed solution (with traffic routing) by $\mathcal{G}$. When the aggregate traffic changes beyond some significant amount over a long time scale (hours or days), the problem is to find a different virtual topology $\mathcal{T}'$ as well a final groomed solution $\mathcal{G}'$ which not only requires the least reconfiguration effort but also gives a good grooming solution under the changed conditions.

The endpoints of a physical link are denoted by l, m and s, d denote the source and destination of the traffic while i, j denotes the endpoints of a lightpath. The *lightpath indicator* b_{ij} is a binary variable that denotes whether a lightpath exits from node i to j, while $b_{ij}(l,m)$ indicates if such a lightpath traverses the physical link from node l to m. The *link indicator* p_{lm} is 1 if a physical link exists between nodes l and m. The *traffic indicator* $d_{ij}^{(sd)}$ is 1 if the lightpath from i to j carries traffic from source s to destination d. The link-path-wavelength indicator $c_{ij}^{(k)}(l,m)$ is 1 if a lightpath from node i to j on wavelength k traverses the physical link from node l to m. The traffic demands between pairs of nodes in the physical topology are given by the traffic matrix $T = [t^{(sd)}]$. We assume that the network supports traffic streams at rates that are a multiple of some basic rate (e.g., OC-3, 155 Mbps). We let C denote the capacity of each wavelength expressed in units of this basic rate and W be the total number of wavelengths supported by each fiber. Thus, C denotes the maximum number of traffic units that can be multiplexed on a WDM channel (wavelength). For example, if each wavelength runs at OC-48 (2.5 Gbps) rates and the basic rate is OC-3, then $C = 16$. Each quantity $t^{(sd)} \in \{0, 1, 2, \cdots, C\}$ is also expressed in terms of the basic rate, and it denotes the number of traffic units that originate at node s and terminate at node d. The lightpath from node i to j carries aggregate traffic t_{ij} and $t_{ij}^{(sd)}$ denotes the amount of this traffic which is due to the demand $t^{(sd)}$. Thus, virtual topology solution $\mathcal{T}$ is denoted by the variables $c_{ij}^{(k)}(l,m)$, b_{ij} and $b_{ij}(l,m)$. While a groomed solution $\mathcal{G}$ is given by $d_{ij}^{(sd)}$, $t^{(sd)}$ and $t_{ij}^{(sd)}$.

Chapter 4

Integrated Approach Formulation

In this chapter, we formulate an integrated approach for performing reconfiguration of a sub-wavelength groomed network topology. Given the physical topology, the initial traffic matrix T and the changed traffic matrix T' and an initial optimal groomed solution $\mathcal{G}$, we calculate a new groomed solution $\mathcal{G}'$ that balances the grooming gain against the reconfiguration cost incurred in the transition from $\mathcal{G}$ to $\mathcal{G}'$.

4.1 Reconfiguration Cost Calculation

As we mentioned in Section 3.2, there are good metrics for grooming cost available in the literature, but no good metric to reflect global reconfiguration cost has been defined. Accordingly, in formulating the integrated problem, we first propose a method to capture network reconfiguration costs. The effort in reconfiguring the network consists of the effort in reconfiguring the network components. Therefore, the reconfiguration cost incurred in the transition from the old grooming solution to the new one is calculated in terms of the number of OXCs and DXCs that need to be reconfigured. This approach is based on the fact that the establishment of a new lightpath absent in the old virtual topology but present in the new virtual topology and/or similar elimination of an old lightpath, requires reconfiguration of the OXC as well as the DXC. The OXC has to be reconfigured because a wavelength channel has to be switched to/from DXC port, and the DXC because it has to switch the tributary traffic components in the wavelength channel. On the other hand, a change in the optical switching of a lightpath in the new virtual topology when compared with the old topology requires only the OXC to be reconfigured. We describe below a procedure for reconfiguration cost calculation by obtaining a matrix representation for each node for the old groomed solution $\mathcal{G}$ and

the new one $\mathcal{G}'$ and then calculating the distance between these two matrices in order to determine the number of OXC and DXC changes.

4.1.1 Matrix Representation of Each Node

As shown in Figure 2.4, an optical node is represented by an OXC attached to a DXC. Here, we represent the same structure of the node as a matrix, in terms of the optical or electronic switchings performed at that node. Each node n in the physical topology is represented by a $W \times N_s$ matrix, where N_s represents the total number of OXC inputs. N_f are the total number of incoming optical fibers for node n, while $N_s - N_f$ are the total number of DXC outputs (i.e. the total number of lightpaths originating from a DXC on each wavelength). The maximum number of useful DXC inputs/outputs for n are $W\times$ {maximum of incoming or outgoing fiber links at n}. Applying the following rules we determine whether a lightpath is being electronically generated or terminated or is being optically switched, using the link-lightpath-wavelength indicator variables $\{c_{ij}^{(k)}(l,m)\}$ from the groomed solution $\mathcal{G}$. To determine the switching state of node n, only those $c_{ij}^{(k)}(l,m)$ values need be considered in which either l or m is equal to the node number n. Moreover, among these, only the indicators which are non-zero (that is, have a value of 1) indicate some actual switching action.

- The electronic generation of a lightpath at node n on wavelength k is indicated by an indicator with $(n = l = i)$.
- Electronic termination of a lightpath at node n on wavelength k is indicated by $(n = m = j)$.
- An optical switching of a lightpath at a node n is indicated by $((n = m) \neq j)$. In that case a pair of non zero $c_{ij}^{(k)}(l,n)$ and $c_{ij}^{(k)}(n,m)$ values exist which can be used to determine that the traffic from fiber link incoming from node l is being switched to the outgoing fiber to node m.

We now represent the switching state of the node n by a matrix. The row of the matrix represents wavelength, and the column represents incoming ports, either from fiber link or the DXC. First, a global identification for each directed fiber link in the network is established.

- *Electronic Termination:* We represent this by a symbol 'O'. The symbol 'O' in the matrix at row k and column f_i, indicates that the traffic on wavelength k for incoming fiber f_i is electronically terminated at the node n.
- *Electronic Generation:* We represent this by a fiber number f_o in one of the columns from $N_f + 1$ to N_s in row k. The columns $N_f + 1$ to N_s represent DXC outputs per wavelength, thus a fiber number f_o in any of these columns means that the traffic on the wavelength k for the outgoing fiber f_o is electronically generated at the node n.

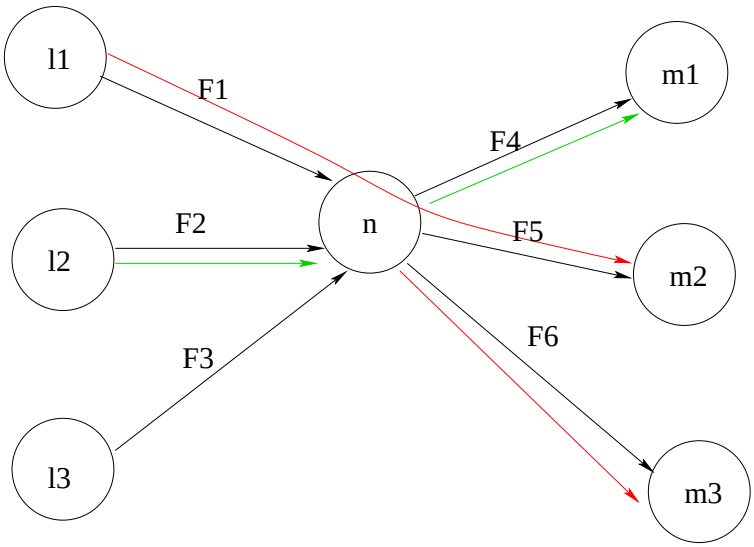

Figure 4.1: Physical and virtual connectivity for node n

- *Optical Switching:* We represent this by a fiber number f_o in one of the columns from 1 to N_f in row k. A fiber number f_o in column f_i, $1 \leq f_i \leq N_f$, represents that the traffic on wavelength k is switched from the fiber f_i to fiber f_o.
- *Don't Care:* We represent this by a symbol 'X'. The symbol 'X' in the matrix at row k and column f_i, represents that the wavelength k on incoming fiber f_i does not carry any traffic for/through the node n.

According to these rules, in each row an outgoing fiber number f_o or an 'X' would be valid entries in any column. The symbol 'O' is valid in only the first N_f columns, but not in the later $N_s - N_f$ columns. We demonstrate the above explained rules using the example shown in Figure 4.1 which shows the physical connectivity and the virtual topology for node n. There are three incoming fiber links $F1$, $F2$ and $F3$ to node n, and three outgoing fiber links $F4$, $F5$ and $F6$ from node n. Let the solid red line represent wavelength number 1, while the dotted green line represents wavelength number 2. We consider only those non zero $c_{ij}^{(k)}(l,m)$ values from the grooming solution, where n is one of the endpoints for the physical link (l,m). $c_{l1,m2}^{(1)}(l1,n) = 1$, $c_{l1,m2}^{(1)}(n,m2) = 1$, $c_{l2,n}^{(2)}(l2,n) = 1$, $c_{n,m1}^{(2)}(n,m1) = 1$ and $c_{n,m3}^{(1)}(n,m3) = 1$, these values are used to obtain the matrix representation for node n, which is given below.

$$\begin{array}{c} \\ 1 \\ 2 \end{array} \begin{array}{c} \begin{array}{ccccc} F1 & F2 & F3 & dxc1 & dxc2 \end{array} \\ \begin{pmatrix} F5 & X & X & F6 & X \\ X & O & X & F4 & X \end{pmatrix} \end{array} \tag{4.1}$$

Here $dxc1$ and $dxc2$ represent the DXC inputs (assuming there are two DXC inputs/outputs per wavelength). The traffic on wavelength 1 on $F1$ is optically switched to $F5$, while traffic for $F6$ is electronically generated at node n. The traffic on wavelength 2 on $F2$ is electronically terminated at node n and traffic for $F4$ is electronically generated on wavelength 2, all the other states are don't cares for node n.

4.1.2 Distance Calculation between Matrices

The matrix representation of the nodes as defined above, changes with any change in the virtual topology. Whenever the traffic matrix changes, requiring a virtual topology change, the amount of reconfiguration effort required in performing this change is obtained by calculating the distance between the old and the new matrix representations for each node. Using the following steps we determine the distance between the matrices for a node n, in terms of whether an OXC and/or DXC would require reconfiguration or not.

In our model, a change from any state to a don't care state does not require any reconfiguration effort while a change from a don't care to any other state always requires reconfiguration. This is a reasonable assumption because when the state of a node for a lightpath changes to don't care it means no one is receiving the data from that lightpath. Hence, we do not need to reconfigure the OXC/DXC for such cases and the OXC/DXC can continue switching the data on the don't care lightpath without affecting any other lightpaths. But every state change from a don't care to any other state must now be assumed to require reconfiguration since we cannot assume anything about the state of the node before the don't care. We found this to be a practical assumption because maintaining all the don't care states would not only increase the complexity, but would also be impossible to implement if the initial grooming solution did not provide the information about the don't care states. We make the following observations regarding reconfiguration which are generally true:

1. Elements in column numbers 1 to N_f : The elements in these columns are compared position wise in the old and new matrices.

   - Any state (fiber number or 'X' or 'O') changing to another fiber number : This is a case of optical switching, thus requires reconfiguration of the OXC only while the DXC is not affected.
   - Any state (fiber number or 'O') changing to a Don't Care X' : This neither affects the OXC nor the DXC.
   - Any state (fiber number or 'X') changing to a 'O' : This is a case of electronic termination, thus requires both the OXC and the DXC to be reconfigured.

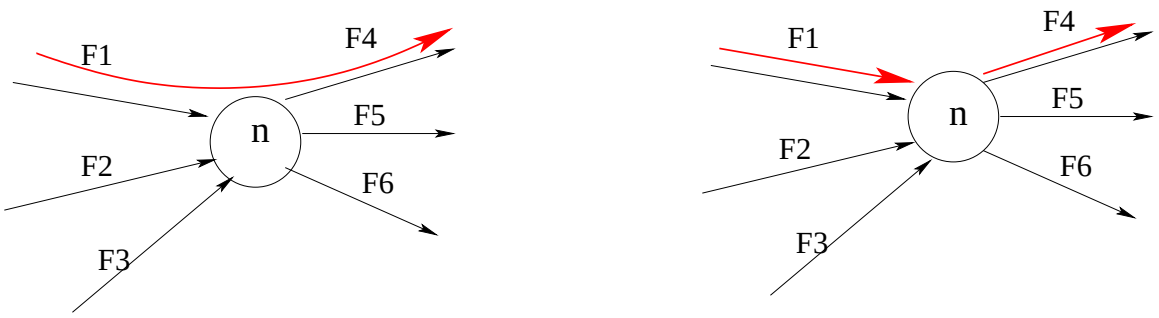

Figure 4.2: Old and new virtual topology for node n

2. Elements in column numbers N_f+1 to N_s : These columns are for DXC outputs (i.e. electronic generation) and contain either a 'X' or a fiber Number. These elements are not compared position wise, as a fiber number in any of these columns means that traffic on that fiber link is generated electronically at node n. The reason for ignoring the exact position for these columns is that the DXC switches the signals sent by the OXC from the input ports to the output ports which are connected to the OXC. The DXC does not switch the signals from an incoming fiber to an outgoing fiber. It is the OXC that is responsible for switching the signals from the DXC to the outgoing fibers.

   - A fiber Number absent in the new matrix (i.e. fiber Number changing to a Don't Care 'X'): This neither affects the OXC nor the DXC.
   - A fiber Number present in the new matrix (but absent in the old matrix): This is a case of electronic generation. Electronic generation will affect both the OXC and the DXC, but if a lightpath is terminating and originating electronically at node n on the same wavelength then we consider this to be a single change to the OXC and a single change for the DXC. We do not consider this as two separate changes because most realistic OXC/DXC fabrics can be reconfigured to implement one lightpath deletion and one establishment in one reconfiguration and this is also in keeping with other relevant work found in literature [37].

   We illustrate the change in the matrix representation for node n corresponding to the old and the new virtual topologies shown in Figure 4.2 by the following matrix representations.

$$\begin{array}{c} \\ 1 \end{array} \begin{array}{c} \begin{array}{ccccc} F1 & F2 & F3 & dxc1 & dxc2 \end{array} \\ \begin{pmatrix} F4 & X & X & X & X \end{pmatrix} \end{array} \Longrightarrow \begin{array}{c} \\ 1 \end{array} \begin{array}{c} \begin{array}{ccccc} F1 & F2 & F3 & dxc1 & dxc2 \end{array} \\ \begin{pmatrix} O & X & X & F4 & X \end{pmatrix} \end{array}$$

The traffic on wavelength 1 on fiber $F1$ is optically switched to fiber $F4$ in the old virtual topology, while in the new virtual topology traffic on wavelength 1 on fiber $F1$ is electronically

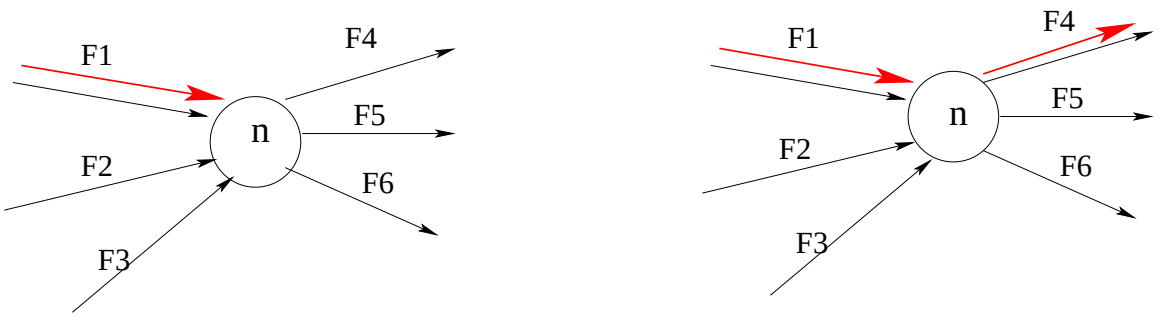

Figure 4.3: Old and new virtual topology for node n

terminated at node n and then electronically generated on fiber $F4$ on wavelength 1. Such a change is counted as a single change to the OXC and the DXC. Another example when the OXC and the DXC are reconfigured just once is shown in Figure 4.3 and the following matrices represent the state of node n in the old and the new virtual topology.

$$\begin{array}{c} \\ 1 \end{array} \begin{array}{c} \begin{array}{ccccc} F1 & F2 & F3 & dxc1 & dxc2 \end{array} \\ \begin{pmatrix} O & X & X & X & X \end{pmatrix} \end{array} \Longrightarrow \begin{array}{c} \\ 1 \end{array} \begin{array}{c} \begin{array}{ccccc} F1 & F2 & F3 & dxc1 & dxc2 \end{array} \\ \begin{pmatrix} O & X & X & F4 & X \end{pmatrix} \end{array}$$

In this example the traffic on wavelength 1 on incoming fiber $F1$ is electronically terminated at node n, while in the new topology the traffic on wavelength 1 on incoming fiber $F1$ is electronically terminated at node n and then electronically generated on outgoing fiber $F4$. Since, this is again an example of a single lightpath deletion and termination, the OXC and the DXC will be reconfigured only once.

We summarize the above as the following distance calculation method between an old and a new matrix.

- Any lightpath established in the new topology, but absent in the old virtual topology requires both the OXC and the DXC to be reconfigured.
- A lightpath optically switched in the new topology, but switched onto a different fiber in the old topology requires only the OXC to be reconfigured.
- A lightpath terminated and originated at a node on a particular wavelength, which was not the case in the old topology requires both the OXC and the DXC to be reconfigured. But this causes a single reconfiguration change to both the OXC and the DXC.

4.1.3 Global Reconfiguration Cost Calculation Methods

As discussed above, we calculate the distance between different matrix representations (under changed traffic conditions) to determine whether an OXC/DXC requires reconfiguration or not. Here we propose four different methods for calculating the reconfiguration cost based on the methodology of reconfiguring the OXC and the DXC.

Different OXC fabrics have different actual implementations [62], which determine what reconfigurations can be performed on the switch simultaneously, or how much reconfiguration effort a given reconfiguration of the switch should be modeled as. For this reason, we try to develop a general framework with several different possible models for the OXC reconfiguration cost. We concentrate more on the OXC reconfiguration cost than on the DXC reconfiguration cost, because signals handled by the OXC, being optically switched, are likely to be more sensitive to the delay caused by the reconfiguration process.

We make the reasonable assumption that the DXC under consideration is a monolithic DXC with limited capacity. Electronic termination and electronic generation result in the same amount of reconfiguration effort/cost. In the first three methods described below we consider a total cost of β for any number of electronic terminations/generations done by a single DXC. That is, the reconfiguration of the DXC is binary in nature. We use the following method in reconfiguration cost calculation I, II and III to determine if the DXC under consideration for node n needs reconfiguration.

- When the total number of 'O' in the new matrix representation is more than that in the old matrix, the DXC needs reconfiguration. We represent this by $is\_dxc\_reconfig_n = 1$; otherwise if DXC does not need reconfiguration, $is\_dxc\_reconfig_n = 0$.
- If $is\_dxc\_reconfig_n = 0$, we now count the total number of fiber numbers in columns from $N_f + 1$ to N_s, if this total is more in the new matrix then the DXC needs reconfiguration. And we represent this the same way, by making $is\_dxc\_reconfig_n = 1$ else $is\_dxc\_reconfig_n = 0$.

Reconfiguration Cost I: In this cost calculation method we consider a cost of α_1 for each OXC reconfiguration. According to this method any number of changes in the new matrix representation for node n, that require OXC reconfigurations would result in a cost of α_1 for node n. This is a coarse-grained representation of the OXC reconfiguration cost. In this case it is assumed that a single reconfiguration change is as costly as multiple reconfiguration changes required by a particular OXC. This is a valid assumption in situations where all the changes required by an OXC are performed simultaneously. If the distance calculation between the old and the new matrix representation for a node n, requires OXC reconfiguration, we represent this by a boolean variable $is\_oxc\_reconfig_n$. If the OXC requires reconfiguration then $is\_oxc\_reconfig_n$ is equal to 1 otherwise its 0. We determine the values of $is\_oxc\_reconfig_n$ and $is\_dxc\_reconfig_n$ for all the N nodes in the network, and then calculate the total reconfiguration cost using the following equation.

$$RC - I = \sum_{n=0}^{N-1} is\_oxc\_reconfig_n \ \alpha_1 + \sum_{n=0}^{N-1} is\_dxc\_reconfig_n \ \beta \tag{4.4}$$

$RC - I$ value varies from zero to $N(\alpha_1 + \beta)$. The cost is zero when no reconfiguration is needed and its maximum when all the N nodes require both the OXC and the DXC to be reconfigured.

Reconfiguration Cost II: We consider the reconfiguration cost of the OXC to be wavelength dependent in this cost calculation method. Any change on a particular wavelength (irrespective of the fiber number) in the new matrix representation that requires OXC reconfiguration will incur cost α_2. This method is based on the assumption that any change on a particular wavelength affects only the $N_s \times N_s$ switch for that wavelength and the other $N_s \times N_s$ switches in the OXC are not affected. This reflects the reconfiguration cost well for an OXC in which the implementation of the fabric is by replicating wavelength planes in the switch, such as the one shown in Figure 2.3. We use the distance calculation rules (Section 4.1.2) to determine if the OXC requires reconfiguration. The boolean variable $is\_oxc\_reconfig\_wavelength_n^k$ is 1 if the OXC at node n requires reconfiguration at wavelength k. The total reconfiguration cost according to this method is given by the following equation.

$$RC - II = \sum_{n=0,k=0}^{n=N-1,k=W-1} is\_oxc\_reconfig\_wavelength_n^k \ \alpha_2 + \sum_{n=0}^{N-1} is\_dxc\_reconfig_n \ \beta \tag{4.5}$$

$RC - II$ varies from zero to $N((W\alpha_2) + \beta)$. The cost is 0 when none of the N nodes require any reconfiguration. The cost is maximum when all the N nodes require both the OXC and the DXC to be reconfigured and the OXCs need reconfiguration on every wavelength.

Reconfiguration Cost III: In this cost calculation method we consider a cost of α_3 for the each OXC reconfiguration. Here the OXC reconfiguration cost is even more fine-grained. Some cost is incurred for every change on each wavelength. This is a suitable method for reconfiguration cost calculation when all the changes in the new virtual topology are implemented individually irrespective of the other changes. The variable used to indicate the total number of changes required by an OXC at node n is $oxc\_reconfig\_count_n$. We increment the value of $oxc\_reconfig\_count_n$ by 1, for every change in the matrix representation that requires reconfiguration of the OXC at node n. The total reconfiguration cost is given by the following equation.

$$RC - III = \sum_{n=0}^{N-1} oxc\_reconfig\_count_n \ \alpha_2 + \sum_{n=0}^{N-1} is\_dxc\_reconfig_n \ \beta \tag{4.6}$$

$RC - III$ varies from zero to $\sum_n N_s W\alpha_3 + N\beta$. The maximum cost is incurred when all the N nodes require the DXC to be reconfigured, and the OXC needs reconfiguration for every wavelength on each column.

Reconfiguration Cost IV: All the three reconfiguration cost calculation methods that we formulated above are good approximations of the total reconfiguration cost incurred by a network in terms of the network components. However, all these cost functions are non linear in nature and

as we discuss in Chapter 6, these cannot be implemented using CPLEX. CPLEX is a software used to solve linear optimization problems and we use this software to prove the correctness and efficiency of our proposed ILP and benchmark the performance of our heuristic against optimal results. Thus, in order to be able to implement the reconfiguration function in the ILP, we have formulated reconfiguration cost IV which is linear in nature. This reconfiguration cost function is used by the ILP formulation in Section 4.3.

In this cost function we count every change in the matrix representations that needs the OXC and/or DXC to be reconfigured as a separate change incurring cost. We assign a cost of α_4 for every OXC reconfiguration and cost β_2 for every DXC reconfiguration. We use the $c_{ij}^{(k)}(l,m)$ values from the new groomed solution and $x_{ij}^{(k)}(l,m)$ values from the old groomed solution to calculate the cost, where $x_{ij}^{(k)}(l,m)$ is 1 is link-path-wavelength indicator from the old virtual topology. As discussed in Section 4.1.3, any lightpath which is electronically terminated in the new virtual topology but was not electronically terminated in the old virtual topology requires both DXC and OXC reconfiguration. Similarly a lightpath which is electronically generated in the new virtual topology but was not electronically generated in the old virtual topology also requires both DXC and OXC reconfiguration. OXC reconfiguration is also required when the optical switching of a lightpath changes from the old virtual topology to the new one.

As discussed in Section 4.1.1, a lightpath is electronically generated at node n on wavelength k when $(n = l = i)$ for a non-zero $c_{ij}^{(k)}(l,m)$ value. Thus, when $c_{nj}^{(k)}(n,m) = 1$ and $x_{nj}^{(k)}(n,m) = 0$, it means a lightpath that is electronically generated in the new virtual topology was absent in the old topology. Based on this observation we formulate the following equation to calculate the total number of electronic generations in the new topology that were absent in the old topology.

$$\sum_{j,k,m} (c_{nj}^{(k)}(n,m) - x_{nj}^{(k)}(n,m))c_{nj}^{(k)}(n,m) \ \forall n \tag{4.7}$$

However, as equation 4.7 is non-linear in nature, thus we formulate the following equivalent of this equation which is linear in nature.

$$r1 = \sum_{j,k,m} (c_{nj}^{(k)}(n,m) - x_{nj}^{(k)}(n,m)) - ((c_{nj}^{(k)}(n,m) - x_{nj}^{(k)}(n,m))x_{nj}^{(k)}(n,m)) \ \forall n \tag{4.8}$$

Similarly, a lightpath is electronically terminated at node n on wavelength k when $(n = m = j)$ for a non-zero $c_{ij}^{(k)}(l,m)$ value. Therefore, when $c_{in}^{(k)}(l,n) = 1$ and $x_{in}^{(k)}(l,n) = 0$, it means a lightpath that is electronically terminated in the new virtual topology was absent in the old topology. Thus, we use the linear equation 4.9 to calculate the total number of electronic terminations in the new topology that were absent in the old topology.

$$r2 = \sum_{i,k,l} (c_{in}^{(k)}(l,n) - x_{in}^{(k)}(l,n)) - ((c_{in}^{(k)}(l,n) - x_{in}^{(k)}(l,n))x_{in}^{(k)}(l,n)) \ \forall n \tag{4.9}$$

Based on our discussion in Section 4.1.1, we know that a lightpath is optically switched at node n when $((n = m) \neq j)$. And a pair of non-zero $c_{ij}^{(k)}(l,n)$ and $c_{ij}^{(k)}(n,m)$ values are used to determine that the traffic from fiber(l,n) is being switched to fiber(n,m). Thus, in the next two linear equations we calculate the total number of optical switchings which have changed i.e. these equations captures all those cases where a lightpath which was being optically switched in the old topology is still being optically switched but on a different fiber in the new topology. It also counts all those cases where a lightpath which was being electronically generated or terminated or both in the old topology is now being optically switched.

$$r3 = \sum_{k,l,m,i\neq n,j\neq n} (c_{ij}^{(k)}(l,n) - x_{ij}^{(k)}(l,n)) - ((c_{ij}^{(k)}(l,n) - x_{ij}^{(k)}(l,n))x_{ij}^{(k)}(l,n))\ \forall n \tag{4.10}$$

$$r4 = \sum_{k,l,m,i\neq n,j\neq n} (c_{ij}^{(k)}(n,m) - x_{ij}^{(k)}(n,m)) - ((c_{ij}^{(k)}(n,m) - x_{ij}^{(k)}(n,m))x_{ij}^{(k)}(n,m))\ \forall n \tag{4.11}$$

The total reconfiguration cost is given by the following linear equation.

$$RC - IV = \phi(\{x_{ij}^{k}(l,m)\}, \{c_{ij}^{(k)}(l,m)\}) = (r1 + r2 + r3 + r4)\ \alpha_4 + (r1 + r2)\ \beta_2 \tag{4.12}$$

$RC - IV$ is linear in nature, and conceptually represents the same effort of reconfiguration that we captured in the previous cost functions. It most closely resembles $RC - III$. However, the "if-then" nature of the OXC/DXC reconfiguration cost calculation that is illustrated in Figures 4.2 and 4.3 is fundamentally non-linearizable. As a consequence, $RC - IV$ has a tendency of over estimating the reconfiguration function when compared to the other three cost functions. $RC - IV$ considers a lightpath termination and generation at a node n on a particular wavelength as two separate changes to both the OXC and the DXC, all the other methods consider this as a single change to both the OXC and the DXC. $RC - IV$ considers every change that requires DXC reconfiguration as a separate change involving a cost β_2 as opposed to the other three cost functions which consider a cost of β for any number of reconfigurations to the same DXC, making further overestimate.

4.2 Traffic Grooming Gain Calculation

As discussed in Section 3.1, the traffic grooming cost can be calculated in various ways. We adopt the "total electronic switching" defined as the grooming cost in [19], because it best reflects the global cost related to total traffic delay which is our focus. This cost function calculates the total amount of traffic that needs to go through OEO conversion at intermediate nodes before reaching the destination. We calculate the grooming cost from the following equation.

$$Grooming\ Cost = \sum_{s,d,i,j} d_{ij}^{(sd)} t^{(sd)} - \sum_{s,d} t^{(sd)} \tag{4.13}$$

This equation calculates the amount of total traffic that flows over intermediate nodes before reaching the destination. The summation of the variable $d_{ij}^{(sd)}$ for a particular s-d pair gives the total number of lightpaths that the traffic from source s to destination d will flow over. We subtract 1 from the total number of lightpaths to give the total number of intermediate nodes between s and d. We then multiply the total number of intermediate nodes with the total amount of traffic from s to d, which gives the total amount of traffic that undergoes OEO conversion i.e. the total amount of electronic switching.

Given an old grooming solution $\mathcal{G}$ and a new grooming solution $\mathcal{G}'$ for new traffic, we formulate the grooming gain in terms of the reduction in the grooming cost achieved by reconfiguring to the new groomed solution from the old one. We calculate this reduction in the grooming cost by subtracting the new grooming cost from the old grooming cost. Here by old grooming cost we mean the total grooming cost incurred when the new traffic is mapped on the old virtual topology. The new grooming cost is given by the equation 4.13, where $t^{(sd)}$ is the new traffic. The old grooming cost is obtained by mapping the new traffic on the old virtual topology because we want to calculate the total amount of electronic switching that would be performed under *current* traffic conditions, if we do not reconfigure to a new virtual topology. We formulate grooming gain as follows.

$$\psi(\{y_{ij}^{(sd)}\},\{d_{ij}^{(sd)}\},\{t^{(sd)}\}) = (\sum_{s,d,i,j} y_{ij}^{(sd)} t^{(sd)} - \sum_{s,d} t^{(sd)}) - (\sum_{s,d,i,j} d_{ij}^{(sd)} t^{(sd)} - \sum_{s,d} t^{(sd)}) \tag{4.14}$$

Here $\{y_{ij}^{(sd)}\}$ are the traffic indicators from $\mathcal{G}$, while $\{d_{ij}^{(sd)}\}$ are the traffic indicators from $\mathcal{G}'$. $[t^{(sd)}]$ is the new traffic matrix. We calculate the grooming gain assuming that the old grooming solution $\mathcal{G}$ can still carry the new traffic. This is because if the traffic pattern changes to such an extent that $\mathcal{G}$ cannot carry the new traffic, then obtaining grooming gain by reconfiguring to $\mathcal{G}'$ is irrelevant as $\mathcal{G}$ has already become infeasible and traffic injected into the network is being lost in the network. We assume that such situations are strictly forbidden and lie outside the scope of the frame work we are addressing. Practically, this can be achieved by planned over-provisioning, which we discuss in the next chapter, and call admission control, which is outside our scope.

For any traffic candidate, the optimal grooming solution $\mathcal{G}$ can be obtained by solving the Integer Linear Program given in [19]. However, we add the following constraints to the ILP according to the assumptions made in our thesis in Section 3.3.

$$b_{ij} \,\epsilon\, 0, 1,\ \forall i, j \tag{4.15}$$

$$t^{sd} \,\epsilon\, 0, 1, \cdots C,\ \forall s, d \tag{4.16}$$

$$d_{ij}^{(sd)} \,\epsilon\, 0, 1,\ \forall s, d, i, j \tag{4.17}$$

$$t_{ij}^{(sd)} = d_{ij}^{(sd)}\ t^{(sd)},\ \forall s, d, i, j \tag{4.18}$$

The constraint 4.15 limits the total number of lightpaths between any two nodes to 1. Constraint 4.16 ensures that the total traffic for any s-d pair cannot be greater than the capacity C of a lightpath.

Constraints 4.18 and 4.17 enforce no bifurcation of traffic for a particular s-d pair i.e. the total traffic from s to d flows over the same sequence of lightpaths.

4.3 Problem Formulation as an Integer Linear Program

The objective of the integrated ILP is to maximize the grooming gain minus the reconfiguration cost as discussed in Section 3.2.2. We formulate the integrated problem as follows.

Maximize:

$$\psi(\{y_{ij}^{(sd)}\},\{d_{ij}^{(sd)}\},\{t^{(sd)}\})\gamma - \phi(\{x_{ij}^{k}(l,m)\},\{c_{ij}^{(k)}(l,m)\}) - \delta \tag{4.19}$$

Subject to:

Integer Constraints:

$$b_{ij}\ \epsilon\ 0,1,\ \forall i,j \tag{4.20}$$

$$t^{sd}\ \epsilon\ 0,1,\cdots C,\ \forall i,j \tag{4.21}$$

$$d_{ij}^{(sd)}\ \epsilon\ 0,1,\ \forall s,d,i,j \tag{4.22}$$

Physical Topology Constraints:

$$b_{ij}(l,m) \leq b_{ij}p_{lm},\ \forall i,j,l,m \tag{4.23}$$

$$c_{ij}^{(k)}(l,m) \leq p_{lm},\ \forall i,j,k,l,m \tag{4.24}$$

Lightpath Routing SP Constraints:

$$\sum_{i,j} b_{ij}(l,m) \leq W,\ \forall l,m \tag{4.25}$$

$$\sum_{l=0}^{N-1} b_{ij}(m,l) - \sum_{l=0}^{N-1} b_{ij}(l,m) = \left\{ \begin{array}{ll} b_{ij}, & m=i \\ -b_{ij}, & m=j \\ 0, & m\neq i, m\neq j \end{array} \right\} \forall m,i,j \tag{4.26}$$

Lightpath Wavelength Assignment SP Constraints:

$$\sum_{k=0}^{W-1} c_{ij}^{(k)}(l,m) = b_{ij}(l,m),\ \forall i,j,l,m \tag{4.27}$$

$$\sum_{i,j} c_{ij}^{(k)}(l,m) \leq 1,\ \forall k,l,m \tag{4.28}$$

$$\sum_{l=0}^{N-1} c_{ij}^{k}(m,l) - \sum_{l=0}^{N-1} c_{ij}^{(k)}(l,m)) \left\{ \begin{array}{ll} \leq b_{ij}, & m=i \\ \geq -b_{ij}, & m=j \\ = 0, & m\neq i, m\neq j \end{array} \right\} \forall i,j,k,m \tag{4.29}$$

Traffic Routing SP Constraints:

$$t_{ij} = \sum_{s,d} t_{ij}^{(sd)},\ \forall i,j \tag{4.30}$$

$$t_{ij} \leq b_{ij} C,\ \forall i,j \tag{4.31}$$

$$t_{ij}^{(sd)} = d_{ij}^{(sd)} t^{(sd)},\ \forall s,d,i,j \tag{4.32}$$

$$\sum_{j=0}^{N-1} t_{ij}^{(sd)} - \sum_{j=0}^{N-1} t_{ji}^{(sd)} = \left\{ \begin{array}{ll} t^{(sd)}, & i = s \\ -t^{(sd)}, & i = d \\ 0, & i \neq s, i \neq d \end{array} \right\} \forall s,d,i \tag{4.33}$$

In the ILP objective 4.19, the function $\psi(\{y_{ij}^{(sd)}\}, \{d_{ij}^{(sd)}\}, \{t^{(sd)}\})$ is the grooming gain discussed in Section 4.2 and $\phi(\{x_{ij}^{k}(l,m)\}, \{c_{ij}^{(k)}(l,m)\})$ is $RC - IV$ as calculated in Section 4.1.3. The factors γ and δ are constants, as explained in Section 3.2.2.

Most of the constraints are self explanatory and have been discussed in [19]. The integer constraints 4.20, 4.21 and 4.22 have been explained in the previous section. The physical topology constraints 4.23 and 4.24, constraint the establishment of the lightpaths to the given physical topology. The lightpath routing constraint 4.25 limits the total number of lightpaths on a physical link to the number of wavelengths W supported by a physical link. Constraint 4.26 asserts the conservation of every lightpath at every physical link endpoint. Constraint 4.27 ensures that a lightpath, if it exits in the virtual topology, has a unique wavelength. Constraint 4.28 enforces that a wavelength can be used at most once in every physical link, avoiding wavelength clash. Constraint 4.29 asserts the conservation of every wavelength at every physical link endpoint for each lightpath. Constraint 4.30 asserts that the total traffic on a lightpath is a sum of all sub-wavelength elements on that lightpath due to different source-destination pairs. Constraint 4.31 asserts that if a lightpath exists, the total traffic on the lightpath cannot exceed the capacity C of the lightpath. We will refer to this constraint as the *capacity constraint* throughout this thesis. Constraint 4.32 is the no bifurcation constraint explained in the previous section and constraint 4.33 is an expression of the conservation of traffic flow at lightpath endpoints.

Chapter 5

Minimal Disturbance Reconfiguration By Over-Provisioning

While the ILP presented in the last Chapter provides a method of solving the integrated grooming and reconfiguration problem we have articulated, it will take significant computation time to actually obtain a solution by employing an ILP solver. Realistically, sometimes this computational effort will have to be invested, but it would be desirable to have to do so only at long intervals, such that the average computational cost incurred is small. Further, some swifter means to react to traffic changes by reconfiguration is desirable than would be the case if the ILP were to be solved. In this chapter, we present a heuristic algorithm designed to serve as an overall reconfiguration control algorithm with these concerns in mind.

A practical strategy to address those concerns must have the following goals. The heuristic approach must (i) avoid resorting to the full ILP whenever possible, (ii) ward off failure of the network, that is a situation in which any increase in the traffic components would cause the current grooming solution $\mathcal{G}$ to fail to carry the new traffic, (iii) avoid adopting grooming solutions that are very much suboptimal, in decreasing order of importance. The rationale is as follows: the first concern must be to minimize ILP use, otherwise this heuristic may become equivalent to simply running the ILP each time. Secondly, we know that if any feasible solutions to the changed traffic conditions exist, the ILP approach would find the optimal such feasible solution. However, feasibility

is more attainable than optimality, so by avoiding the ILP, we must lower our sights to at least remain feasible. Since the heuristic will run repeatedly on the same network (with successively changing traffic patterns), the feasibility concern is represented by the next goal, which makes the heuristic somewhat proactive in seeking continued feasibility. Finally, we still want to address the grooming cost concern, even though it has to be subjugated to the first two needs. Briefly, the second and third goals describe the reconfiguration and the grooming goals respectively.

Since the problem is intractable in nature, a tractable heuristic is unlikely to attain globally optimal solutions. For this reason, we attempt to apply heuristic methods when it appears that the new traffic matrix represents a local change in traffic, in some sense. When the reach of the change in traffic is close to global, the heuristic simply indicates that the ILP solution approach should be employed.

As we remarked in Section 4.2, warding off failure can be done by planned over-provisioning. In the next section, we present this important idea of over-provisioning at the traffic demand level, and present the heuristic approach in the following sections.

5.1 Over-Provisioning of Traffic Matrices

By over-provisioning at the traffic demand level, we mean that given a groomed solution $\mathcal{G}$, we apply a modified version of this solution, $\mathcal{G}_{op}$ to the actual network. In $\mathcal{G}_{op}$, some or all traffic elements are increased by some amount such that the network is made to carry as much traffic as possible on all the lightpaths without violating the capacity constraint. This method increases the probability of a new traffic matrix being supported on the existing virtual topology, as we demonstrate below in Section 5.1.1.

We then propose various methods of calculating the total amount of over-provisioned traffic that can be assigned to each *s-d* pair in this section. Once the amount of over-provisioning that needs to be done is determined from an optimal groomed solution $\mathcal{G}$, the OXCs and the DXCs are configured to carry this over-provisioned traffic instead of the original traffic. However, it has been assumed that the control information about the actual amount of traffic and other management related functions are handled by a separate control unit. The control unit is responsible for directly communicating with the OXCs and the DXCs and issuing configuration commands. The study of the control unit and the flow of control messages is beyond the scope of our work.

The proposed method of over-provisioning will not result in a bad grooming solution i.e. an increase in the grooming cost. This is because we start with an optimal groomed solution $\mathcal{G}$ and we calculate the grooming cost from the original traffic. We do not calculate the grooming cost from the over-provisioned traffic because it is the OEO conversion that the original traffic undergoes determines that total delay. It does not matter if the over-provisioned traffic is delayed or is transmitted erroneously. By over-provisioning we only try to utilize the extra capacity in the

network, which would otherwise be wasted. However, over-provisioning can only be done if there is some extra capacity on the established lightpaths. Thus, if $b_{ij} = 1$ and $C - \sum_{s,d} t_{ij}^{(sd)} > 0$, then there is some capacity on the lightpath ij which is not used by the current traffic. We can over-provision such a lightpath only to a limit where $\sum_{s,d} t_{ij}^{(sd)}$ becomes equal to C, otherwise the capacity constraint would be violated.

5.1.1 Family of Traffic Matrices Supported

To show that over-provisioning helps in continued feasibility, we determine the set of all traffic matrices that can be carried on a given virtual topology $\mathcal{T}$ using a grooming solution $\mathcal{G}$ *without* requiring any reconfiguration of the nodes (OXCs/DXCs).

Given a groomed solution $\mathcal{G}$ for a traffic matrix $T_{old} = [t_{old}^{(sd)}]$, any new traffic matrix $T_{new} = [t_{new}^{(sd)}]$ whose elements can be mapped on the existing lightpaths given by $\mathcal{G}$ without violating the capacity constraint (see Section 4.3), will be supported. Thus, all the traffic matrices which are a subset of the current traffic matrix can be carried on the existing virtual topology.

$$T_{\triangle} = T_{new} - T_{old} \tag{5.1}$$

The new traffic can be carried on the existing virtual topology if all the elements of $T_{\triangle}$ are $<= 0$, although the mapping of the traffic on the lightpaths might not be optimal with respect to the grooming cost. Any new traffic demand for a particular *s-d* pair which does not exceed the total traffic for that *s-d* pair can be routed on the existing topology. We determine this based on the architecture and functioning of the OXC and the DXC. When an OXC/DXC has been configured to switch a certain amount of traffic on its input/output links, it will continue to switch that same amount of traffic until its reconfigured. Thus, if the traffic demand on any link reduces from its original assignment and the OXC/DXC has not been reconfigured then the OXC/DXC will continue to switch the original amount of traffic, the extra traffic could be all 0 bits or don't cares. On the other hand, if the traffic demand on any link increases more than the initial assignment, then the OXC/DXC need to be reconfigured to be able to switch the total new traffic. Based on this fact, we determine that a grooming solution $\mathcal{G}$ that can carry C amount of traffic on all its lightpaths can support more number of traffic matrices as compared to a grooming solution that carries less than C amount of traffic on its lightpaths. In general, of two grooming solutions which differ only in some spare capacity on some lightpath that is (1) left unassigned by one solution $\mathcal{G}_1$, and (2) assigned as over-provisioning to one or more traffic elements traversing that lightpath by another solution $\mathcal{G}_2$, the latter solution supports a strictly superset of possible traffic changes than the former one. This demonstrates that an over-provisioned grooming solution has a higher chance of supporting an arbitrary change in the traffic matrix.

5.1.2 Initial Over-Provisioned Amount

In order to decide by how much to over-provision each sub-wavelength element on a particular lightpath, we determine the sub-wavelength elements on all the lightpaths using the variable values given by $\mathcal{G}$. We represent the total number of sub-wavelength elements on an existing lightpath from i to j (i.e. for each $b_{ij} = 1$) by s_{ij}. We increment the value of s_{ij} by 1 for each $t_{ij}^{(sd)} \neq 0$ for the lightpath from i to j. To be able to over-provision each traffic element, we not only need to determine the total number of sub-wavelength elements on a lightpath but also to uniquely identify them on each lightpath. Since as per the assumption made in this thesis there can only be a single lightpath from i to j, every non-zero $t_{ij}^{(sd)}$ can be mapped only to one lightpath, uniquely identifying the sub-wavelength elements. We propose three different methods of calculating the initial over-provisioned amount i.e. the maximum possible amount which can be allocated to each sub-wavelength element on a per lightpath basis. As discussed earlier, this over-provisioned amount should not violate the capacity constraint on any lightpath. After determining the initial maximum possible over-provisioned amount, we calculate the actual amount by which each traffic element will be over-provisioned in Section 5.1.3. This actual amount of over-provisioning might be different from the maximum value calculated below due to the capacity constraint.

We propose the following three methods for calculating the initial over-provisioned amount for each sub-wavelength element on each lightpath.

1. *Equal allocation:* In this method we allocate the following equal amount of over-provisioned traffic to each $t_{ij}^{(sd)}$ on a lightpath from i to j.

$$t\_over_{ij}^{(sd)} = \left\{ \begin{array}{ll} (C - \sum_{s,d} t_{ij}^{(sd)})/s_{ij}, & t^{(sd)} \neq 0 \\ 0, & \text{otherwise} \end{array} \right\} \forall s, d, i, j \tag{5.2}$$

   This is a fair allocation method, as each sub-wavelength element is over-provisioned by the same amount, irrespective of the actual amount of traffic carried by them.

2. *Prorated allocation:* In this method we assign the maximum amount of over-provisioned traffic to a sub-wavelength element that carries the maximum amount of original traffic $t^{(sd)}$. This method is based on the assumption that the current magnitude of a sub-wavelength element is representative of its possible future growth; thus a larger traffic element has greater potential/possibility of having increased traffic demand in the future. Thus, in this method we the allow the maximum sub-wavelength element to be over-provisioned the most. The over-provisioned value to be assigned to each sub-wavelength element on a particular lightpath is given by the following equation.

$$t\_over_{ij}^{(sd)} = \left\{ \begin{array}{ll} (t_{ij}^{(sd)} / \sum_{s,d} t_{ij}^{(sd)}) \, (C - \sum_{s,d} t_{ij}^{(sd)}), & t^{(sd)} \neq 0 \\ 0, & \text{otherwise} \end{array} \right\} \forall s, d, i, j \tag{5.3}$$

3. *Inverse allocation:* In this method we assign the minimum amount of over-provisioned traffic to a sub-wavelength element that carries the maximum amount of original traffic $t^{(sd)}$. This method is based on the assumption that all sub-wavelength elements have similar averages over time, so one that is already carrying the maximum traffic will probably have the least increase in traffic demand in future. In the following equation we make sure that the smallest sub-wavelength element (in terms of amount of traffic carried by it) on that particular lightpath gets the maximum amount of over-provisioned traffic. The total over-provisioned traffic is given by the following equations.

$$t\_over_{ij}^{(sd)} = \left\{ \begin{array}{ll} ((\sum_{s,d} t_{ij}^{(sd)} - t^{(sd)}) / \sum_{s,d} t_{ij}^{(sd)}(s_{ij} - 1))\,(C - \sum_{s,d} t_{ij}^{(sd)}), & t^{(sd)} \neq 0 \\ 0, & \text{otherwise} \end{array} \right\} \forall s, d, i, j \tag{5.4}$$

Any of the above three methods can be used to determine the initial over-provisioned amount. The administrator's knowledge of network and traffic characteristics would determine the appropriate one.

5.1.3 Different Methods of Over-Provisioning

Traffic from a particular source to destination may flow over many lightpaths. The amount of over-provisioned traffic $t\_over_{ij}^{(sd)}$ as calculated above may be different on different lightpaths for the same traffic element $t_{ij}^{(sd)}$. The maximum amount of over-provisioned traffic that we can actually assign to a particular *s-d* pair is the minimum of all the over-provisioned values calculated on all the lightpaths where the traffic for this *s-d* pair flows. We choose the minimum value to guarantee that the capacity constraint is not violated on any of the lightpaths. We propose the following methods for calculating the final over-provisioned traffic by which each *s-d* pair traffic will be incremented.

Equal Over-provisioning Method: In this method, we select the minimum of all $t\_over_{ij}^{(sd)}$ values calculated for all the *s-d* pairs, using any of the three methods described in the previous section. We denote this minimum value by t_{min} and all the *s-d* pairs carrying traffic are over-provisioned by this value. Thus, in this method each sub-wavelength element is over-provisioned by an equal amount. We denote the final traffic matrix after over-provisioning by $T_{prov} = [t_{prov}^{(sd)}]$ and calculate this matrix according to the following equations.

$$t_{min} = Minimum(t\_over_{ij}^{(sd)}),\ \forall t^{(sd)} \neq 0, s, d, i, j \tag{5.5}$$

$$t_{prov}^{(sd)} = \left\{ \begin{array}{ll} t^{(sd)} + t_{min}, & t^{(sd)} \neq 0 \\ 0, & \text{otherwise} \end{array} \right\} \forall s, d \tag{5.6}$$

Selective Over-provisioning Method: In this method, instead of selecting a minimum t_{min} value and over-provisioning each *s-d* pair carrying traffic by this value we determine a different

over-provisioned amount for each s-d pair. In this method we over-provision each s-d pair by the minimum of all the $t\_over_{ij}^{(sd)}$ values on all the lightpaths that the traffic from s to d flows. Thus, we over-provision all the s-d pairs by a different value, denoted by $t_{min}^{(sd)}$. We use the following equations in Selective method to calculate the final traffic matrix after over-provisioning.

$$t_{min}^{(sd)} = \left. Minimum(t\_over_{ij}^{(sd)}),\ \forall t^{(sd)} \neq 0, i, j \right\} \forall s, d \tag{5.7}$$

$$t_{prov}^{(sd)} = \left\{ \begin{array}{ll} t^{(sd)} + t_{min}^{(sd)}, & t^{(sd)} \neq 0 \\ 0, & \text{otherwise} \end{array} \right\} \forall s, d \tag{5.8}$$

Iterative Over-provisioning Method: In this method, we try to over-provision each s-d pair by a maximum amount, though all the lightpaths may still not be over-provisioned upto the maximum capacity C. Using any method it may not be possible to over-provision all the sub-wavelength elements to such an extent that all lightpaths carry total traffic equal to C. This is due to certain sub-wavelength elements flowing over bottle neck lightpaths, these are the lightpaths carrying the maximum total traffic. In Iterative method, we calculate the $t_{min}^{(sd)}$ values as done in the Selective method. Then we over-provision one of the s-d pairs by this amount given by the following equation.

$$t_{prov}^{(sd)} = t^{(sd)} + t_{min}^{(sd)} \tag{5.9}$$

After fixing the over-provisioned amount for a particular s-d pair, we recalculate $t\_over_{ij}^{(sd)}$ for all the other s-d pairs. The only difference this time in calculating $t\_over_{ij}^{(sd)}$ is that, the s-d pair which we have already over-provisioned will not be considered as a sub-wavelength element on any of the lightpaths that it flows. And the total capacity of all the lightpaths over which $t^{(sd)}$ flows will now be considered as $C - t_{prov}^{(sd)}$ instead of C. After recalculating $t\_over_{ij}^{(sd)}$ for rest of the s-d pairs, we recalculate the $t_{min}^{(sd)}$ values also. We keep iterating through these steps until all the non-zero $t^{(sd)}$ values have been over-provisioned.

We have also determined the following five variations to the Iterative method based on the order in which each s-d pair is selected and over-provisioned.

- *Iterative-Min:* The s-d pair that carries the minimum traffic is over-provisioned before the others. Thus, at each iteration we select the minimum $t^{(sd)}$ which has not yet been over-provisioned.

- *Iterative-Max:* The s-d pair that carries the maximum traffic is over-provisioned before the others. Thus, at each iteration we select the maximum $t^{(sd)}$ which has not yet been over-provisioned.

- *Iterative-Ratio:* The s-d pair that has the minimum ratio given by $t_{min}^{(sd)}/t^{(sd)}$, is over-provisioned before the others. Thus, at each iteration we select that s-d pair which has the minimum value of $t_{min}^{(sd)}/t^{(sd)}$ and has not yet been over-provisioned.

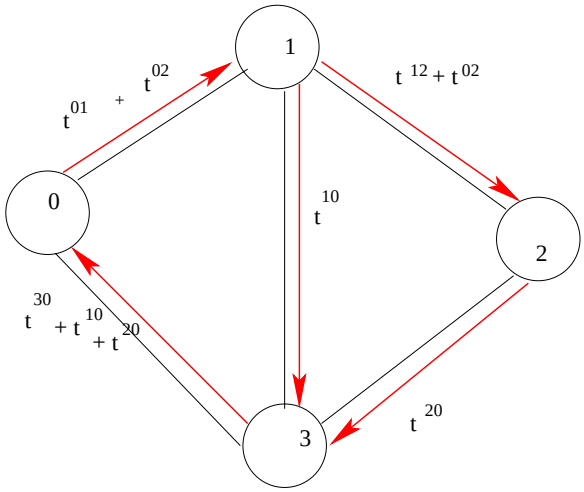

Figure 5.1: A Virtual Topology Representation Carrying Sub-wavelength Traffic

- *Iterative-Max-lightpath:* The *s-d* pair that has traffic flowing over the maximum number of lightpaths, is over-provisioned before the others. Thus, at each iteration we select the $t^{(sd)}$ that flows over the maximum number of lightpaths and has not yet been over-provisioned. We formulated this method based on the observation that the more number of lightpaths a particular traffic elements flows on, the higher is the probability that one of those lightpaths would be a bottle neck lightpath. By a bottle neck lightpath we mean, a lightpath which carries maximum traffic and has very little extra capacity that can be used for over-provisioning as compared to the other lightpaths.

- *Iterative-Min-Max:* The *s-d* pair that has the maximum $t_{min}^{(sd)}$ value is over-provisioned before the others. Thus, at each iteration we select the *s-d* pair that has the maximum $t_{min}^{(sd)}$ value and has not yet been over-provisioned.

We use the virtual topology given in Figure 5.1 to explain the Equal, Selective and Iterative methods. Let, $C = 15$ and the traffic matrix T to be over-provisioned is given below. This instance has been constructed with only single-hop lightpaths, and only one loop-free routing available for any traffic element; thus the $t_{ij}^{(sd)}$ are obvious from the $t^{(sd)}$.

$$T = \begin{pmatrix} 0 & 3 & 4 & 0 \\ 7 & 0 & 5 & 0 \\ 2 & 0 & 0 & 0 \\ 3 & 0 & 0 & 0 \end{pmatrix}$$

We use the Equal Allocation method defined in Section 5.1.2 to calculate $t\_over_{ij}^{(sd)}$ values. These values are $t\_over_{01}^{(01)} = 4, t\_over_{01}^{(02)} = 4, t\_over_{12}^{(12)} = 3, t\_over_{12}^{(02)} = 3, t\_over_{13}^{(10)} =$

$8, t\_over_{23}^{(20)} = 13, t\_over_{30}^{(10)} = 1, t\_over_{30}^{(20)} = 1, t\_over_{30}^{(30)} = 1$.

For the Equal method we select the minimum $t\_over_{ij}^{(sd)}$ value which is 1 for this example, thus $t_{min} = 1$. Following is the over-provisioned traffic matrix calculated using the Equal method.

$$T_{prov} = \begin{pmatrix} 0 & 4 & 5 & 0 \\ 8 & 0 & 6 & 0 \\ 3 & 0 & 0 & 0 \\ 4 & 0 & 0 & 0 \end{pmatrix}$$

In order to calculate the over-provisioned matrix according to the Selective method, we calculate the following $t_{min}^{(sd)}$ values, $t_{min}^{(01)} = 4$, $t_{min}^{(02)} = 3$, $t_{min}^{(12)} = 3$, $t_{min}^{(10)} = 1$, $t_{min}^{(20)} = 1$, $t_{min}^{(30)} = 1$. Following is the over-provisioned traffic matrix calculated using Selective method.

$$T_{prov} = \begin{pmatrix} 0 & 7 & 7 & 0 \\ 8 & 0 & 8 & 0 \\ 3 & 0 & 0 & 0 \\ 4 & 0 & 0 & 0 \end{pmatrix}$$

To calculate the over-provisioned matrix by the Iterative-Max method, we first calculate all the $t_{min}^{(sd)}$ values as calculated in Selective method. The maximum traffic element is $t^{(10)}$, thus we over-provision it by $t_{min}^{(10)}$ and we recalculate all the $t\_over_{ij}^{(sd)}$ and $t_{min}^{(sd)}$ values. Similarly fixing $t^{(12)}$, then $t^{(02)}$, then $t^{(01)}$ and then $t^{(30)}$, we get the following over-provisioned traffic matrix.

$$T_{prov} = \begin{pmatrix} 0 & 8 & 7 & 0 \\ 8 & 0 & 8 & 0 \\ 3 & 0 & 0 & 0 \\ 4 & 0 & 0 & 0 \end{pmatrix}$$

For this example the total amount of traffic after over-provisioning is maximum in the Iterative-Max method, then in the Selective method, followed by the Equal method.

5.1.4 Comparative Analysis of Over-Provisioning Methods

We make a comparative study of all the over-provisioning methods discussed above. We also mathematically formulate the goal of the over-provisioning algorithms. Finally in this section we briefly discuss the method of over-provisioning for those groomed solutions where there are more than one lightpaths between a pair of nodes.

Using logical deductions we prove that the Iterative method of over-provisioning is better than both the Equal method and the Selective method in terms of the total amount of over-provisioned traffic calculated. Further, the Selective method is better than the Equal method in terms of the total amount of over-provisioned traffic calculated.

Proposition: The Selective method performs at least as well as the Equal method in terms of total amount of over-provisioned traffic.

Proof: In the Equal method we select the minimum of all the $t\_over_{ij}^{(sd)}$ values. This minimum value is represented by t_{min} and each $t^{(sd)}$ element is over-provisioned by this amount. On the other hand, in the Selective method we select the minimum of only the $t\_over_{ij}^{(sd)}$ values for a particular *s-d* pair, represented by $t_{min}^{(sd)}$. Each $t^{(sd)}$ is over-provisioned by the respective $t_{min}^{(sd)}$ value. The t_{min} selected in the Equal method is equal to at least one of the $t_{min}^{(sd)}$ values selected by Selective method. Thus the *s-d* pair for which $t_{min}^{(sd)} = t_{min}$, would be over-provisioned by the same amount in both the methods. However, the other $t_{min}^{(sd)}$ values which are not equal to t_{min}, would always be greater than t_{min} because t_{min} was the minimum of all $t\_over_{ij}^{(sd)}$ values. Thus, using the Selective method we can never over-provision by a value which is less than that done in the Equal method.

Proposition: The Iterative method performs at least as well as the Selective method in terms of total amount of over-provisioning.

Proof: In both the Selective and the Iterative methods we initially calculate $t\_over_{ij}^{(sd)}$ and $t_{min}^{(sd)}$ for all non-zero $t^{(sd)}$ values. This $t_{min}^{(sd)}$ is the final over-provisioned value for each non-zero $t^{(sd)}$ according to the Selective method. However, in the Iterative method after this initial calculation we fix a $t^{(sd)}$ and over-provision it by the respective $t_{min}^{(sd)}$ value. Until this step the Selective and the Iterative methods are the same. At every step after this, in the Iterative method we fix a particular $t^{(sd)}$ value and recalculate the total extra capacity on all the lightpaths where $t^{(sd)}$ flows. We then redistribute this total extra capacity among rest of the sub-wavelength elements which have not yet been over-provisioned. At every iterative step in Iterative method we try to reassign any extra capacity over the initial assignment. In the worst case, if there is no extra capacity on the lightpath after fixing one of the sub-wavelength elements, we will assign all $t^{(sd)}$ elements the $t_{min}^{(sd)}$ values initially calculated. In this worst case scenario the over-provisioned values will be same for both methods. Thus, Iterative method can never assign an over-provisioned value which is less than that assigned by Selective method.

Experimental Results for Comparative Study: We performed experiments for a mesh physical topology of 9 nodes with $C = 32$ and $W = 4$ to calculate the total amount of over-provisioned traffic. The traffic matrices in the experiment were generated randomly having an integer uniform distribution bounded between 0 and 8 and a groomed solution used to calculate the over-provisioned traffic was obtained by solving the ILP in [19], which we describe in Section 4.2. The ILP was coded using CPLEX, which is a software designed to solve linear optimization problems.

Figure 5.2 shows the total amount of traffic for all the *s-d* pairs after over-provisioning. The experiments were run for 55 random instances of the input traffic and at each instance the over-provisioned traffic matrix was calculated using all the over-provisioning methods described in the previous section. For the experiments, the initial over-provisioned amount was calculated using equal allocation method (Section 5.1.2). The results shown in Figure 5.2 prove that the Iterative

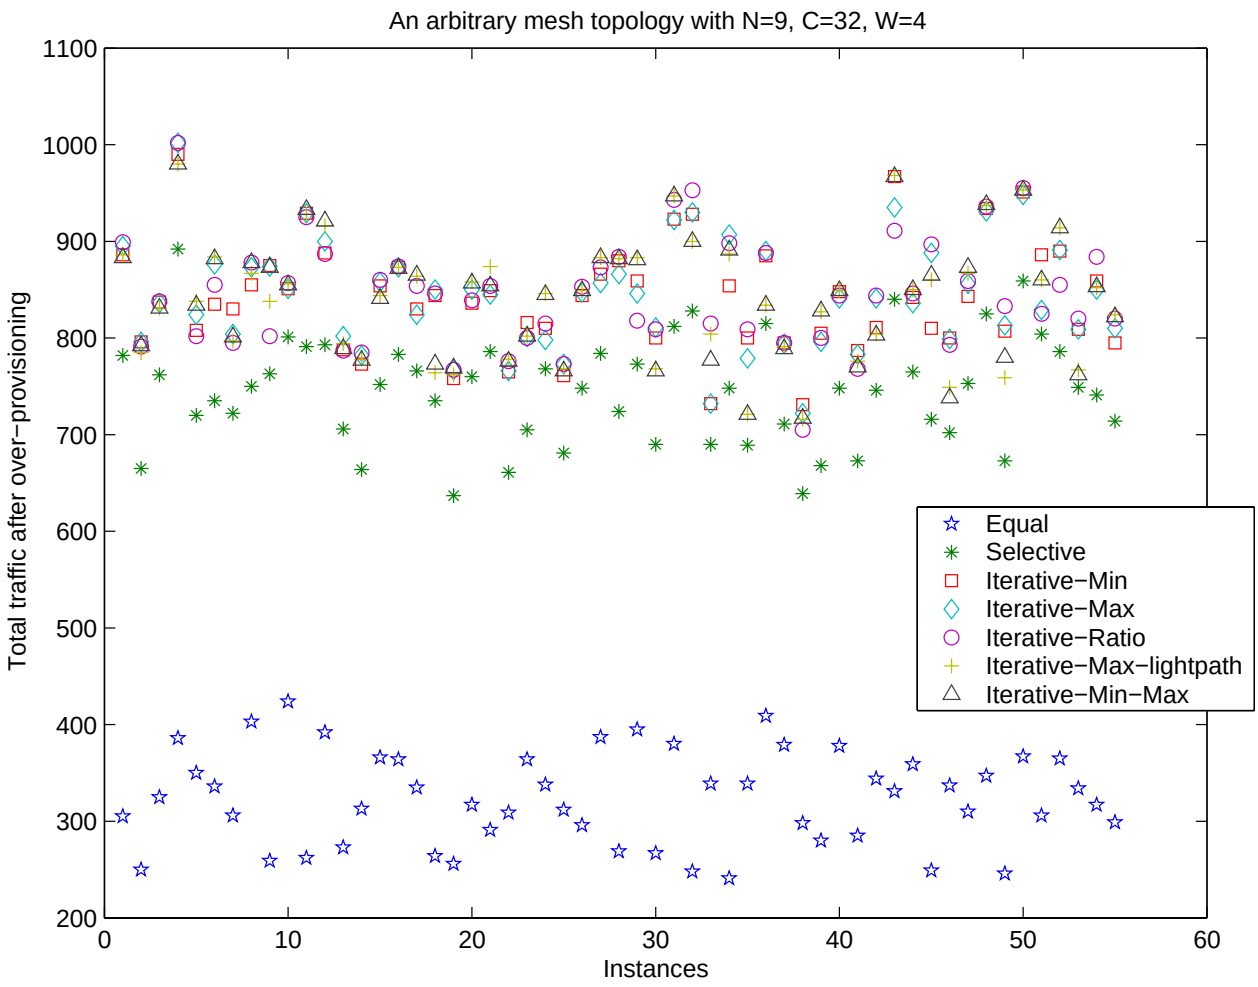

Figure 5.2: Total Over-Provisioned Traffic for N=9, C=32, W=4

method is better than the other two methods in terms of the total amount of traffic over-provisioned. However, as observed in the Figure 5.2 the different Iterative methods over-provision by almost the same amount, with none of them out-performing the rest over all the instances.

Figure 5.3 shows the variance of the ratio of the over-provisioned traffic to the original traffic. As seen in the figure, the variance is very high for certain instances and this is because of a few *s-d* pairs which are over-provisioned by very large amounts.

Formulating the goal of the over-provisioning algorithm: The above discussion helps us formulate an appropriate and precise goal of the over-provisioning process. As discussed in Section 5.1.1 over-provisioning is done to be able to support more number of traffic matrices, without performing reconfiguration. Thus, over-provisioning algorithm aims at allocating the max-

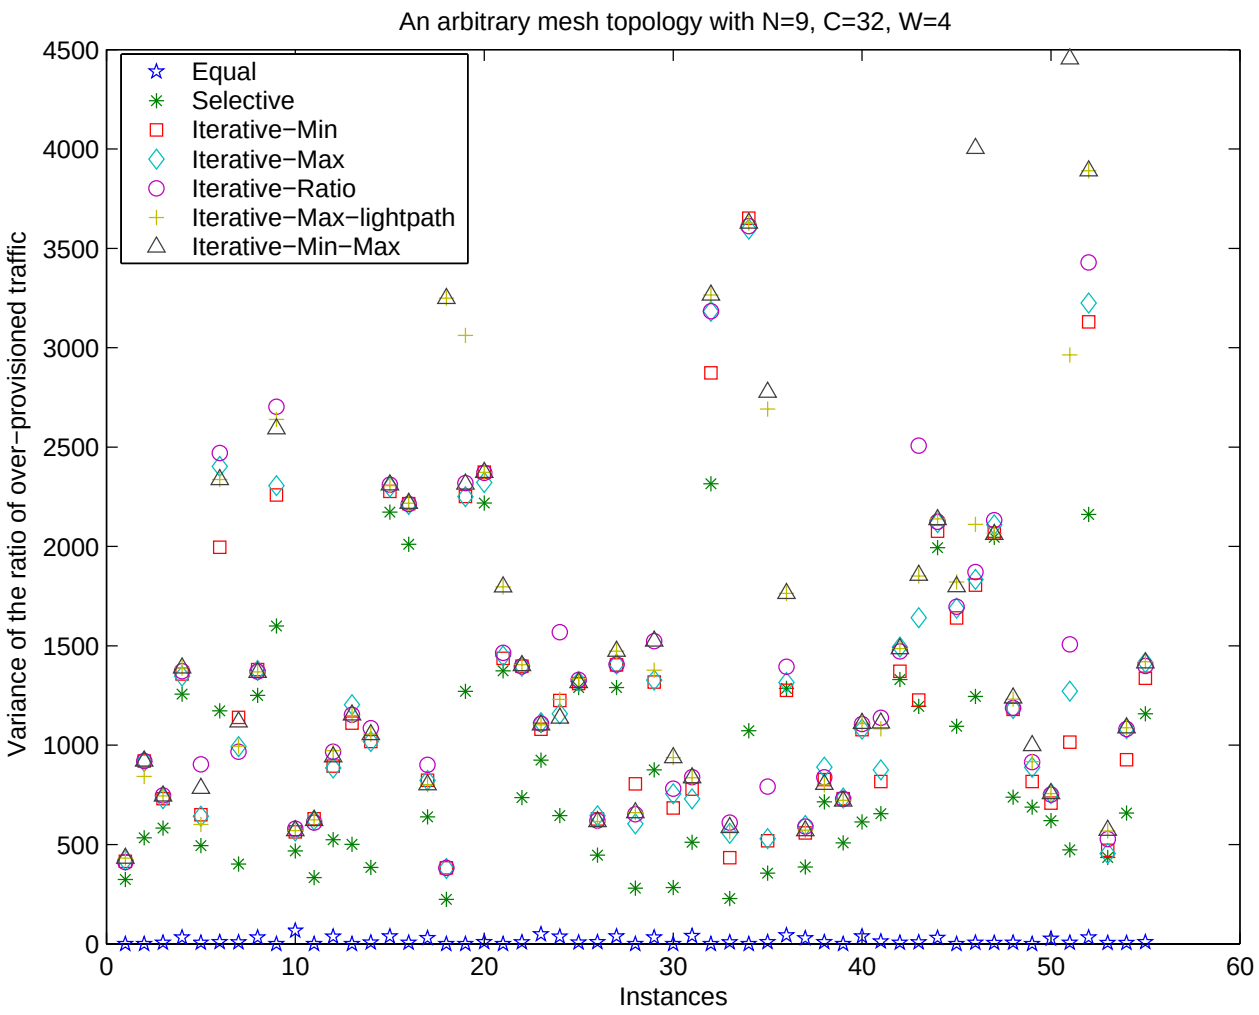

Figure 5.3: Variance of the Ratio of Over-Provisioned Traffic for N=9, C=32, W=4

imum possible extra capacity to each of the sub-wavelength elements on all the existent lightpaths. Mathematically the goal of the over-provisioning algorithm can be formulated as follows.

Objective: Maximize:

$$\sum_{s,d} t_{prov}^{(sd)} - \sum_{s,d} t^{(sd)} \tag{5.14}$$

Subject to:

$$\sum_{s,d} t_{prov}^{(sd)} \leq \; b_{ij}C, \; \forall i,j \tag{5.15}$$

The objective 5.14 maximizes the difference between the total traffic calculated after over-provisioning and the initial traffic without any over-provisioning. Constraint 5.15 assures that the

total traffic after over-provisioning never exceeds the capacity of any lightpath. However, this goal does not allow any inclusion of max-min fairness of the over-provision amounts.

5.1.5 The Case of Multiple Lightpaths

In this thesis, we have considered only those grooming solutions which allow a single lightpath between any two nodes (i.e. $0 \leq b_{ij} \leq 1$). However, here we briefly discuss some of the issues involved with a more general grooming solution $\mathcal{G}$ where b_{ij} is allowed to be greater than 1. As discussed in Section 5.1.2, the over-provisioned traffic matrix can only be obtained after uniquely identifying the sub-wavelength elements on each lightpath. The problem with $\mathcal{G}$ where $b_{ij} > 1$ is that the traffic can be mapped onto the lightpaths in more than one ways. We identify the following possibilities of solving this issue and obtaining a unique mapping of the traffic onto the lightpaths so that a over-provisioned traffic matrix can then be calculated. This is an important issue because grooming formulations in the literature do not provide this mapping information in the grooming solutions.

1. From the many possible mappings of the traffic on the existing lightpaths, we select that solution which maximizes the minimum slack on all the lightpaths. By slack we mean the amount of capacity not used by the total current traffic flowing on a particular lightpath. We try to maximize the minimum slack because a lightpath with the least slack usually becomes the bottle neck link and forces less amount of over-provisioning. For example: Assume a groomed solution where $b_{12} = 3$ and there is a feasible solution that has a slack of $(0, 4, 4)$ on each of these three lightpaths, while there is another solution that has slack of $(2, 2, 2)$. If these are the only two possible solutions, then the second one which gives a slack of $(2, 2, 2)$, is the one we select because it maximizes the minimum slack. However, we prove below that the problem of maximizing the minimum slack is NP complete.

2. Another possibility is that we can randomly select any one of the feasible solutions, probably the first feasible solution that is calculated.

3. We propose a third possibility such that the groomed solution itself is modified to calculate the value of a variable $t_{ij}^{(k)(sd)}$. The variable $t_{ij}^{(k)(sd)}$ represents the total traffic from source s to destination d on the lightpath i-j which is identified by k. The lightpath identifier, k uniquely identifies every lightpath from i to j, $k \epsilon 1, 2 ..... b_{ij}$. However, since the existing groomed solution does not include the variable $t_{ij}^{(k)(sd)}$, there is no way of distinguishing between more than one lightpaths from i to j.

Proposition: Maximizing the minimum slack is an NP complete problem

Proof: The following proof considers just two nodes in the network, that are connected by N number of lightpaths. As discussed above, slack is the difference between the capacity C

of a lightpath and the sum of all the sub-wavelength elements on that lightpath. The problem of maximizing the minimum slack (max-min problem) is defined below and we prove the max-min problem to be NP complete by reducing an instance of the bin-packing problem to it.

GIVEN: N, C, a finite set S of items, a size $size(x)\epsilon Z^+$ and $size(x) \leq C$, for each $x\epsilon S$

FIND: A Partition $R = S_1, S_2 \cdots S_m$, $S_i \subseteq S$, $m \leq N$, which satisfies

$$\sum_{x_j \epsilon S_i} x_j \leq C \tag{5.16}$$

i.e. a partition where the sum of all the elements is less than or equal to C and the total number of partitions is less than or equal to the number of lightpaths available.

Max-Min Problem: Given an instance of $P < N, C, S >$ and a goal l, find a solution such that for each $i \leq m$,

$$\sum_{x_j \epsilon S_i} x_j \leq C - l \tag{5.17}$$

BIN-PACKING: Finite set U of items, a size $S(u)\epsilon Z^+$, for each $u\epsilon U$, a positive integer bin capacity B, and a positive integer K.

DECISION: Is there a partition of U into disjoint sets $U_1, U_2 \cdots U_K$ such that the sum of sizes of the items in each U_i, is B or less?

Reduction of the Bin-packing problem, which is a proved NP complete problem, to max-min problem: U from bin-packing is mapped to S of max-min problem, similarly K is mapped to N, $S(u)$ mapped to $size(x)$ where $u\epsilon U$ and $x\epsilon S$, B mapped to C. When $l = 0$, the bin-packing problem is reduced to the max-min problem. Thus, the problem of maximizing minimum slack is also an NP complete problem.

5.2 Reconfiguration Heuristic

With change in the traffic conditions, reconfiguration is needed either to reoptimize the virtual topology or to support the new traffic pattern that can no longer be carried on the existing virtual topology. As discussed in Section 2.4, in this thesis we have considered the need for reconfiguration to reoptimize the virtual topology. The undesirable effect of reconfiguration is the disturbance caused to the network. The process of adapting to the new optimal virtual topology under the changed traffic conditions involves the reconfiguration of the network components (OXC/DXC) that introduces delay, variability of delay, loss of throughput and disrupts the traffic. Although reconfiguration is needed to optimally carry the changed traffic, it is also important to keep the amount of reconfiguration as well as the frequency of reconfiguration at its minimum. The method of over-provisioning, discussed in the previous section, is a way of delaying the need for extensive reconfiguration as far as possible.

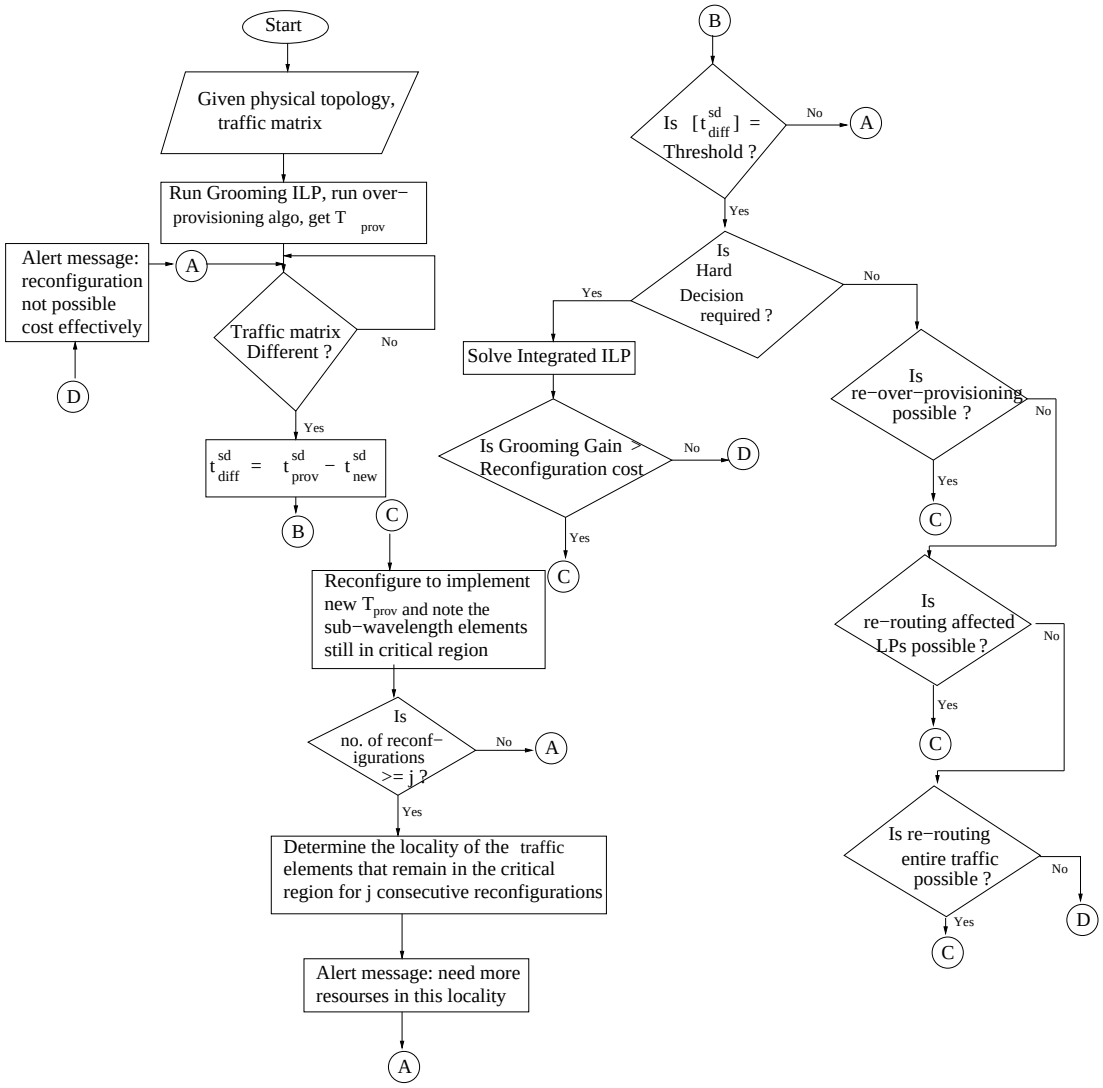

Figure 5.4: Flowchart

In this section we propose a reconfiguration heuristic, which is designed with the three goals mentioned at the beginning of this Chapter in mind. We show the entire flow of the reconfiguration heuristic in Figure 5.4. In the remaining sections of this chapter, we describe in detail the various aspects of this control flow.

The reconfiguration heuristic is triggered for every change in the traffic matrix. Given a physical topology, an initial traffic matrix T and a new traffic matrix T_{new}, we obtain an optimal groomed solution $\mathcal{G}$ by solving the ILP discussed in Section 4.2. We then calculate the over-provisioned traffic matrix T_{prov} (refer to Section 5.1.3) from $\mathcal{G}$ and configure the network components to carry the traffic given by T_{prov}. In order to determine if the new traffic would require the virtual

topology to be reconfigured or not, we calculate the difference matrix T_{diff} given by the following equation.

$$t_{diff}^{(sd)} = t_{prov}^{(sd)} - t_{new}^{(sd)} \; \forall s, d \tag{5.18}$$

From the discussion about the family of traffic matrices supported in Section 5.1.1, we know that any new traffic matrix which is a subset of the over-provisioned traffic matrix can be carried on the current virtual topology but the grooming objective can become suboptimal if no reconfiguration is performed. Thus, for all the T_{new} matrices which result in a T_{diff} matrix having all non-negative values, can be carried on the current virtual topology. It is for such traffic matrices that we trigger the reconfiguration heuristic to calculate an optimal grooming objective weighted against the reconfiguration cost. As mentioned above, the aim of our reconfiguration heuristic is reoptimization under the changed traffic conditions, hence we consider only those traffic matrices that can be supported by the current virtual topology. If the current virtual topology cannot carry the new traffic, we might lose traffic and this is an highly undesirable situation. The aim of our reconfiguration heuristic is to avoid such situations.

We do not reconfigure any network component unless any of the T_{diff} matrix elements touch the **threshold value**. The threshold value is a limit that we use to determine the *s-d* pairs that are carrying traffic very close to the current over-provisioned traffic. The threshold value is just an indicator used to determine the elements that are very close to the supported/over-provisioned amount of traffic for which the OXCs and the DXCs have been configured. The value of this threshold limit can be determined by the network administrator based on the general characteristics of the input traffic pattern. If the traffic changes slowly and steadily then a small value for threshold could be sufficient. On the other hand, if the traffic pattern changes abruptly and to a great extent, then it would be advisable to have a higher threshold limit. A high threshold limit would result in more and frequent reconfigurations but would less frequently result in an unsupported traffic pattern. Any difference matrix T_{diff} element which touches the threshold value, falls in the **critical region**. We perform reconfiguration whenever any traffic element falls in the critical region, to bring the elements below the threshold limit. As long as the difference matrix elements are in the critical region, the current virtual topology can carry the new traffic. However, if these elements are not brought below the threshold, the virtual topology could fail to support them for a small increase in traffic in near future.

As we mentioned at the beginning of the chapter, our reconfiguration heuristic distinguishes between two types of reconfiguration efforts, namely local reconfiguration and global reconfiguration. When the change in the traffic matrix is small and affects few nodes, then local reconfiguration is performed reconfiguring very few OXCs and DXCs. The condition when local reconfiguration is needed is called *Soft Decision Criterion*, while *Hard Decision Criterion* requires global reconfiguration. We determine soft and hard decision criteria based on the ratio of the total number of lightpaths that carry sub-wavelength elements which are in the critical region. This is because the

number of lightpaths carrying elements in the critical region determines the extent of traffic change in the network and also the amount of reconfiguration effort required to handle such a change. We discuss the soft and the hard decision criterion in the following sections.

5.3 Hard Decision Criterion

Hard decision criterion is a situation where the traffic matrix has changed to such an extent that many lightpaths now carry sub-wavelength elements that are in the critical region. Whenever a hard decision is made we perform global reconfiguration of the network requiring many network components to be reconfigured. Thus, if the ratio of the total number of lightpaths that carry sub-wavelength elements which are in the critical region is large (say > 0.5), the hard decision criterion is met. (We discuss this ratio in more detail in the next section.) Whenever hard decision is hit, we solve the integrated ILP described in Section 4.3 to obtain a new groomed solution $\mathcal{G}'$. However, we perform reconfiguration to adapt the new $\mathcal{G}'$ only if the integrated ILP results in an objective function which is greater than 0. A positive objective function value means the grooming gain exceeds atleast δ times the reconfiguration cost and hence performing reconfiguration of the network components to implement $\mathcal{G}'$ would result in a significant reduction in the grooming cost weighted against the reconfiguration cost. On the other hand, if the objective function is negative or zero then we do not perform any reconfiguration because such a reconfiguration would not be justified in terms of the integrated objective. In situations when hard decision criterion is made but the integrated ILP results in a negative objective value, in a real-life scenario it would be appropriate to issue an alert message to the network administrator, informing him/her that many traffic demands are in the critical region but no reconfiguration can be performed cost effectively. This could act as a trigger for the administrator/ operator to consider deploying more network resources.

We perform reconfiguration for hard decision only when the reconfiguration results in a better grooming solution involving comparatively less cost of reconfiguration. However, every time we make a hard decision, we have to calculate the integrated ILP which is computationally expensive. Therefore, every time the traffic matrix changes, we do not call the integrated ILP since that would result in wasting the resources. We call the integrated ILP only when the hard decision is hit, while for soft decision we perform local reconfiguration and do not calculate the integrated ILP. We describe this next.

5.4 Soft Decision Criterion

We make a soft decision whenever the ratio of the total number of lightpaths carrying sub-wavelength elements that are in the critical region is small. The determination of the value

of $LPlimit$, the ratio that distinguishes between hard and soft criteria, is upto the discretion of the network administrator. We experiment with a few different values of $LPlimit$ to determine the effect of variations in $LPlimit$, and present our results in Chapter 6. We make a soft decision whenever comparatively less lightpaths carry the sub-wavelength elements that are in the critical region i.e. it is of local nature. Because of the local nature of the problem we require very few network components to be reconfigured, to get the sub-wavelength elements in the critical region below the threshold. We propose three different methods to perform local reconfiguration below. In making a soft decision we have reasonably assumed that when the change is of local nature, the grooming gain would also not be much. Therefore, if we call the integrated ILP for every local change, it will not only result in extra computational effort but probably the integrated ILP will also result in a negative objective function value resulting in no reconfiguration. When we make a soft decision, we try to minimize the amount of disturbance caused to the network by performing local reconfiguration. All these local reconfiguration methods limit the number of network components that have to be reconfigured and result in the reconfiguration of only the DXCs and not the OXCs. Thus in all the local reconfigurations the virtual topology is not changed, only the grooming of traffic on the existing lightpaths is changed which does not affect the OXCs. This is desirable since OXC reconfigurations are expected to be more significant disruptions, as we mentioned before.

5.4.1 Recalculation of Over-Provisioned Traffic Matrix

This is a local reconfiguration method where we recalculate the over-provisioned traffic matrix (refer to Section 5.1.3) to get the sub-wavelength elements in the critical region below the threshold limit. The over-provisioned traffic matrix is recalculated with the aim of increasing the over-provisioned share of those sub-wavelength elements which are in the critical region, either by utilizing any extra capacity on the lightpaths or by reducing the over-provisioned share of other sub-wavelength elements that are not in the critical region. We use the term **affected lightpaths** for all those lightpaths carrying sub-wavelength traffic elements that are in the critical region. This method of reconfiguration results in the reconfiguration of only the DXCs on the affected lightpaths. The OXCs will not require any reconfiguration because the only change that we make is in the amount of traffic each sub-wavelength element now carries. The DXCs need to be reconfigured to reflect such a change while the OXCs remain unaffected since they are still switching the same lightpaths. We propose the following two methods of recalculating the over-provisioned values when the traffic conditions in the network change, demanding reconfiguration.

Method 1: Recalculation of the over-provisioned traffic matrix only on the affected lightpaths. In this method, we recalculate the over-provisioned values for all the sub-wavelength elements only on the affected lightpaths, thus limiting the number of DXCs that would be reconfigured. Given the initial groomed solution $\mathcal{G}$, the over-provisioned traffic matrix T_{prov} and the new traffic matrix T_{new}, we calculate the new over-provisioned matrix $T_{newprov}$. In this method, we select the sub-

wavelength elements on the affected lightpaths and run the over-provisioning algorithm, discussed in Section 5.1.3, only on these elements. Here, we also make sure that the elements which were not in the critical region, do not enter the critical region even after the re-over-provisioning. We assure this by assigning the sub-wavelength elements on the affected lightpaths but not in the critical region, total traffic at least one unit more than the threshold value before using the over-provisioning algorithm. The detailed steps used in the implementation of this method are listed below.

1. For all the sub-wavelength elements $t_{new}^{(sd)}$ on all the affected lightpaths, we recalculate the traffic carried by them using the following equations:

   If $t_{prov}^{(sd)} - t_{new}^{(sd)} \neq threshold$
   then $t_{newtemp}^{(sd)} = t_{new}^{(sd)} + threshold + 1$
   else $t_{newtemp}^{(sd)} = t_{new}^{(sd)}$

   In this step we free any extra capacity on the affected lightpaths, making sure that the sub-wavelength elements that are not in the critical region, do not enter the critical region even after recalculations.

2. We then construct a new traffic matrix with these $t_{newtemp}^{(sd)}$ values, rest of the elements in the matrix are 0. After this we run the over-provisioning algorithm discussed in Section 5.1.3 on this new matrix $T_{newtemp}$ and get the new over-provisioned matrix $T_{tempprov}$.

3. The final recalculated over-provisioned matrix $T_{newprov}$ consists of all those non-zero elements from matrix $T_{tempprov}$ which are zero in matrix T_{prov}. Matrix $T_{newprov}$ has all those non-zero elements from matrix T_{prov} which are zero in matrix $T_{tempprov}$. While traffic elements which are non-zero in both the matrices $T_{tempprov}$ as well as T_{prov}, then we chose the smaller of the two values for the matrix $T_{newprov}$. All the elements that are zero in both $T_{tempprov}$ and T_{prov} remain zero in $T_{newprov}$ also.

We illustrate this method using the virtual topology given by Figure 5.5. Let the capacity C be 10, and the initial traffic matrix is given by T.

$$T = \begin{pmatrix} 0 & 1 & 3 & 0 \\ 0 & 0 & 2 & 1 \\ 5 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \end{pmatrix}$$

We calculate the over-provisioned traffic matrix using Iterative-Min method.

$$T_{prov} = \begin{pmatrix} 0 & 6 & 4 & 0 \\ 0 & 0 & 3 & 2 \\ 6 & 0 & 0 & 2 \\ 4 & 0 & 0 & 0 \end{pmatrix}$$

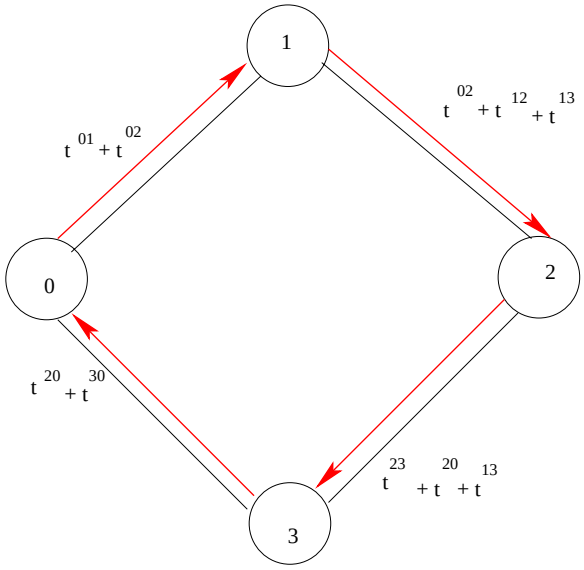

Figure 5.5: A Virtual Topology Representation Carrying Sub-Wavelength Traffic

We use a threshold limit equal to 0 in this example. The new traffic matrix is given by T_{new}, where $t_{new}^{(02)}$ touches the threshold and enters the critical region.

$$T_{new} = \begin{pmatrix} 0 & 1 & 4 & 0 \\ 0 & 0 & 2 & 1 \\ 5 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \end{pmatrix}$$

In this case the affected lightpaths are 01 and 12 as $t_{new}^{(02)}$ flows over these lightpaths. Let this be a soft decision criterion, in the following steps we show the calculation of new over-provisioned traffic matrix $T_{newprov}$. The intermediate traffic matrices calculated as described by the above steps are:

$$T_{newtemp} = \begin{pmatrix} 0 & 1 & 4 & 0 \\ 0 & 0 & 2 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \quad T_{tempprov} = \begin{pmatrix} 0 & 5 & 5 & 0 \\ 0 & 0 & 3 & 2 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}$$

The final over-provisioned matrix calculated is $T_{newprov}$. This recalculation gets $t_{new}^{(02)}$ out of the critical region by performing local reconfiguration.

$$T_{newprov} = \begin{pmatrix} 0 & 5 & 5 & 0 \\ 0 & 0 & 3 & 2 \\ 6 & 0 & 0 & 2 \\ 4 & 0 & 0 & 0 \end{pmatrix}$$

Method 2: Recalculation of the over-provisioned traffic matrix on a subset of lightpaths. This second method of recalculating the over-provisioned traffic matrix differs from the previous one only in terms of number of the sub-wavelength elements that are re-over-provisioned. In this method we do not confine our selection to all the sub-wavelength elements on the affected lightpaths. In addition, all the sub-wavelength elements on the lightpaths that have sub-wavelength elements of the affected lightpaths are also selected. After selecting the sub-wavelength elements to be re-over-provisioned, rest of the steps in this method are the same as the previous one. Usually, the total amount of re-over-provisioning done in this method is greater and therefore there is better chance of successful over-provisioning, but in this method the number of DXCs that need to be reconfigured are also typically more.

We illustrate this second method of re-over-provisioning using the example shown in Figure 5.5. We select all the sub-wavelength elements on the lightpaths 01, 12, 23 and 30 to be re-over-provisioned and following are the two intermediate matrices obtained in this method.

$$T_{newtemp} = \begin{pmatrix} 0 & 1 & 4 & 0 \\ 0 & 0 & 2 & 1 \\ 5 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \end{pmatrix} \quad T_{tempprov} = \begin{pmatrix} 0 & 5 & 5 & 0 \\ 0 & 0 & 3 & 2 \\ 6 & 0 & 0 & 2 \\ 4 & 0 & 0 & 0 \end{pmatrix}$$

The final over-provisioned traffic matrix $T_{newprov}$ obtained by this method also gets the element $t_{new}^{(02)}$ out of the critical region.

$$T_{newprov} = \begin{pmatrix} 0 & 5 & 5 & 0 \\ 0 & 0 & 3 & 2 \\ 6 & 0 & 0 & 2 \\ 4 & 0 & 0 & 0 \end{pmatrix}$$

The above two methods of recalculating the over-provisioned traffic result in a new traffic matrix $T_{newprov}$. But if recalculation of the over-provisioned traffic matrix does not result in a better solution in terms of the number of sub-wavelength elements that are still in the critical region, then we try another local reconfiguration method described in the next section.

5.4.2 Rerouting Traffic in the Critical Region

In this method of local reconfiguration, we try to reroute all the sub-wavelength elements that are in the critical region. We select all the sub-wavelength elements that are in the critical region, then we use Dijkstra's algorithm to find an alternate route for these sub-wavelength elements on the existing lightpaths. After determining the alternate route for a sub-wavelength element, we also make sure that the capacity constraint is not violated on any of the lightpaths in the alternate route. The detailed steps used in the implementation of this method are listed below.

1. We select a $t_{new}^{(sd)}$ element that is in the critical region, and trace all the lightpaths it flows on.

2. We use the Dijkstra's algorithm to find the shortest path from this s to d, on the virtual topology without the bottleneck lightpath (i.e. a lightpath on which the sum of sub-wavelength elements is the greatest for this s-d pair). If we cannot find an alternate route or the route length is greater than y hops/lightpaths then we try the next local reconfiguration method given in Section 5.4.3. We do not implement alternate routes that have more than y hops (say y is 70% of the total number of lightpaths in the network) because this requires at least y number of DXCs to be reconfigured, changing the nature of this problem from local to global.

3. If an alternate route is found, then we recalculate the traffic carried by the sub-wavelength elements in the alternate route by using the following equations:

   If $t_{prov}^{(sd)} - t_{new}^{(sd)} \neq threshold$
   then $t_{newtemp}^{(sd)} = t_{new}^{(sd)} + threshold + 1$
   else $t_{newtemp}^{(sd)} = t_{new}^{(sd)}$

   We perform this step to free any extra capacity on the lightpaths and also to make sure that the sub-wavelength elements that are not in the critical region remain out of the critical region even after rerouting.

4. We sum all the sub-wavelength elements $\sum t_{newtemp}^{(sd)}$, for all the lightpaths in the alternate route. If the sub-wavelength element $t_{new}^{(sd)}$ that was in the critical region does not already pass over this lightpath then we add that $t_{new}^{(sd)}$ value also in the summation.

5. If the summation on each of the lightpaths is less than or equal to C, then we implement this alternate route by deleting the sub-wavelength element, for which the alternate route is computed, from all the other lightpaths except the ones in this new route.

6. We perform the above steps for all the sub-wavelength elements in the critical region and recalculate the over-provisioned matrix only if an alternate route is found for at least one of sub-wavelength elements that was in the critical region.

The method of rerouting is illustrated using the example given in Figure 5.6. The virtual topology and the initial routing of traffic has been depicted in Figure 5.6. Say the new traffic matrix causes $t_{new}^{(02)}$ to enter the critical region, and the lightpath 12 is the bottle neck link because of which the over-provisioned traffic could not be increased for $t_{new}^{(02)}$. We run Dijkstra's algorithm on the virtual connection graph formed by all the lightpaths in the virtual topology except the lightpath 12. The output of Dijkstra's algorithm results in an alternate route for $t_{new}^{(02)}$, passing over lightpaths 01, then 13 and then 32 and if this path obeys the capacity constraint, then we implement this new route as shown in Figure 5.7.

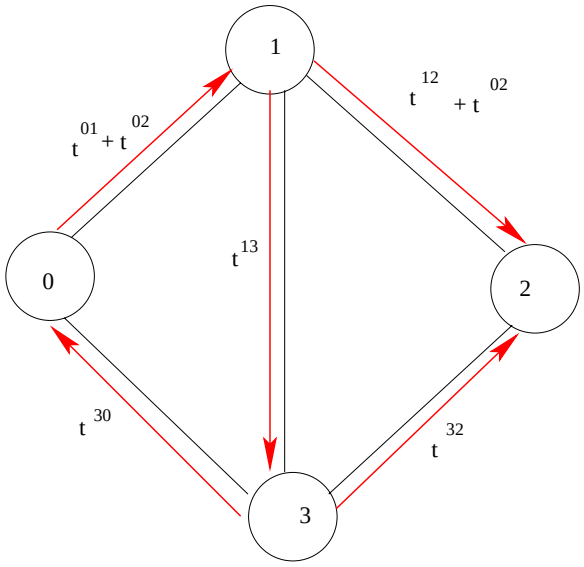

Figure 5.6: Virtual Topology Representation Carrying Sub-Wavelength Traffic

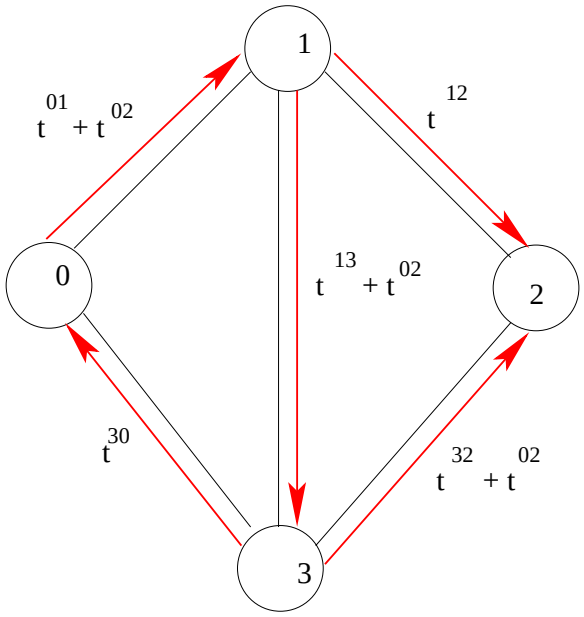

Figure 5.7: Virtual Topology After Rerouting was Performed

If this method of rerouting does not result in a solution where the number of sub-wavelength elements that are in the critical region is reduced, then we try the last local reconfiguration method explained in the next section.

5.4.3 Rerouting the Entire Traffic

After all the local reconfiguration methods fail to result in a new over-provisioned traffic matrix where the total number of sub-wavelength elements that are in the critical region are reduced, we try rerouting all the sub-wavelength elements on all the lightpaths. This method reroutes the entire traffic on all the established lightpaths by solving the following ILP. Clearly, rerouting less traffic components is more desirable, not only because of the smaller reconfiguration cost, but also because less unaffected traffic elements need to be rearranged. Hence the order of applying these three methods.

Given the virtual topology, the only sub-problem to be solved is the traffic routing sub-problem. The constraints that we use in this local rerouting ILP are the traffic routing sub-problem constraints from the integrated ILP given in Section 4.3.

Objective :

Minimize:

$$Maximum(\sum_{s,d} t_{ij}^{(sd)}), \forall i, j \tag{5.26}$$

Subject to:

$$t_{ij} = \sum_{s,d} t_{ij}^{(sd)}, \forall i, j \tag{5.27}$$

$$t_{ij} \leq b_{ij} C, \forall i, j \tag{5.28}$$

$$\sum_{j=0}^{N-1} t_{ij}^{(sd)} - \sum_{j=0}^{N-1} t_{ji}^{(sd)} = \left\{ \begin{array}{ll} t^{(sd)} & i = s \\ -t^{(sd)} & i = d \\ 0 & i \neq s, i \neq d \end{array} \right\} \forall s, d, i \tag{5.29}$$

The objective function of this ILP is to minimize the maximum traffic being carried by a lightpath. This objective function tries to increase the slack on a lightpath which carries the maximum traffic, this slack can now be used for over-provisioning. We propose this objective because usually it is the most loaded lightpath whose sub-wavelength elements are in the critical region as there is very little extra capacity on that lightpath to over-provision. After solving the ILP, we calculate the over-provisioned traffic matrix and implement the new matrix only if that results in fewer number of sub-wavelength elements in the critical region.

All the three local reconfiguration methods that we presented, aim at getting the maximum number of sub-wavelength elements out of the critical region. Because of being local in nature, we incur little reconfiguration cost in these methods. All these methods require only a few DXCs to be reconfigured, while the OXCs are never reconfigured. However, being local in nature and limited in their scope, these local methods may not always be able get the sub-wavelength elements out of the critical region. In situations where a soft decision criterion is hit but none of the local reconfigurations methods can reduce the number of elements in the critical region, then no reconfiguration is performed. In such situations it is again appropriate to generate an alert message to the network administrator informing that soft decision criterion was hit but no reconfiguration can be under taken cost effectively. In this case the appropriate action might also include examination of the algorithm parameters: for example, the hard/soft decision criterion parameter $LPlimit$ may need adjustment if this situation recurs frequently.

Any reconfiguration method, whether local or global, may not be able to get all the sub-wavelength elements out of the critical region because there might not be enough extra capacity in the network to do that. Some of the sub-wavelength elements that remain in the critical region even after a number of reconfigurations, could be an indication of scarcity of resources. After determining the locality of such elements the network administrator could also be informed in this case, to consider deploying more resources in that region.

Chapter 6

Numerical Results

In this chapter we discuss the numerical results of some of the experiments carried out. In all the experiments, we compare the performance our integrated approach with the "grooming-only" approach that only considers an optimal grooming objective and not the reconfiguration cost. We also perform experiments to compare the efficiency and performance of our heuristic, discussed in the previous chapter, as opposed to the "grooming-only" and the integrated approaches.

For the purpose of the experiments, the reconfiguration heuristic has been coded in C/C++ (discussed in Section 5.2). The exact solutions to the grooming approach and the integrated approach have been formulated as ILPs, discussed in Sections 4.2 and 4.3 respectively. We have coded the ILPs using ILOG CPLEX 7.1, a software designed to solve linear optimization problems. The concert technology in CPLEX provides a C++ modelling layer for linear and mixed integer programs, and a C++ interface for solving these problems. However, due to the complexity and the intractable nature of the problem, each instance that we run with CPLEX took from few minutes to several hours. For larger instances CPLEX took several days to give optimum results, hence for the exhaustive nature of our experiments we have used physical topologies of 6 nodes.

6.1 Instance Generation

An instance of the problem is given by a physical topology and an initial traffic matrix and a series of changing traffic matrices for which we observe the performance of the three approaches. For a given physical topology and a traffic matrix, we calculate the initial optimal groomed solution using the grooming ILP. All the three approaches start from this optimal groomed solution. Here, we discuss the various physical topologies and the evolution of traffic matrices that were used to

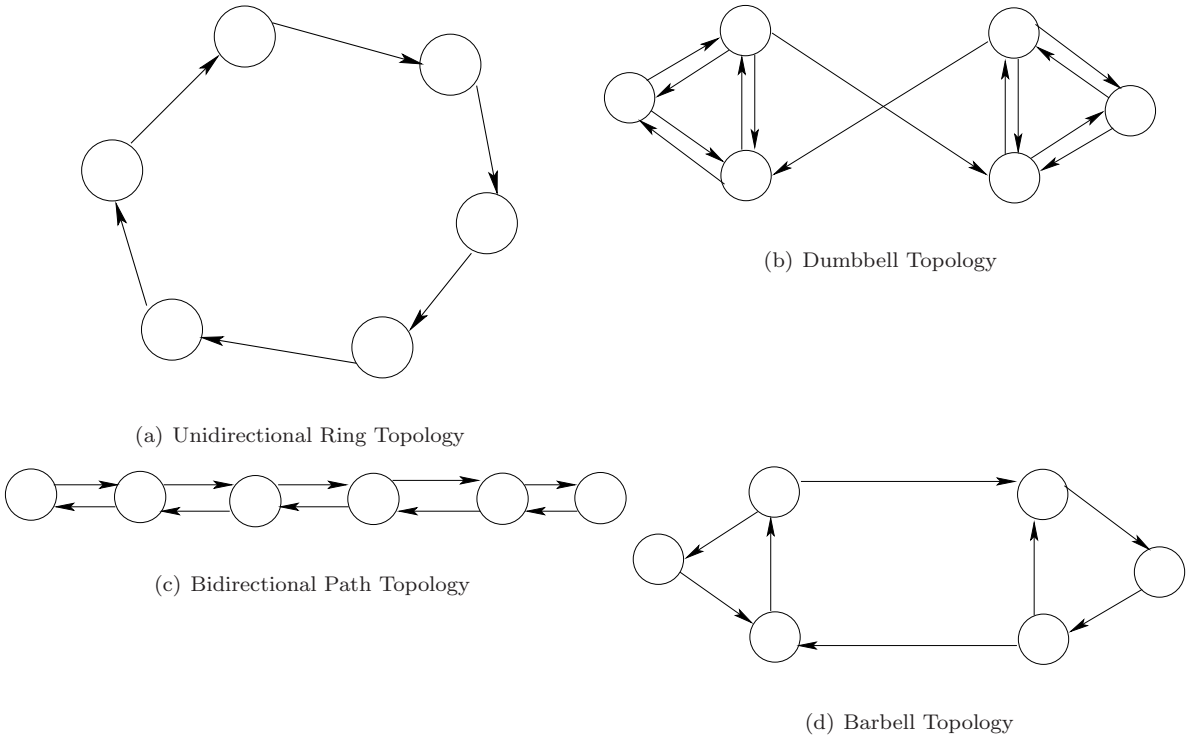

(a) Unidirectional Ring Topology

(b) Dumbbell Topology

(c) Bidirectional Path Topology

(d) Barbell Topology

Figure 6.1: Different Physical Topologies used for the Experiments

prove our assertions, made in the previous chapters.

6.1.1 Physical Topology

We performed the experiments for different physical topologies, in order to prove the generality of our proposed approach. The four physical topologies that we used for most of our experiments, have been shown in Figure 6.1. The directed arrows in the Figure 6.1 are used to represent the directed physical links.

6.1.2 Traffic Matrix Generation

We have generated the traffic matrices randomly having uniform distribution. The diagonal elements of the traffic matrix are always zero as there cannot be any traffic from a node to itself. All the other traffic matrix elements were generated having integer uniform distribution bounded between 0 and C. Here, we refer to a *feasible traffic matrix* as a traffic matrix that can be mapped onto the old virtual topology. Any traffic matrix that cannot be carried on the existing lightpaths without exceeding the capacity C of the lightpaths is said to be infeasible. As discussed in Section 5.2,

the reconfiguration heuristic is designed for feasible traffic matrices. Thus, each random instance of the traffic matrix we generated for the experiments was a feasible traffic matrix.

Traffic Matrix Evolution: The following three evolutions of the traffic matrices have been considered for experimentation.

Rising Traffic Evolution: In this rising traffic evolution all the feasible traffic matrices were generated in a way that the total traffic for any source-destination pair either increases or remains the same for all consecutive traffic matrix generations. All the traffic instances are generated having a uniform distribution with a slowly rising traffic. The initial traffic matrix was generated having integer uniform distribution bounded between 0 and $C/2$ and the rise in the traffic for a particular *s-d* pair has been generated having integer uniform distribution bounded between 0 and 2, which is added to the total current traffic for this *s-d* pair.

Falling Traffic Evolution: In this falling traffic evolution all the feasible traffic matrices were generated in a way that the total traffic for any source-destination pair either decreases or remains the same for all consecutive traffic matrix generations. All the traffic instances are generated having uniform distribution with a slowly falling traffic. The initial traffic matrix was generated having integer uniform distribution bounded between 0 and C and the fall in the traffic for a particular *s-d* pair has been generated having integer uniform distribution bounded between 0 and 2, which is subtracted from the total current traffic for this *s-d* pair.

Rising and Falling Traffic Evolution: In this rising and falling traffic evolution all the feasible traffic matrices were generated randomly having a uniform distribution. In each consecutive traffic matrix instance, some of the elements rise, while some fall randomly. This traffic evolution has been generated using integer uniform distribution bounded between 0 and C.

6.1.3 Experimental Parameters

For the purpose of the experiments, we have used reconfiguration cost function $RC - IV$ (Section 4.1.3) because this is the only linear cost function. Thus, all the reconfiguration cost calculations in the experiment are made using this cost function. The over provisioning method used is Iterative-Max-lightpath, with equal allocation as the initial over-provisioning amount. The experiments were carried out with the Iterative method since it has better performance than the Equal and Selective methods of over-provisioning (Section 5.1.4). The parameter δ (Section 3.2.2) was chosen to be equal to 5% of the total amount of initial traffic matrix. The reconfiguration cost of a single OXC α_4 and the reconfiguration cost of a single DXC β_2 (discussed in Section 4.1.3) is considered to be 1. The threshold value (Section 5.2) which determines whether a traffic matrix element is in the critical region or not, is 0 for the experiments. The grooming cost that we have considered throughout is the total electronic switching cost and is given by the following equation.

$$\sum_{s,d,i,j} d_{ij}^{(sd)} t^{(sd)} - \sum_{s,d} t^{(sd)} \tag{6.1}$$

The integrated objective function calculated is the objective of the integrated ILP (Section 4.3. These are all the values which remain constant for all the experiments, however we vary the following two parameters in the experiments to analyze their effect on the performance of our proposed methods.

- We have experimented with variations in the value of γ (Section 3.2.2), which is a parameter related to the average delay between reconfigurations. We have experimented with very high and low values of γ (2,7,15,200), though $\gamma = 7$ turned out to be a suitable value for most of our instances.

- The ratio of the total number of lightpaths that carry sub-wavelength elements in the critical region is referred as LPlimit (Section 5.4). We have experimented with two different LPlimit values, 70% and 30%. The value of LPlimit determines whether its a hard decision criterion or a soft decision criterion (Section 5.2). The values of LPlimit for the experiments were chosen to be far apart, so that we could clearly demonstrate the effect of this parameter on the performance of our reconfiguration heuristic.

6.2 Comparative Analysis

In this section, we analyze some of the results from the experiments on the basis of various input parameters. However, a more detailed and exhaustive set of results has been provided in Appendix A, which exhibit similar behavior. The results are plotted as graphs of four kinds. Each instance of the problem provides the following four graphs and the three lines in each graph represent the respective quantities for the "grooming-only" approach, the integrated approach and our heuristic.

- *Reconfiguration Cost:* This graph represents the total reconfiguration cost incurred in reconfiguring to the new virtual topology. The reconfiguration cost has been shown in Figure 6.4.

- *Grooming Cost:* This graph represents the total grooming cost in terms of the total amount of electronic switching performed at each instance. The grooming cost has been shown in Figure 6.5.

- *Integrated Objective:* This is the integrated objective for grooming and reconfiguration as proposed by us. Figure 6.9 shows the integrated objective, without plotting the constant factor δ, for all the three approaches.

- *Cumulation of the Integrated Objective:* We have also calculated the cumulation of the integrated objective because as mentioned earlier our heuristic proactively delays reconfiguration by planned over-provisioning, hence we observe its performance over the evolution of traffic matrices. The cumulation of the integrated cost has been plotted for all the three approaches as seen in Figure 6.2.

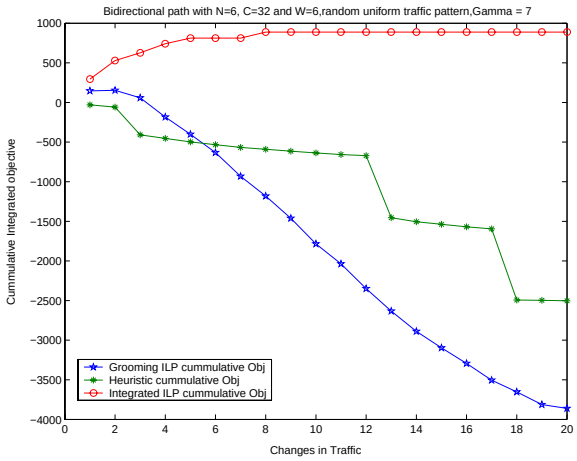

Figure 6.2: Cumulative Integrated objective for bidirectional path topology, under rising and falling traffic, for N=6, C=32, W=6, $LPlimit = 70\%$, $\gamma = 7$

We expect that the "grooming-only" approach will give the least/optimal grooming cost as it pays no attention to the reconfiguration cost. But we also expect that the integrated approach will not give very high grooming cost as compared to the optimal, however the heuristic might give high grooming cost at times. For the reconfiguration cost we expect that the "grooming-only" approach will have the maximum cost, while the integrated approach and the heuristic will incur less reconfiguration cost. We also expect that the integrated approach will give the maximum value for the integrated objective, with the heuristic following closely. However, we expect the "grooming-only" approach to give the least value for the integrated objective because it does not consider reconfiguration cost while calculating the groomed solution. Finally, we also expect that the integrated approach will give the maximum value for the cumulation of the integrated objective, with the heuristic following it closely, while the "grooming only" approach giving very less values.

6.2.1 Comparative Study on the basis of Physical Topology

In this section we analyze the effect of the physical topology on the performance of all the three approaches. As seen in Figure 6.2, the cumulative integrated objective for bidirectional path topology is the highest for the integrated ILP, while the performance of the grooming ILP is very low. This is because the grooming ILP performs reconfigurations involving high costs, though it always maintains an optimal grooming solution as can been seen in Figure 6.7. The reconfiguration heuristic performs local reconfigurations which some times results in a high grooming cost as seen for the initial instances in Figure 6.2. However, over a period of time the reconfiguration heuristic performs better than the grooming ILP in terms of the integrated objective. Similar results can be seen for a unidirectional ring topology in Figure A.5(d). The integrated ILP has high integrated

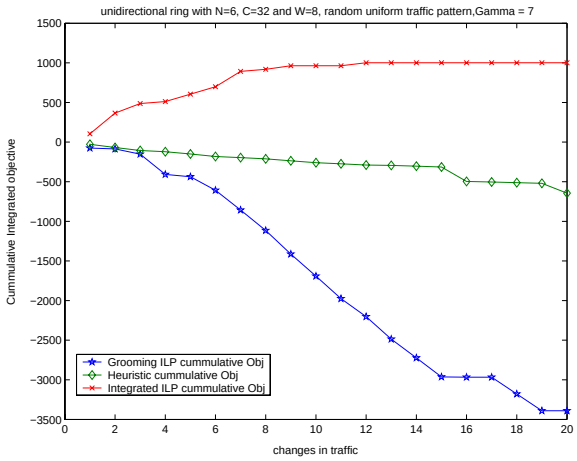

Figure 6.3: Cumulative Integrated objective for unidirectional ring topology , under rising and falling traffic, for N=6, C=32, W=8, $LPlimit = 70\%$, $\gamma = 7$

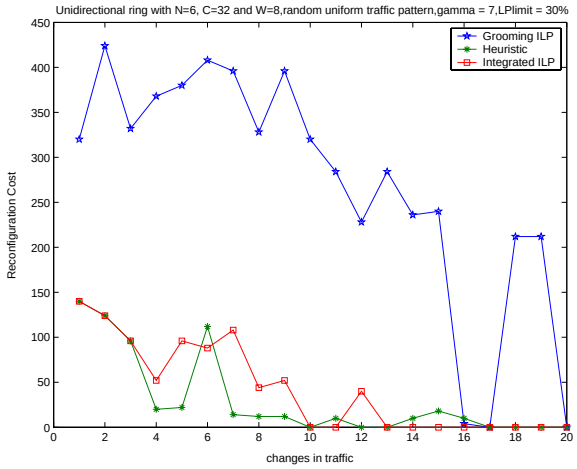

Figure 6.4: Reconfiguration cost for unidirectional ring topology , under rising and falling traffic, for N=6, C=32, W=8, $LPlimit = 30\%$, $\gamma = 7$

objective values because it performs less costly reconfigurations for more gain in grooming. We observe similar results for other virtual topologies as well in Figures A.9(d) and 6.2.

These results show the generality of our approach and indicate that our approach is robust to varying physical topologies. The performances of the integrated approach and the reconfiguration heuristic are good and consistent for all the physical topologies under study.

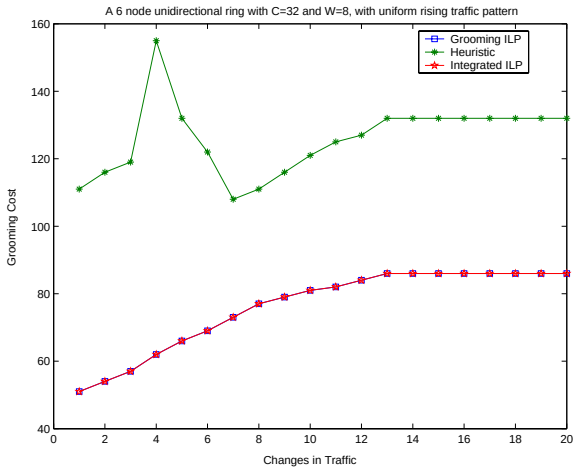

Figure 6.5: Grooming cost for unidirectional ring topology , under rising traffic, for N=6, C=32, W=8, $LPlimit = 70\%$, $\gamma = 200$

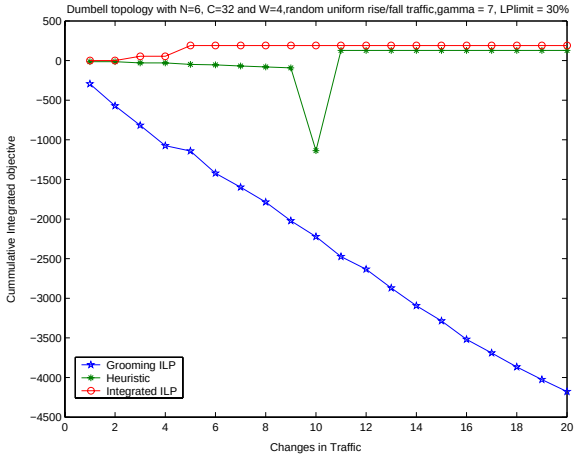

Figure 6.6: Cumulative Integrated objective for dumbbell topology , under rising and falling traffic, for N=6, C=32, W=4, $LPlimit = 30\%$, $\gamma = 7$

6.2.2 Comparative Study on the basis of Traffic Evolution

We observe the performance of the approaches for different physical topologies and traffic evolutions. Here we describe the performance of the different approaches under the following traffic evolutions.

Rising Traffic Evolution We studied the rising traffic evolution for all the physical topologies and observed that with the rise in traffic the grooming cost increases over the period of time. However, the integrated ILP follows very closely the optimal grooming cost given by the grooming ILP as seen in Figure A.5(b). The reconfiguration heuristic performs local reconfigurations

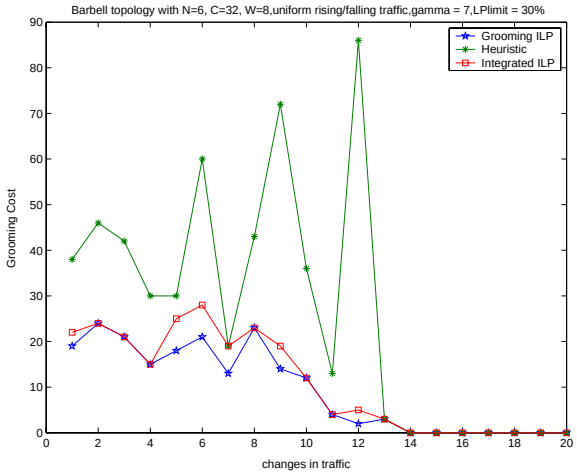

Figure 6.7: Grooming cost for barbell topology , under rising and falling traffic, for N=6, C=32, W=8, $LPlimit = 30\%$, $\gamma = 7$

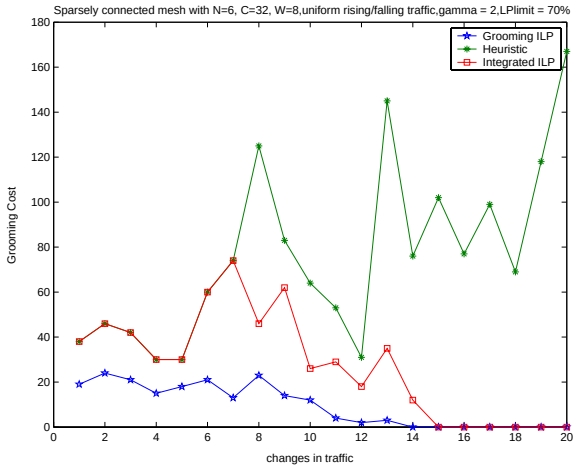

Figure 6.8: Grooming cost for barbell topology , under rising and falling traffic, for N=6, C=32, W=8, $LPlimit = 70\%$, $\gamma = 2$

at some instances due to sub-wavelength elements entering the critical region, and thus incurs a little higher grooming cost. However, by limiting the reconfiguration cost the heuristic performs better than the grooming ILP with respect to the integrated objective. The heuristic does not behave as good as the integrated ILP because of performing local reconfigurations, however the heuristic is a more practical solution as it calculates the ILP for much less instances. These observations have been made from various results shown in Figures A.1, A.7, A.11, A.17.

Falling Traffic Evolution We observe quite similar falling traffic evolution results for the different physical topologies, shown in Figures A.4, A.8 and A.14. The grooming cost of the

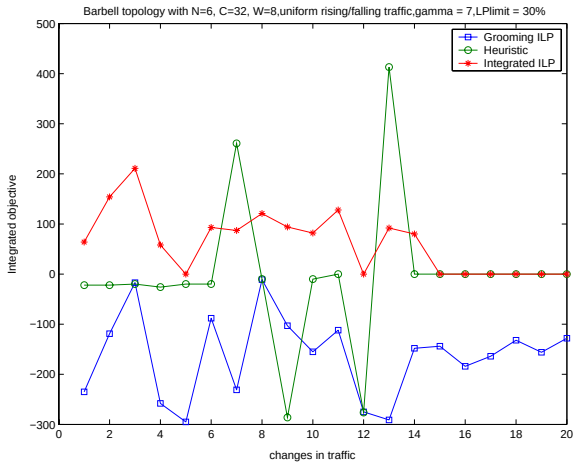

Figure 6.9: Integrated objective for barbell topology , under rising and falling traffic, for N=6, C=32, W=8, $LPlimit = 30\%$, $\gamma = 7$

integrated ILP and the heuristic are very close to the optimal grooming cost given by the grooming ILP. For the reconfiguration heuristic none of the elements entered the critical region after the initial over-provisioned traffic matrix was implemented, thus the reconfiguration cost is 0 for all the instances. We observed that the integrated ILP never performed reconfiguration because the grooming gain never justified the reconfiguration cost. On the hand, the grooming ILP frequently performed costly reconfigurations, resulting in low integrated objective values. However, the observation in Figure A.19 for the falling traffic was a little different. We observed that the heuristic performs some reconfigurations initially, this is because some of the elements remained in the critical region after the initial over-provisioned traffic matrix was implemented. And later when the traffic for some other sub-wavelength elements reduced, it was possible to get these elements out of the critical region by re-over-provisioning (Section 5.4.1), after which none of the elements entered the critical region as the traffic was falling continously.

Rising and Falling Traffic Evolution As shown in Figure 6.4, the reconfiguration cost for the grooming ILP remained high, while the heuristic and the integrated ILP resulted in comparatively less reconfiguration costs for the rising and falling traffic. The integrated objective function values observed in Figure 6.9 show that the integrated objective for the integrated ILP is always greater than that for the grooming ILP. However, the heuristic has a few troughs and crests in the integrated objective value, this is because the heuristic suffers from bad grooming cost at times but gradually the hard decision criterion is reached and the heuristic improves the grooming objective.

6.2.3 Comparative Study on the basis of Lplimit

As observed in Figure 6.6, when the LPlimit is 30% the performance of the heuristic is close to that of the integrated ILP with respect to the integrated objective. However, when Lplimit is 70% as seen in Figure 6.2, the performance of the heuristic deviates further away from the integrated ILP. The heuristic still performs better than the grooming ILP for the integrated objective. These observations are are validated from the fact that when the LPlimit is reduced from 70% to 30%, there are more instances when the heuristic makes a hard decision criterion and solves the integrated ILP, reducing the grooming cost. For the result shown in Figure 6.2, when LPlimit was 70% the heuristic calculated the local ILP for rerouting 4 times and the integrated ILP was never calculated. On the other hand, when LPlimit was 30% shown in Figure A.16(d) the heuristic calculated the local ILP for rerouting 4 times and the integrated ILP was calculated 3 times. However, the heuristic still successfully reduces the number of ILP computations even when LPlimit is 30%.

6.2.4 Comparative Study on the basis of γ

As seen in Figure 6.5, when γ is very high equal to 200 here, the integrated ILP gives optimal grooming costs as given by the grooming ILP. However, as seen in Figure A.24(b) with $\gamma = 15$ also the integrated ILP gives optimal grooming cost for almost all the instances. On the other hand, when γ is very low as seen in Figure 6.8, the integrated ILP gives higher grooming cost as opposed to the grooming ILP. For the value of γ intermediate between these two extreme values, as seen in Figure 6.7, the grooming cost calculated by the integrated ILP is closer to the optimal grooming cost as opposed to when γ is very low.

These observations were as expected because when γ is made very high, then reconfiguration cost is given very less consideration and the integrated objective tries to achieve optimal grooming costs. However, when γ is made very small then the integrated objective gives more importance to the reconfiguration cost and tries to reduce the reconfiguration cost at the expense of higher grooming cost. Thus, we selected an intermediate value of γ equal to 7 for most of our experiments as it seemed to balance the grooming cost and the reconfiguration cost for our instances.

Performance of our heuristic for a very long evolution of rising and falling traffic matrices: We performed an experiment for an evolution of 142 traffic matrices with the heuristic, to observe if its a feasible solution for such long evolutions. However, due to the limitation of CPLEX and time, every time the heuristic made a hard decision the integrated ILP was not solved. The results of this experiment have been shown in Figure A.29.

6.3 Summary of Results

In this chapter we presented some results that verify the claims made in the previous chapters. Below we summarize these results:

- The performance of the reconfiguration heuristic is robust to the variations in the underlying physical topology and the traffic in the network.
- The integrated approach and the reconfiguration heuristic perform better than the "grooming-only" approach which respect to the integrated objective for reconfiguration of sub-wavelength groomed optical networks.
- The reconfiguration heuristic effectively reduces the number of ILP calculations, making it a more practical approach.
- The integrated ILP gives maximum value for cumulation of the integrated objective, with the reconfiguration heuristic following closely the optimum value. However, the "grooming-only" approach deviates a lot from the optimal.
- The performance of the heuristic for the integrated objective improves by reducing the value of LPlimit, but at the expense of increased ILP computations. This is still significantly less than the computational effort in calculating the ILP for each instance.
- The grooming approach results in high reconfiguration cost as compared to our proposed approach, hence is not suitable for reconfiguration of sub-wavelength groomed optical networks.

Chapter 7

Conclusion and Future Work

We have proposed a new problem in optical network design and control, namely joint grooming and reconfiguration. We have argued the motivation for such an integrated approach, and formulated the problem as an Integer Linear Program. Since solving the ILP is not computationally attractive, we have designed a heuristic that minimizes the need for running the ILP, but achieves some of the benefits of the integrated approach nevertheless. Our numerical results validate our expectations both regarding the exact and the heuristic approaches. To our knowledge, this is the first attempt at an integrated treatment of these problems.

While we believe that our study is an important first step, much useful research remains to be done in this area. In the first place, more study is required of the effect of various parameters on the performance and behavior of our algorithms than we have done in this thesis. The various parameters, such as the relative cost weightage between optical switch reconfiguration and digital switch reconfiguration, or the threshold δ, are primarily meant to allow tuning to network characteristics. However, it would be interesting to see if the behavior of the algorithm changes qualitatively with different combinations of values for these parameters.

The basic algorithm itself can be enhanced in various ways. For example, all the local reconfiguration methods we use involve only the reconfiguration of the DXCs. It is conceptually straightforward to extend this, when those methods fail, a reconfiguration of both the OXCs and the DXCs. In such an extension we would attempt to establish new lightpaths, and perhaps delete the old ones (not carrying traffic any longer), only if needed to establish new ones. This also involves solving the complete grooming ILP, only the range of the variables decreases. However, heuristic grooming methods, many of which are available in literature for our cost function, could be used.

In conclusion, we have made a beginning on an interesting new problem area in optical networking. We believe much useful research will be performed in this area in the near future.

Bibliography

[1] R.K. Ahuja, T.L. Magnanti, and J.B. Orlin. *Network Flows*. Prentice Hall, 1993.

[2] I. Alfouzan and A. Jayasumana. Dynamic reconfiguration of wavelength-routed wdm networks. In *Proceedings of IEEE Conference on Local Computer Networks*, pages 477–85, 2001.

[3] D. O. Awduche. MPLS and traffic engineering in IP networks. *IEEE Communications*, 37(12):42–47, December 1999.

[4] K. Bala, G. Ellinas, M. Post, C-C. Shen, J. Wei, and N. Antoniades. Towards hitless reconfiguration in wdm optical networks for atm transport. In *Proceedings of IEEE GLOBECOM*, pages 316–20 vol.1, 1996.

[5] I. Baldine and G.N. Rouskas. Dynamic load balancing in broadcast wdm networks with tuning latencies. In *Proceedings IEEE INFOCOM'98 Conference on Computer Communications*, pages 78–85, 1998.

[6] I. Baldine and G.N. Rouskas. Traffic adaptive wdm networks: a study of reconfiguration issues. *Journal-of-Lightwave-Technology*, 19(4):433–55, 2001.

[7] D. Banerjee and B. Mukherjee. Wavelength-routed optical networks: linear formulation, resource budgeting tradeoffs, and a reconfiguration study. *IEEE/ACM Transactions on Networking*, 8(5):598–607, Oct 2000.

[8] R.J. Bates. *Optical switching and networking handbook*, pages 171–184. McGraw-Hill Professional, 2001.

[9] R. Berry and E. Modiano. Reducing electronic multiplexing costs in SONET/WDM rings with dynamically changing traffic. *IEEE Journal on Selected Areas in Communications*, 18(10):1961–1971, Oct 2000.

[10] A.L. Chiu and E.H. Modiano. Reducing electronic multiplexing costs in unidirectional SONET/WDM ring networks via efficient traffic grooming. In *Proceedings of IEEE Globecom*, volume 1, pages 322–327, Nov 1998.

[11] A.L. Chiu and E.H. Modiano. Traffic grooming algorithms for reducing electronic multiplexing costs in WDM ring networks. *Journal of Lightwave Technology*, 18(1):2–12, Jan 2000.

[12] C.J. Colbourn and A.C.H. Ling. Graph decompositions with application to wavelength add-drop multiplexing for minimizing SONET ADMs. *Discrete Mathematics*, 261:141–156, January 2003.

[13] ILOG CPLEX 7.1. http://www.iro.umontreal.ca/ gendron/IFT6551/CPLEX/.

[14] B. Davie and Y. Rekhter. *MPLS Technology and Applications*. Morgan Kaufmann Publishers, San Diego, California, 2000.

[15] R. Dutta, S. Huang, and G.N. Rouskas. On optimal traffic grooming in elemental network topologies. In *Proceedings of OPTICOMM*, 2003. (To appear).

[16] R. Dutta, S. Huang, and G.N. Rouskas. Traffic grooming in path, star and tree networks: NP-Completeness results and approximation schemes. In *Proceedings of ACM SIGMETRICS*, pages 298–299, 2003.

[17] R. Dutta and G.N. Rouskas. A survey of virtual topology design algorithms for wavelength routed optical networks. *Optical Networks Magazine*, 1(1):73–89, Jan 2000.

[18] R. Dutta and G.N. Rouskas. On optimal traffic grooming in WDM rings. *IEEE Journal of Selected Areas in Communications*, 20(1):110–121, Jan 2002.

[19] R. Dutta and G.N. Rouskas. Traffic grooming in WDM networks: Past and future. *IEEE Network*, 16(6):46–56, Nov/Dec 2002.

[20] P.H.H. Ernest, G. Mohan, and V. Bharadwaj. An efficient algorithm for virtual topology reconfiguration in wdm optical ring networks. In *Proceedings Tenth International Conference on Computer Communications and Networks*, pages 55–60, 2001.

[21] Fai-Siu and R.K.C. Chang. Optimal node assignment in reconfigurable wdm lightwave networks with regular virtual topologies. In *Proceedings of 3rd IEEE Symposium on Computers and Communications*, pages 375–9, 1998.

[22] M.R. Garey and D.S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W.H.Freeman and Company.

[23] A. Gencata and B. Mukherjee. Virtual-topology adaptation for wdm mesh networks under dynamic traffic. In *Proceedings of IEEE INFOCOM*, volume 1, pages 48–56, 2002.

[24] O. Gerstel, P. Lin, and G. Sasaki. Combined WDM and SONET network design. In *Proceedings of IEEE INFOCOM*, pages 734–743, 1999.

[25] O. Gerstel, R. Ramaswami, and G. Sasaki. Cost effective traffic grooming in WDM rings. In *Proceedings of IEEE INFOCOM*, pages 69–77, 1998.

[26] O. Gerstel, R. Ramaswami, and G. Sasaki. Cost effective traffic grooming in WDM rings. *IEEE/ACM Transactions on Networking*, 8(5):618–630, Oct 2000.

[27] W. Goralski. *Optical Networking & WDM*. McGraw-Hill.

[28] D. Guo and Z. Zhang. A scalable and reconfigurable wdm network architecture. In *Proceedings of the SPIE - The International Society for Optical Engineering*, volume 3552, pages 65–73, 1998.

[29] U. Gupta, D. Lee, and J. Leung. Efficient algorithms for interval graphs and circular-arc graphs. *Networks*, 12:459–467, 1982.

[30] G. Huiban, S. Perennes, and M. Syska. Traffic grooming in wdm networks with multi-layer switches. *IEEE*, pages 2896–2901, 2002.

[31] Wu. Kai, Zeng. Qingji, Xiong. Yizhi, and Yang. Xudong. Dynamic configuration of virtual topology through joint operation of packet-routing and wavelength-routing network. In *Proceedings of the SPIE - The International Society for Optical Engineering*, volume 4225, pages 181–5, 2000.

[32] A. Kershenbaum. *Telecommunications Network Design Algorithms*. McGraw-Hill, 1993.

[33] B. Kim, J-H. Eom, J. Park, and C-J. Chae. A lightpath routing reconfiguration algorithm for wdm optical networks. In *Proceedings of Fifth Asia Pacific Conference on Communications and Fourth Optoelectronics and Communications Conference*, pages 11–14 vol.1, 1999.

[34] V. Kuleshov and S. Banerjee. Minimal-disturbance topology reconfiguration in all-optical networks. In *Proceedings of the SPIE - The International Society for Optical Engineering*, volume 3230, pages 146–57, 1997.

[35] J-F.P. Labourdette. Traffic optimization and reconfiguration management of multiwavelength multihop broadcast lightwave networks. *Computer-Networks-and-ISDN-Systems*, 30(9-10):981–98, May 1998.

[36] J-F.P. Labourdette and A.S. Acampora. Logically rearrangeable multihop lightwave networks. *IEEE-Transactions-on-Communications*, 39(8):1223–30, Aug 1991.

[37] J-F.P. Labourdette, G.W. Hart, and A.S. Acampora. Branch-exchange sequences for reconfiguration of lightwave networks. *IEEE/ACM Transactions on Communications*, 42(10):2822–32, Oct 1994.

[38] J. Lacey. Optical switching and its impact on optical networks. *Optical Society of America*, 2000.

[39] G. Li and R. Simha. Efficient routing to reduce the cost of add-drop multiplexers in WDM optical ring networks. In *Proceedings of the SPIE*, volume 3843, pages 258–267, 1999.

[40] K.H. Liu, B.J. Wilson, and J.Y. Wei. A management and visualization framework for reconfigurable wdm optical networks. *IEEE Network*, 14(6):8–15, 2000.

[41] L. Liu, X. Li, P-J. Wan, and O. Frieder. Wavelength assignment in WDM rings to minimize SONET ADMs. In *Proceedings of IEEE INFOCOM*, pages 1020–1025, 2000.

[42] Lucent - WaveStar® OLS 1.6T Product Catalog. http://www.lucent.com/products/solution/0,,CTID+2021-STID+10482-SOID+392-LOCL+1,00.html.

[43] M. Mellia, E. Leonardi, M. Feletig, R. Gaudino, and F. Neri. Exploiting OTDM technology in WDM networks. In *Proceedings of IEEE INFOCOM*, pages 1822–1831, 2002.

[44] A. Narula-Tam, P.J. Lin, and E. Modiano. Wavelength requirements for virtual topology reconfiguration in wdm ring networks. In *Proceedings of IEEE International Conference on Communications*, pages 1650–4 vol.3, 2000.

[45] A. Narula-Tam and E. Modiano. Dynamic load balancing for wdm-based packet networks. In *Proceedings of IEEE INFOCOM*, pages 1010–19 vol.2, 2000.

[46] A. Narula-Tam and E. Modiano. Dynamic load balancing in wdm packet networks with and without wavelength constraints. *IEEE Journal on Selected Areas in Communications*, 18(10):1972–9, Oct 2000.

[47] A. Narula-Tam and E. Modiano. Dynamic reconfiguration in wdm packet networks with wavelength limitations. In *Proceedings of the Optical Fiber Communication Conference*, pages 210–12 vol.3, 2000.

[48] M.H.M. Nizam, D.K. Hunter, and D.G. Smith. A dynamic reconfiguring tool for improving multiwavelength transport network robustness. In *Proceedings of IEEE International Conference on Communications*, pages 246–50 vol.1, 1997.

[49] P.To Philip. Reconfigurability of shufflenets in multi-star implementation. In *Proceedings IEEE INFOCOM*, pages 1252–9 vol.3, 1994.

[50] C. Qiao, Y. Mei, M. Yoo, and Y. Zhang. Polymorphic control for cost-effective design of optical networks. *European-Transactions-on-Telecommunications*, 11(1):Jan.–Feb, 2000.

[51] B. Ramamurthy and A. Ramakrishnan. Virtual topology reconfiguration of wavelength-routed optical wdm networks. In *Proceedings of IEEE Global Telecommunications Conference*, pages 1269–75 vol.2, 2000.

[52] R. Ramaswami and K.N. Sivarajan. *Optical networks: a practical perspective*. Morgan Kaufmann Publishers, 1998.

[53] G. Sai Kiran Reddy, G. Manimaran, and C. Siva Ram Murthy. Reconfiguration based failure restoration in wavelength-routed wdm networks. In *Proceedings of International Conference on Dependable Systems and Networks*, pages 543–52, 2000.

[54] J. Simmons, E. Goldstein, and A. Saleh. Quantifying the benefit of wavelength add-drop in WDM rings with distance-independent and dependent traffic. *Journal of Lightwave Technology*, pages 48–57, Jan 1999.

[55] J. Simmons and A. Saleh. The value of optical bypass in reducing router size in gigabit networks. In *Proceedings of IEEE International Conference on Communications*, pages 591–596, Jun 1999.

[56] N. Sreenath, C.S.R. Murthy, B.H. Gurucharan, and G. Mohan. A two-stage approach for virtual topology reconfiguration of wdm optical networks. *Optical Networks Magazine*, 2(3):58–71, 2001.

[57] N. Sreenath, G.R. Panesar, and C.S.R. Murthy. A two-phase approach for virtual topology reconfiguration of wavelength-routed wdm optical networks. In *Proceedings of IEEE International Conference on Networks*, pages 371–6, 2001.

[58] Koundinya B. Srinivasarao. Traffic grooming in translucent optical ring networks. Master's thesis, North Carolina State University, August 2003.

[59] H. Takagi, Y. Zhang, X. Jia, and H. Takagi. Virtual topology reconfiguration for wide-area wdm networks. In *Proceedings IEEE ICCCAS*, pages 835–839, 2002.

[60] P-J. Wan, G. Calinescu, L. Liu, , and O. Frieder. Grooming of arbitrary traffic in SONET/WDM BLSRs. *IEEE Journal on Selected Areas in Communications*, 18(10):1995–2003, Oct 2000.

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

[62] Y. Wang and X. Chen. Comprehensive analysis of different optical interconnecting architectures. In *Proceedings of the SPIE The International Society for Optical Engineering*, volume 4582, pages 205–215, 2001.

[63] D.B West. *Introduction to Graph Theory*. Prentice Hall, 1996.

[64] G. Wilfong and P. Winkler. Ring routing and wavelength translation. In *Proceedings of the 9th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 333–341, 1998.

[65] X. Zhang and C. Qiao. On scheduling all-to-all personalized connections and cost-effective designs in WDM rings. *IEEE/ACM Transactions on Networking*, 7(3):435–445, Jun 1999.

[66] K. Zhu and B. Mukherjee. Traffic grooming in an optical WDM mesh network. *IEEE Journal on Selected Areas in Communications*, 20(1):122 –133, Jan 2002.

Appendix A

Experimental Results

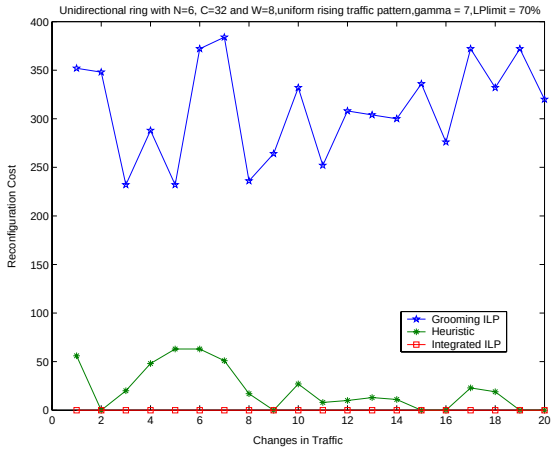

(a) Reconfiguration cost for unidirectional ring topology for N=6, C=32, W=8

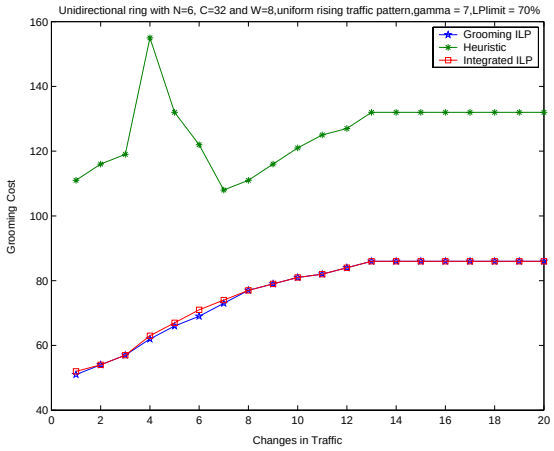

(b) Grooming cost for unidirectional ring topology for N=6, C=32, W=8

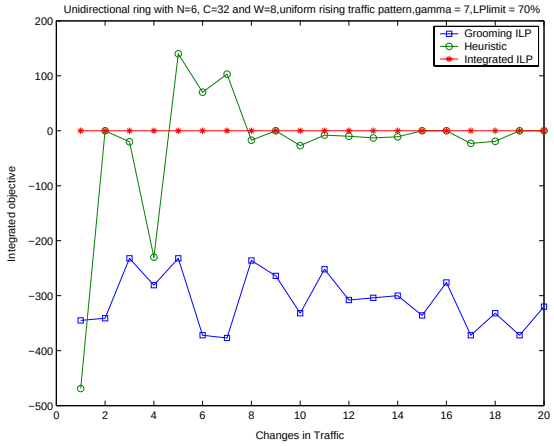

(c) Integrated objective function for unidirectional ring topology for N=6, C=32, W=8

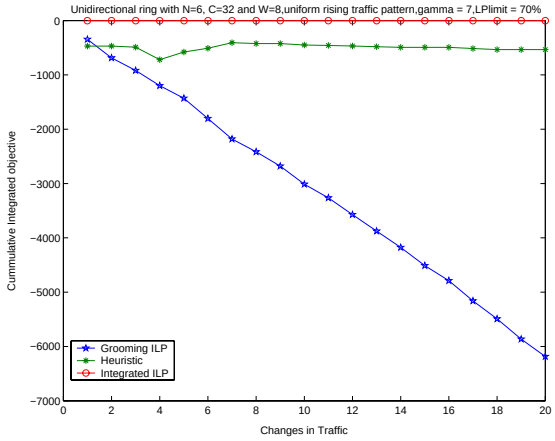

(d) Cumulative Integrated objective value for unidirectional ring topology for N=6, C=32, W=8

Figure A.1: Unidirectional ring performance for uniform rising traffic evolution with $LPlimit = 70\%$, $\gamma = 7$

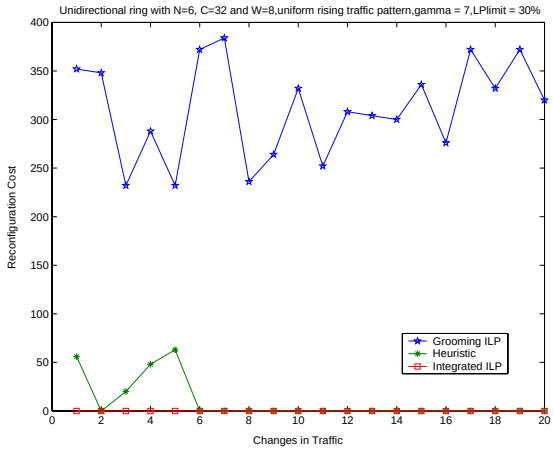

(a) Reconfiguration cost for unidirectional ring topology for N=6, C=32, W=8

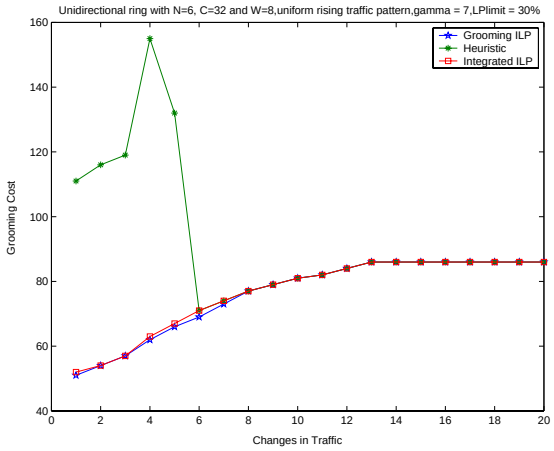

(b) Grooming cost for unidirectional ring topology for N=6, C=32, W=8

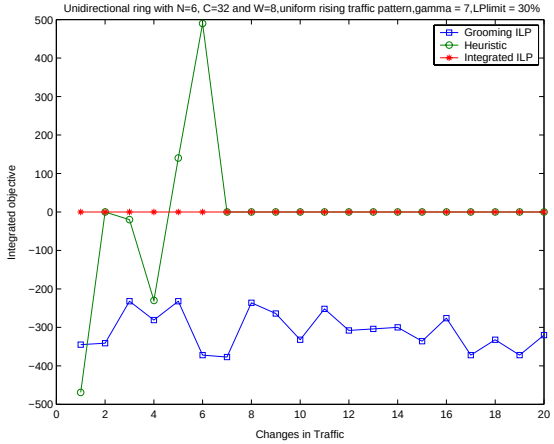

(c) Integrated objective function for unidirectional ring topology for N=6, C=32, W=8

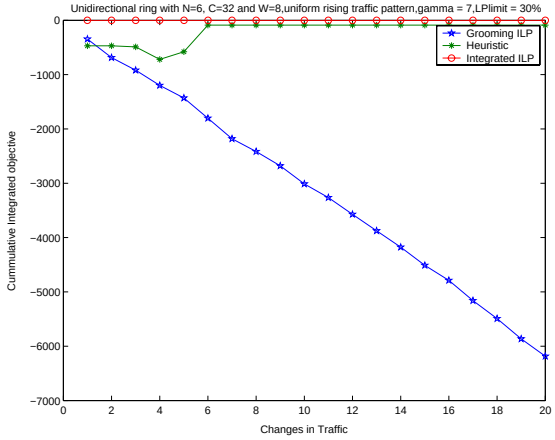

(d) Cumulative Integrated objective value for unidirectional ring topology for N=6, C=32, W=8

Figure A.2: Unidirectional ring performance under uniform rising traffic evolution with $LPlimit = 30\%$, $\gamma = 7$

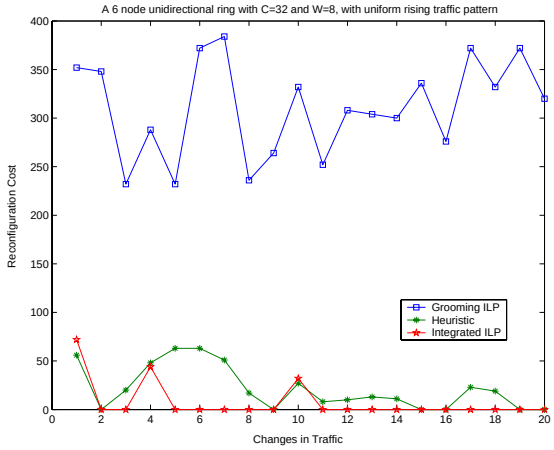

(a) Reconfiguration cost for unidirectional ring topology for N=6, C=32, W=8

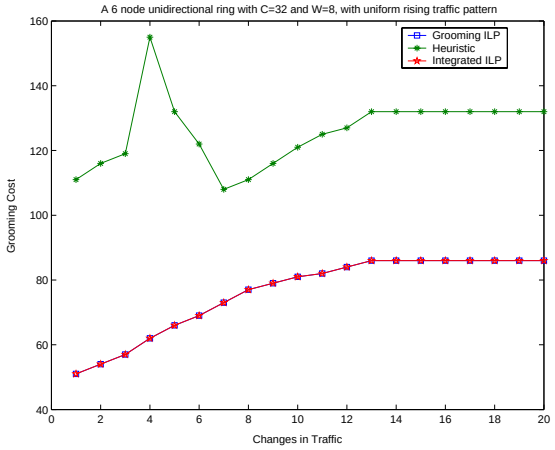

(b) Grooming cost for unidirectional ring topology for N=6, C=32, W=8

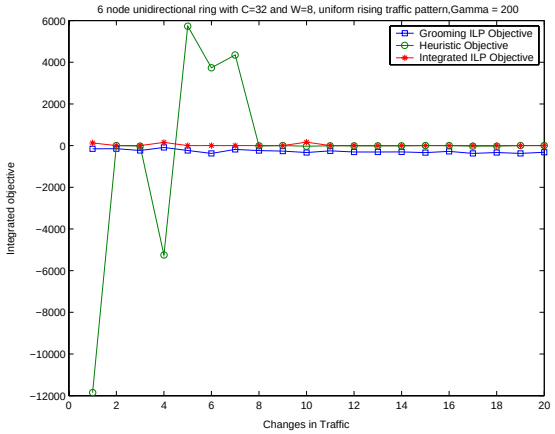

(c) Integrated objective function for unidirectional ring topology for N=6, C=32, W=8

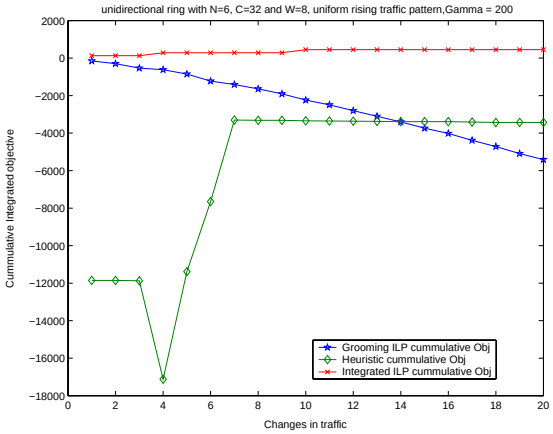

(d) Cumulative Integrated objective value for unidirectional ring topology for N=6, C=32, W=8

Figure A.3: Unidirectional ring performance under uniform rising traffic evolution with $LPlimit = 70\%$, $\gamma = 200$

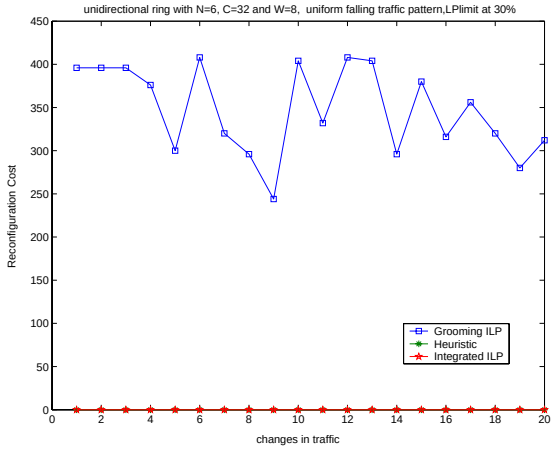

(a) Reconfiguration cost for unidirectional ring topology for N=6, C=32, W=8

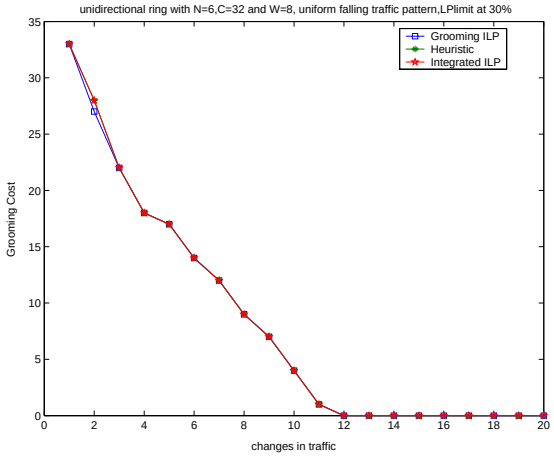

(b) Grooming cost for unidirectional ring topology for N=6, C=32, W=8

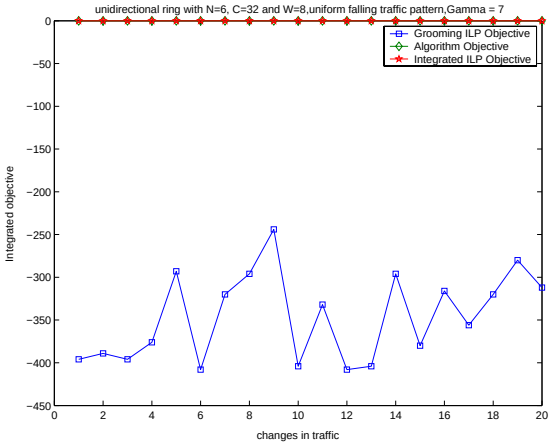

(c) Integrated objective function for unidirectional ring topology for N=6, C=32, W=8

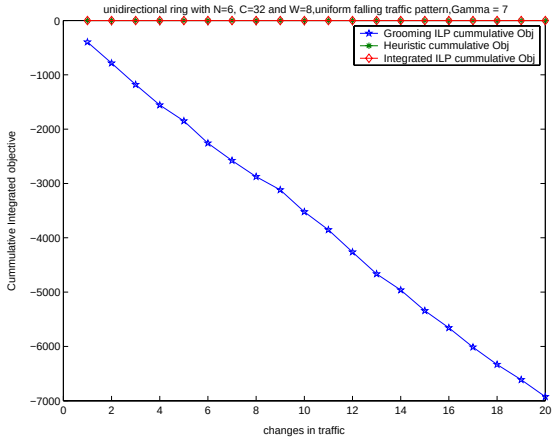

(d) Cumulative Integrated objective value for unidirectional ring topology for N=6, C=32, W=8

Figure A.4: Unidirectional ring performance under uniform falling traffic evolution with $LPlimit = 70\%$, $\gamma = 7$

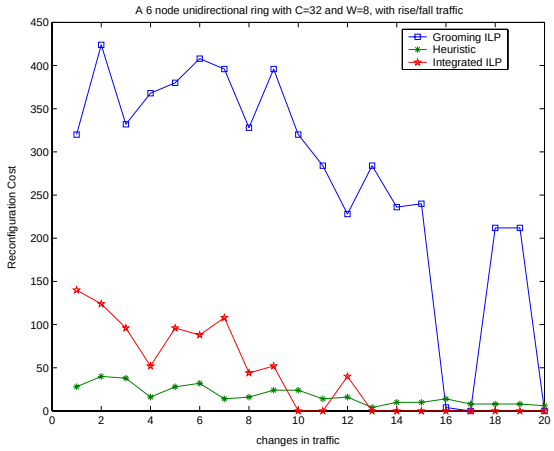

(a) Reconfiguration cost for unidirectional ring topology for N=6, C=32, W=8

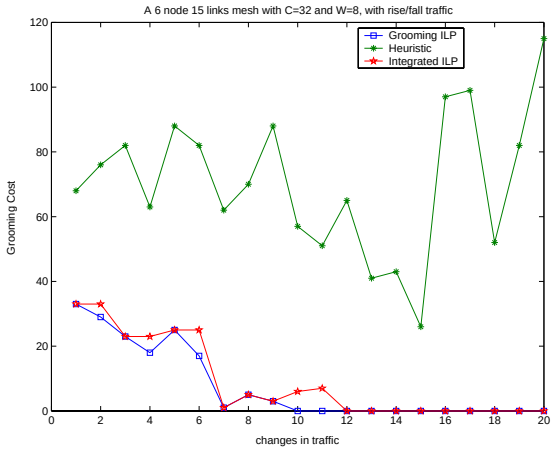

(b) Grooming cost for unidirectional ring topology for N=6, C=32, W=8

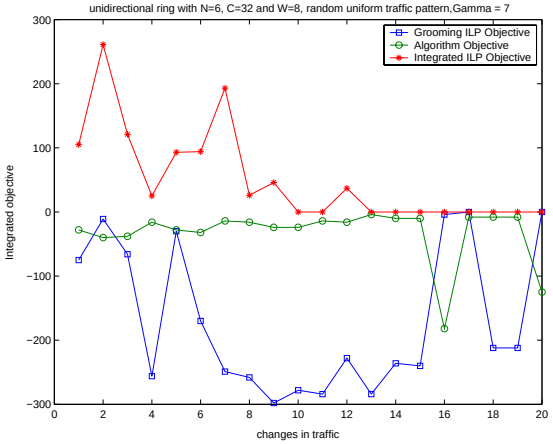

(c) Integrated objective function for unidirectional ring topology for N=6, C=32, W=8

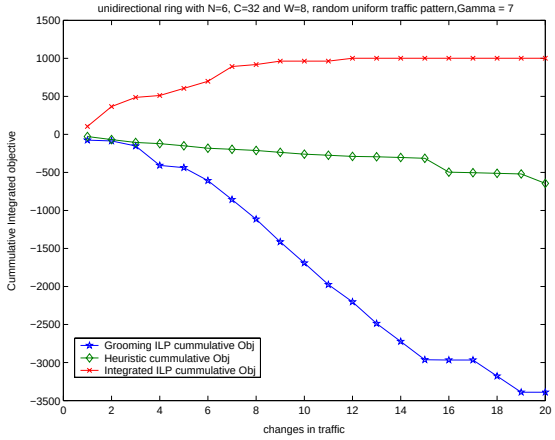

(d) Cumulative Integrated objective value for unidirectional ring topology for N=6, C=32, W=8

Figure A.5: Unidirectional ring performance under uniform rising and falling traffic evolution with $LPlimit = 70\%$, $\gamma = 7$

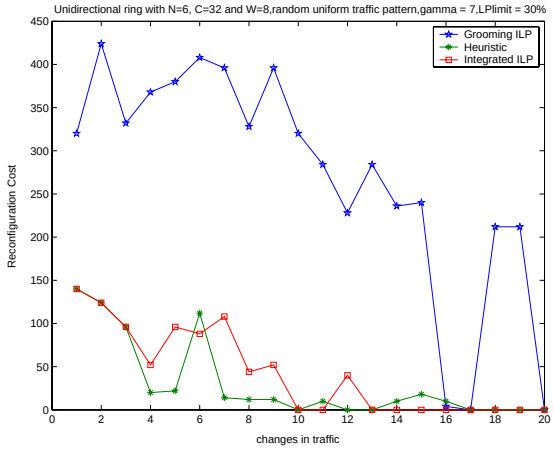

(a) Reconfiguration cost for unidirectional ring topology for N=6, C=32, W=8

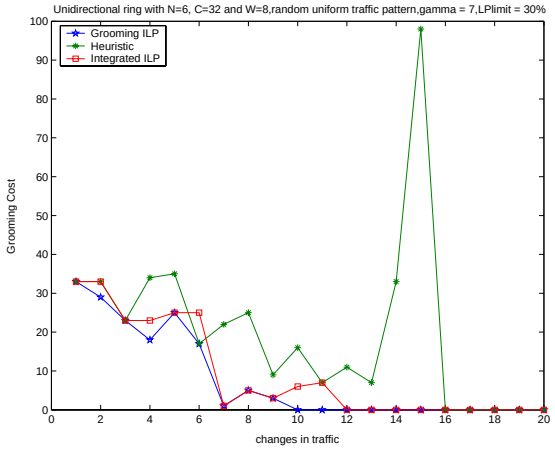

(b) Grooming cost for unidirectional ring topology for N=6, C=32, W=8

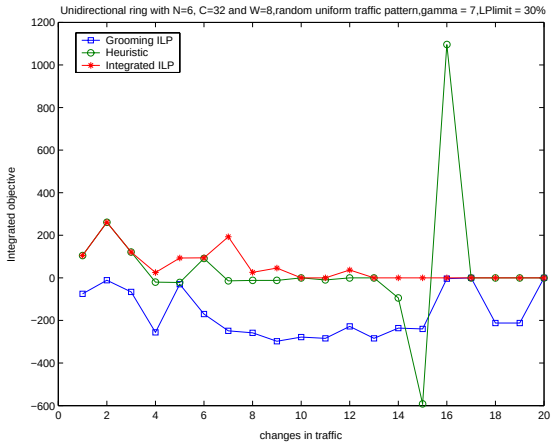

(c) Integrated objective function for unidirectional ring topology for N=6, C=32, W=8

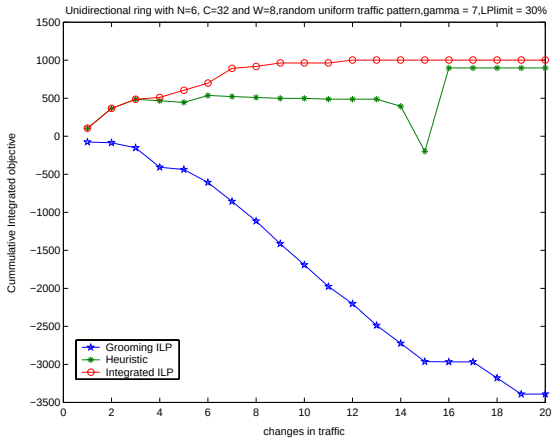

(d) Cumulative Integrated objective value for unidirectional ring topology for N=6, C=32, W=8

Figure A.6: Unidirectional ring performance under uniform rising and falling traffic evolution with $LPlimit = 30\%$, $\gamma = 7$

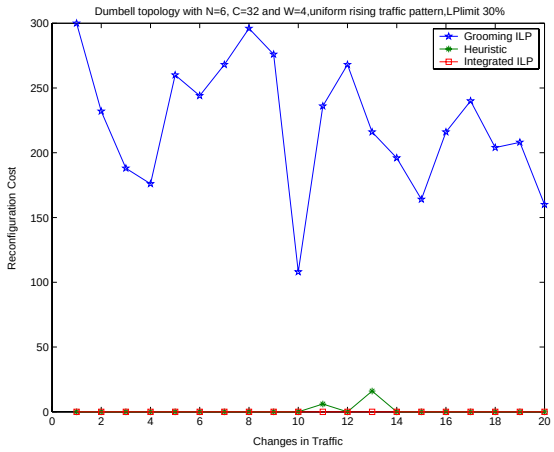

(a) Reconfiguration cost for dumbbell topology for N=6, C=32, W=4

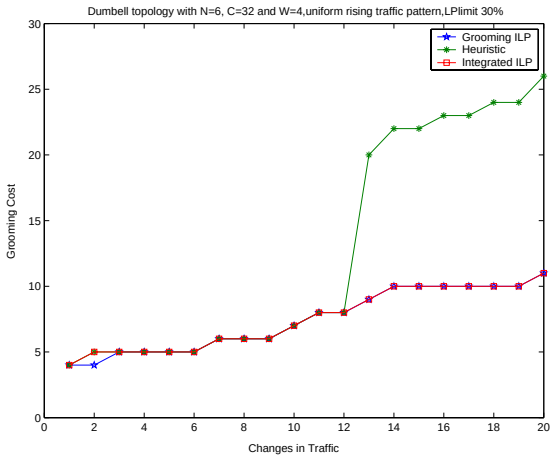

(b) Grooming cost for dumbbell topology for N=6, C=32, W=4

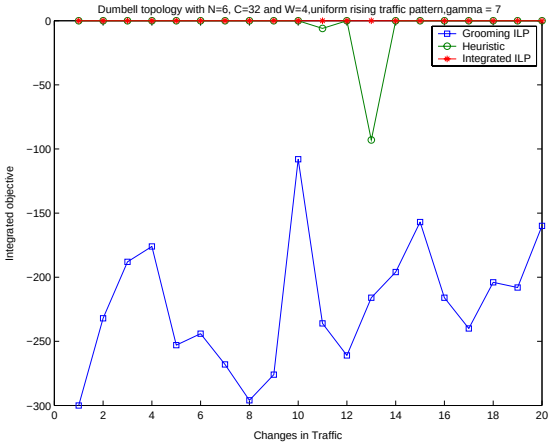

(c) Integrated objective function for dumbbell topology for N=6, C=32, W=4

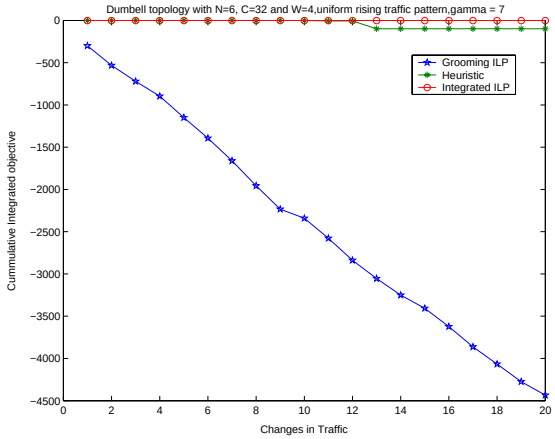

(d) Cumulative Integrated objective value for dumbbell topology for N=6, C=32, W=4

Figure A.7: Dumbbell topology performance under uniform rising traffic evolution with $LPlimit = 70\%$, $\gamma = 7$, same result obtained for $LPlimit = 30\%$

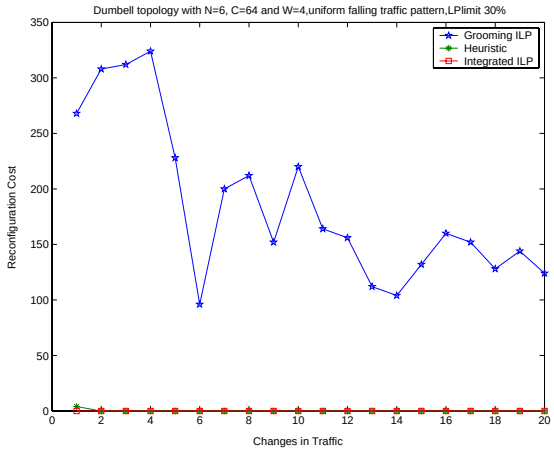

(a) Reconfiguration cost for dumbbell topology for N=6, C=32, W=4

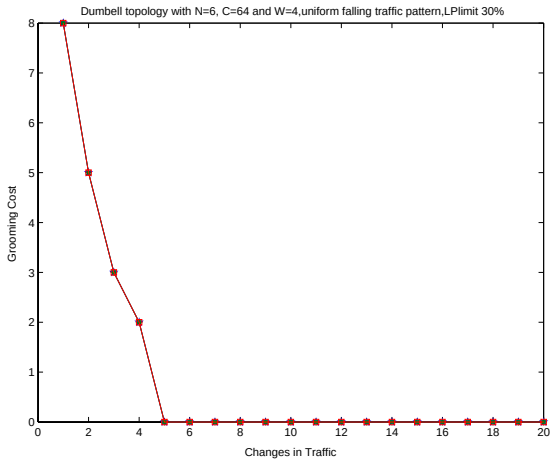

(b) Grooming cost for dumbbell topology for N=6, C=32, W=4

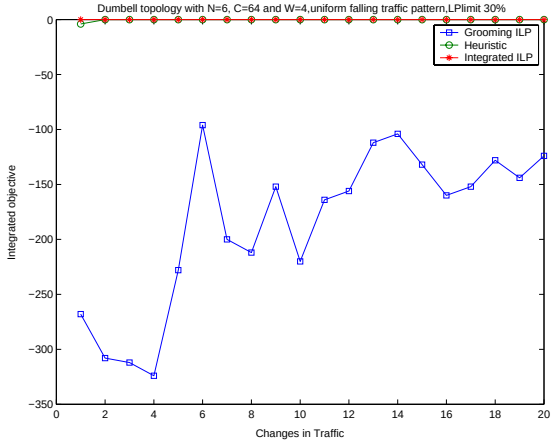

(c) Integrated objective function for dumbbell topology for N=6, C=32, W=4

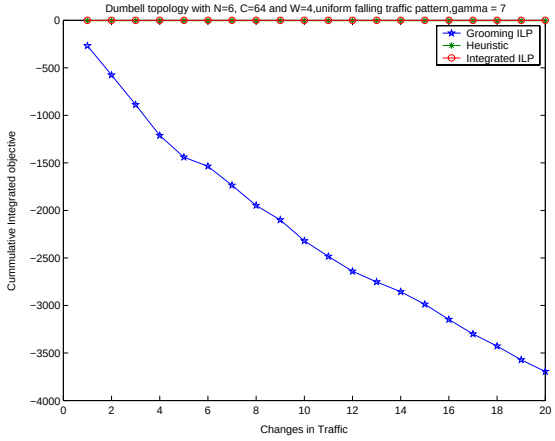

(d) Cumulative Integrated objective value for dumbbell topology for N=6, C=32, W=4

Figure A.8: Dumbbell topology performance under uniform falling traffic evolution with $LPlimit = 30\%$, $\gamma = 7$

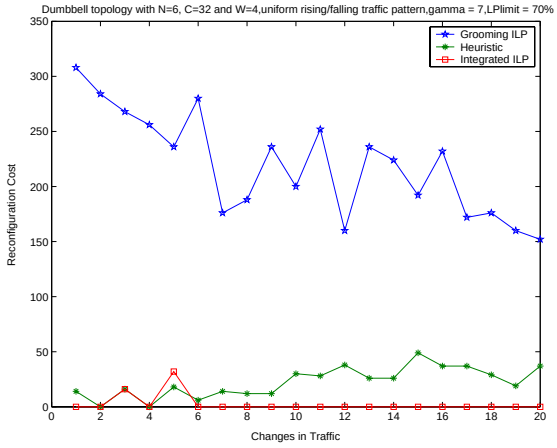

(a) Reconfiguration cost for dumbbell topology for N=6, C=32, W=4

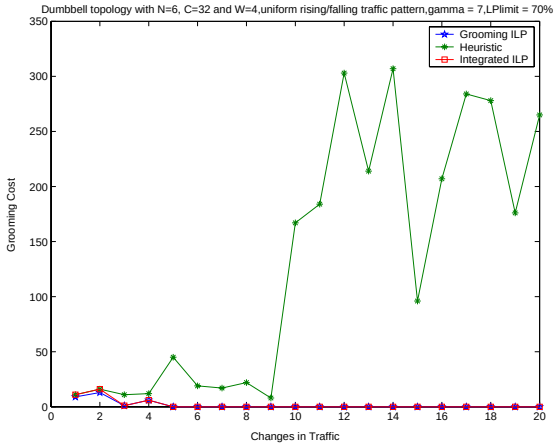

(b) Grooming cost for dumbbell topology for N=6, C=32, W=4

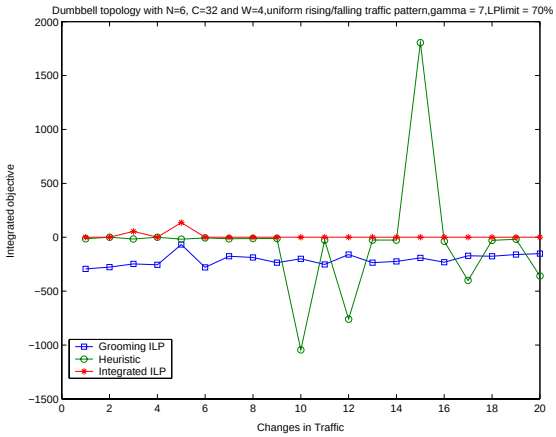

(c) Integrated objective function for dumbbell topology for N=6, C=32, W=4

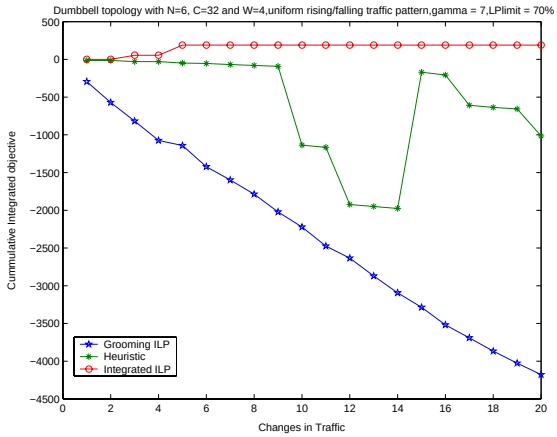

(d) Cumulative Integrated objective value for dumbbell topology for N=6, C=32, W=4

Figure A.9: Dumbbell topology performance under uniform rising and falling traffic evolution with $LPlimit = 70\%$, $\gamma = 7$

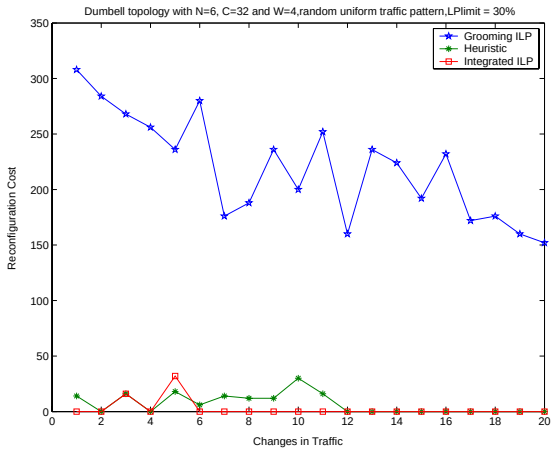

(a) Reconfiguration cost for dumbbell topology for N=6, C=32, W=4

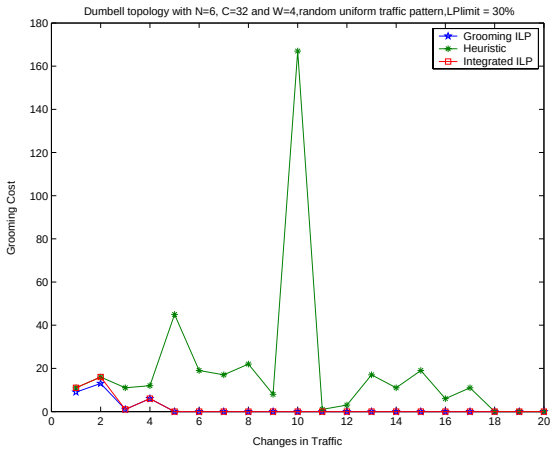

(b) Grooming cost for dumbbell topology for N=6, C=32, W=4

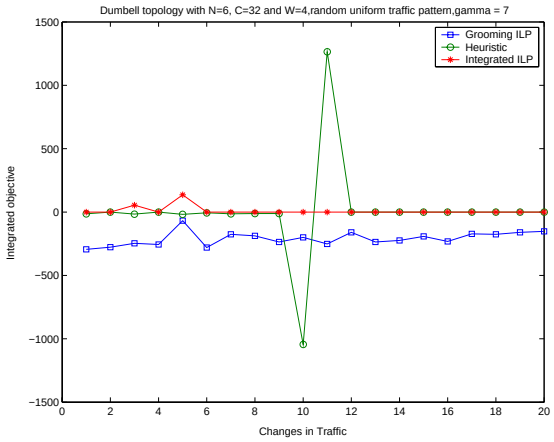

(c) Integrated objective function for dumbbell topology for N=6, C=32, W=4

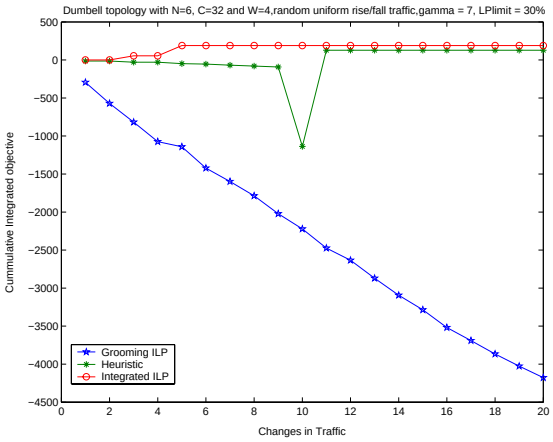

(d) Cumulative Integrated objective value for dumbbell topology for N=6, C=32, W=4

Figure A.10: Dumbbell topology performance under uniform rising and falling traffic evolution with $LPlimit = 30\%$, $\gamma = 7$

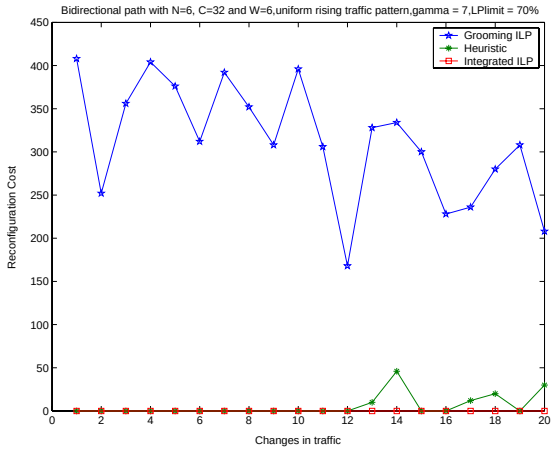

(a) Reconfiguration cost for bidirectional path topology for N=6, C=32, W=6

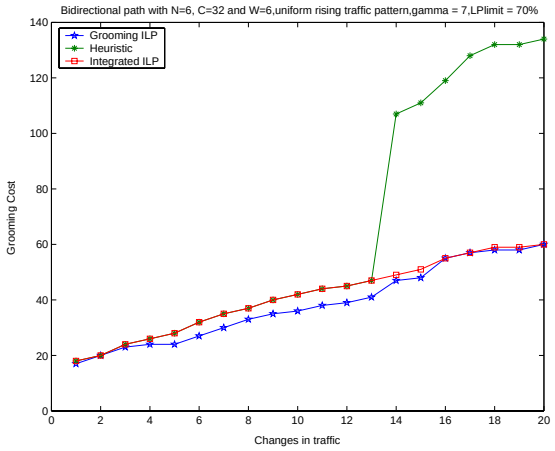

(b) Grooming cost for bidirectional path topology for N=6, C=32, W=6

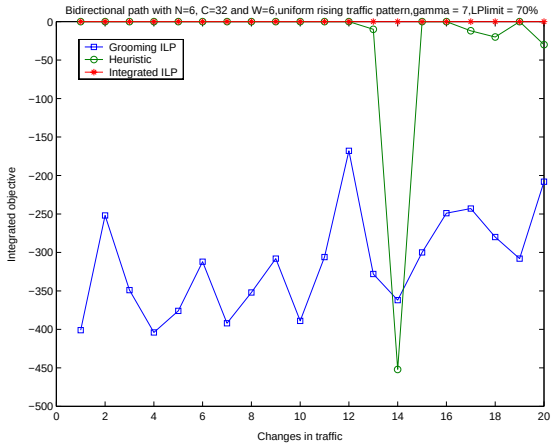

(c) Integrated objective function for bidirectional path topology for N=6, C=32, W=6

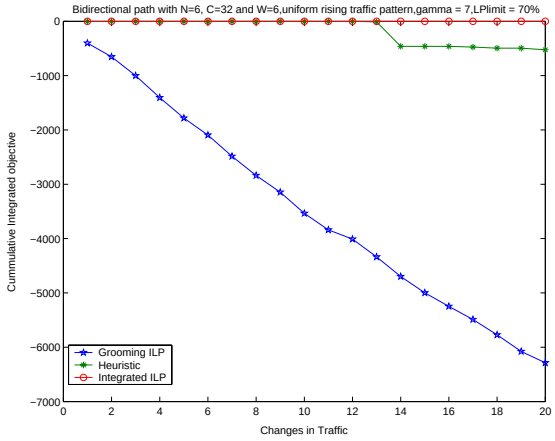

(d) Cumulative Integrated objective value for bidirectional path topology for N=6, C=32, W=6

Figure A.11: Bidirectional path performance under uniform rising traffic evolution with $LPlimit = 70\%$, $\gamma = 7$

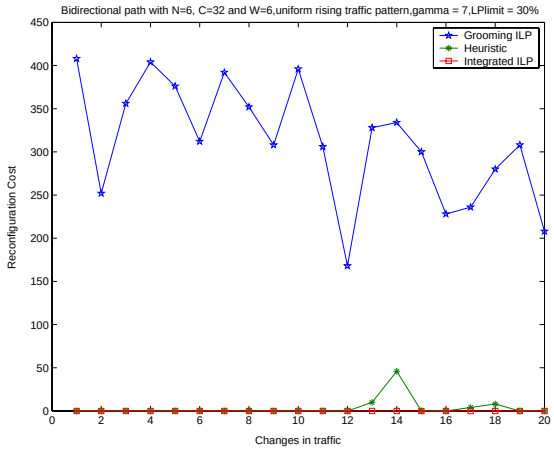

(a) Reconfiguration cost for bidirectional path topology for N=6, C=32, W=6

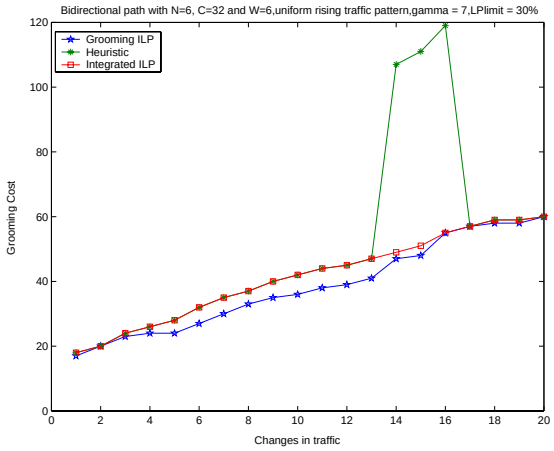

(b) Grooming cost for bidirectional path topology for N=6, C=32, W=6

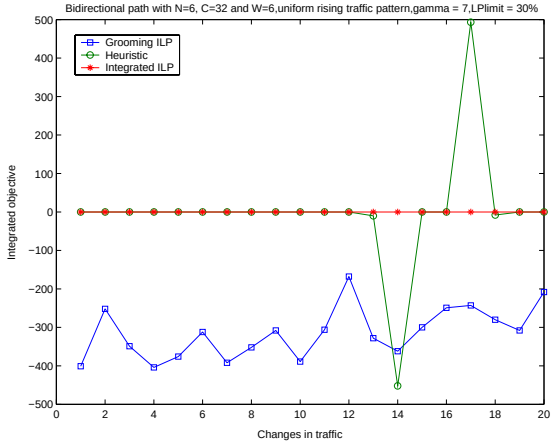

(c) Integrated objective function for bidirectional path topology for N=6, C=32, W=6

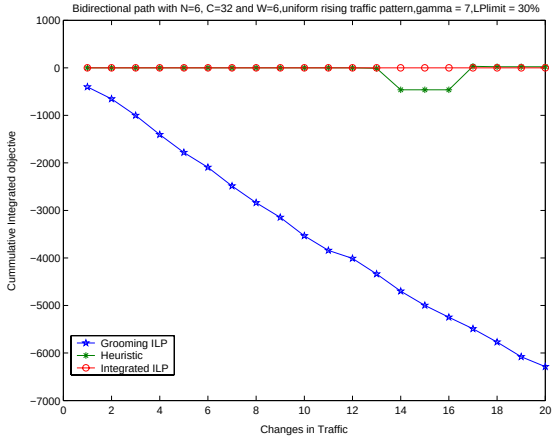

(d) Cumulative Integrated objective value for bidirectional path topology for N=6, C=32, W=6

Figure A.12: Bidirectional path performance under uniform rising traffic evolution with $LPlimit = 30\%$, $\gamma = 7$

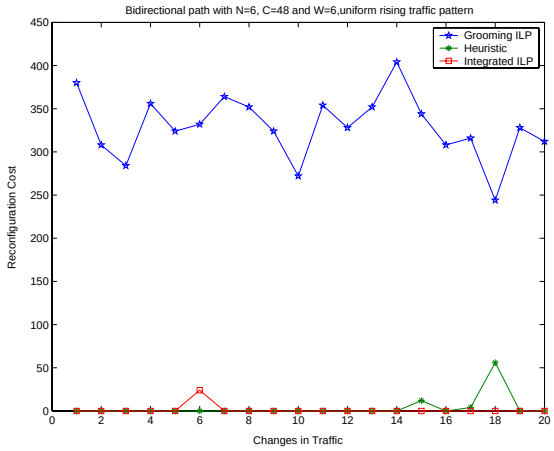

(a) Reconfiguration cost for bidirectional path topology for N=6, C=48, W=6

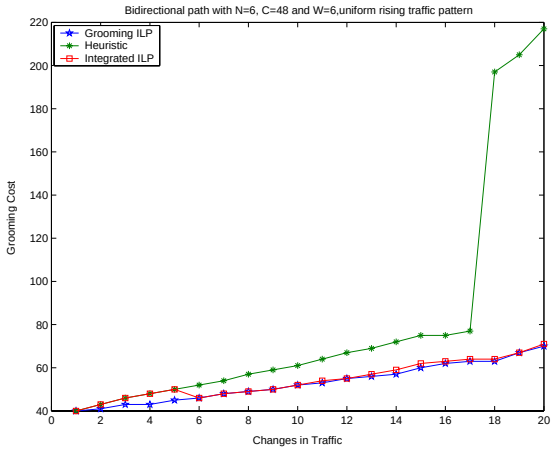

(b) Grooming cost for bidirectional path topology for N=6, C=48, W=6

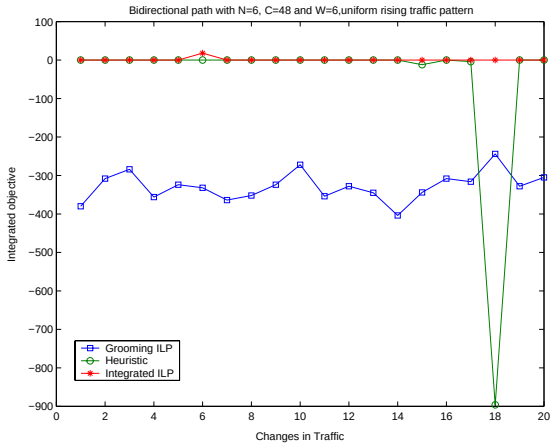

(c) Integrated objective function for bidirectional path topology for N=6, C=48, W=6

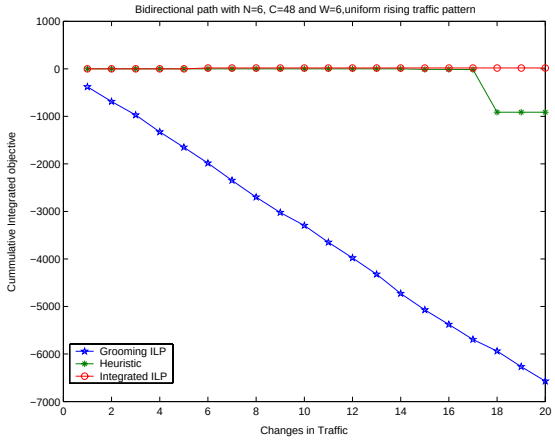

(d) Cumulative Integrated objective value for bidirectional path topology for N=6, C=48, W=6

Figure A.13: Bidirectional path performance under uniform rising traffic evolution with $LPlimit = 70\%$, $\gamma = 7$, same result obtained for $LPlimit = 30\%$

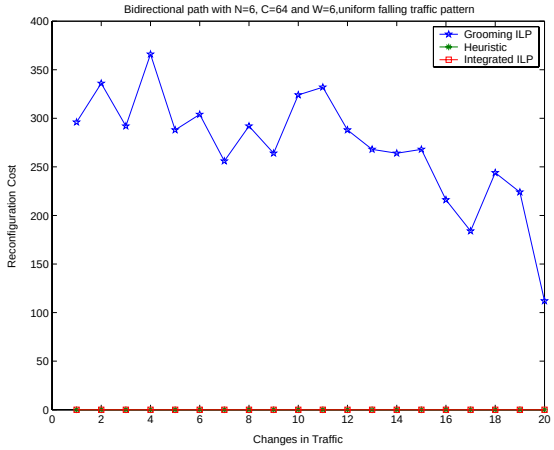

(a) Reconfiguration cost for bidirectional path topology for N=6, C=64, W=6

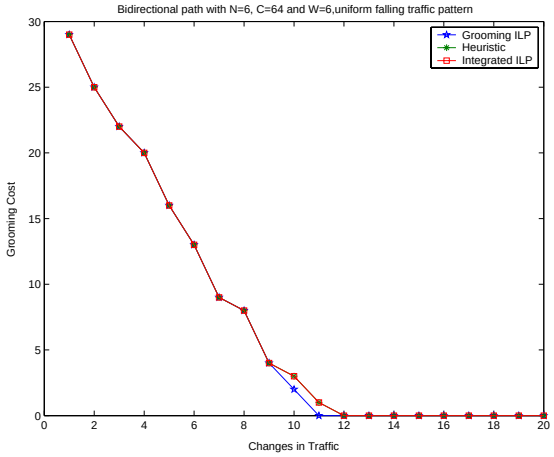

(b) Grooming cost for bidirectional path topology for N=6, C=64, W=6

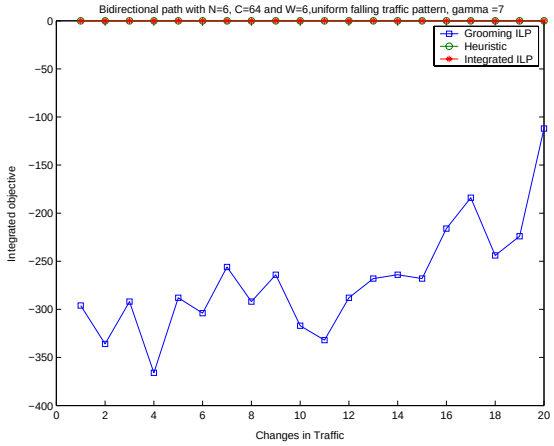

(c) Integrated objective function for bidirectional path topology for N=6, C=64, W=6

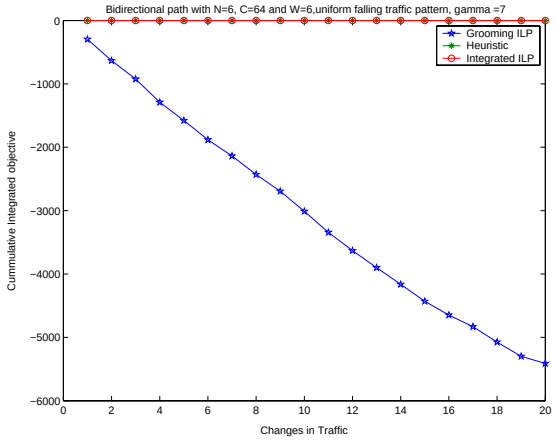

(d) Cumulative Integrated objective value for bidirectional path topology for N=6, C=64, W=6

Figure A.14: Bidirectional path performance under uniform falling traffic evolution with $LPlimit = 30\%$, $\gamma = 7$

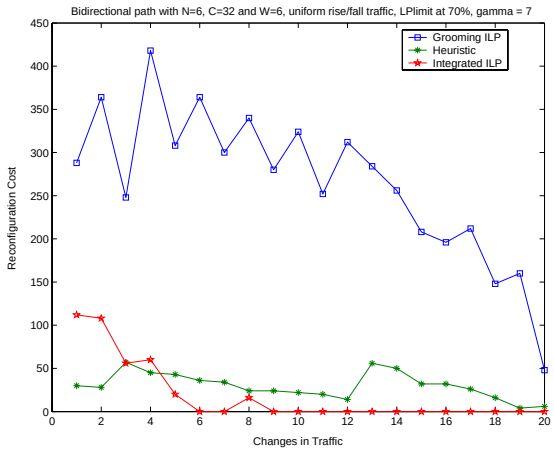

(a) Reconfiguration cost for bidirectional path topology for N=6, C=32, W=6

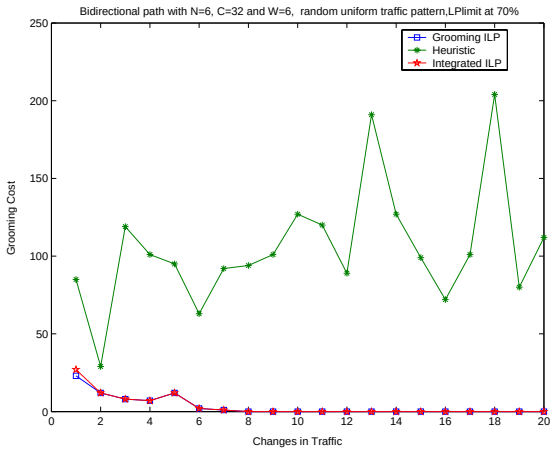

(b) Grooming cost for bidirectional path topology for N=6, C=32, W=6

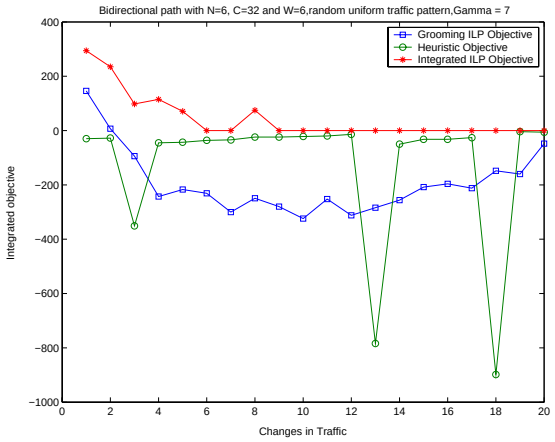

(c) Integrated objective function for bidirectional path topology for N=6, C=32, W=6

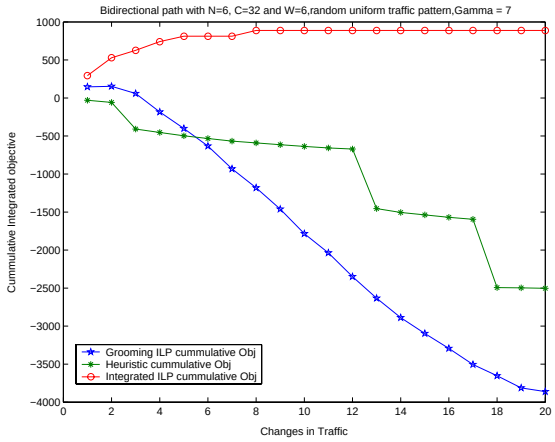

(d) Cumulative Integrated objective value for bidirectional path topology for N=6, C=32, W=6

Figure A.15: Bidirectional path performance under uniform rising and falling traffic evolution with $LPlimit = 70\%$, $\gamma = 7$

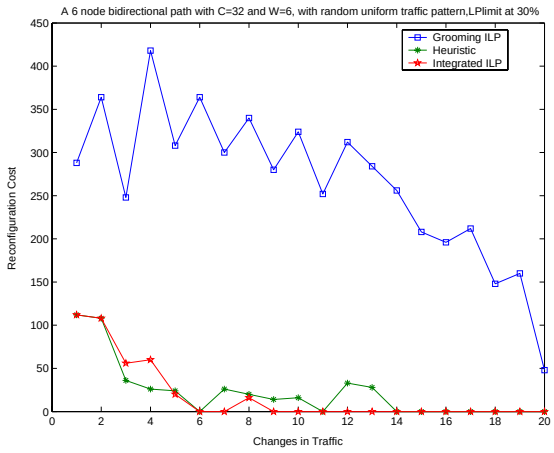

(a) Reconfiguration cost for bidirectional path topology for N=6, C=32, W=6

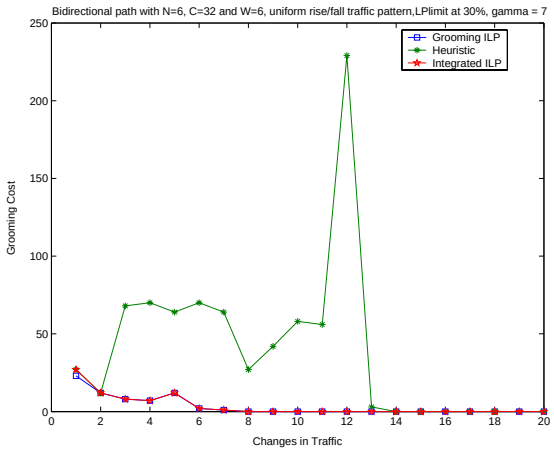

(b) Grooming cost for bidirectional path topology for N=6, C=32, W=6

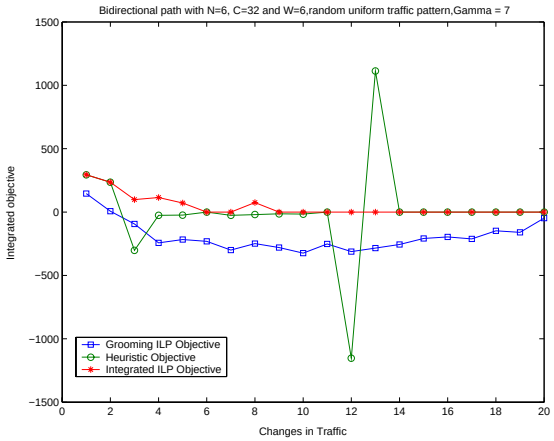

(c) Integrated objective function for bidirectional path topology for N=6, C=32, W=6

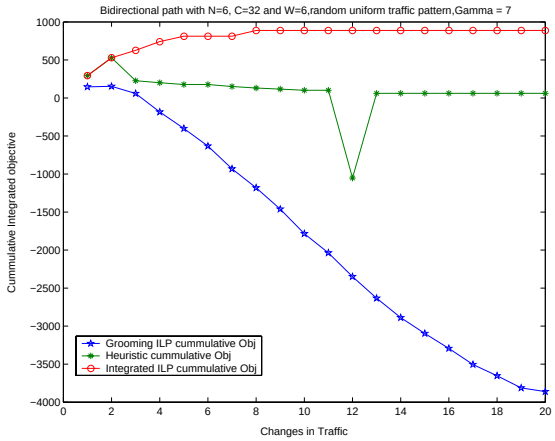

(d) Cumulative Integrated objective value for bidirectional path topology for N=6, C=32, W=6

Figure A.16: Bidirectional path performance under uniform rising and falling traffic evolution with $LPlimit = 30\%$, $\gamma = 7$

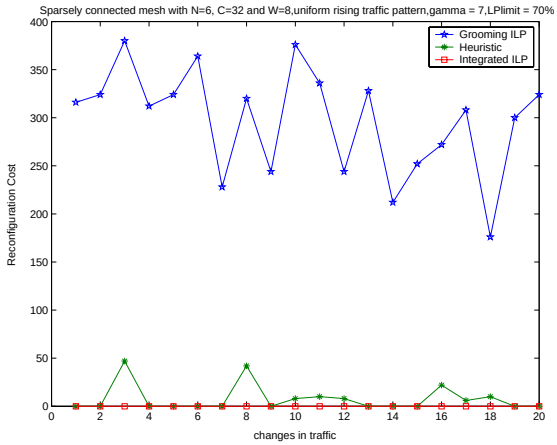

(a) Reconfiguration cost for barbell topology for N=6, C=32, W=8

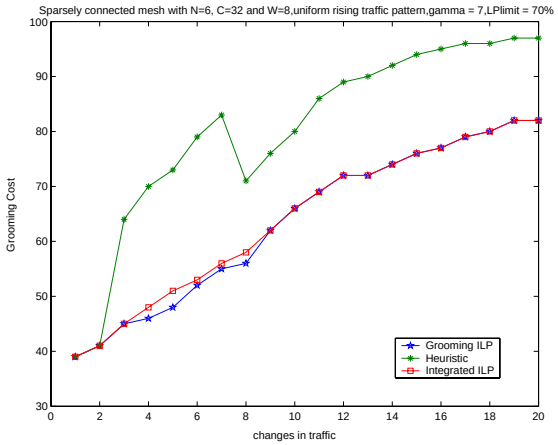

(b) Grooming cost for barbell topology for N=6, C=32, W=8

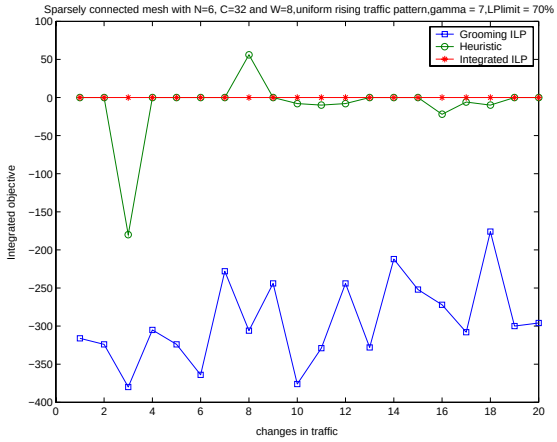

(c) Integrated objective function for barbell topology for N=6, C=32, W=8

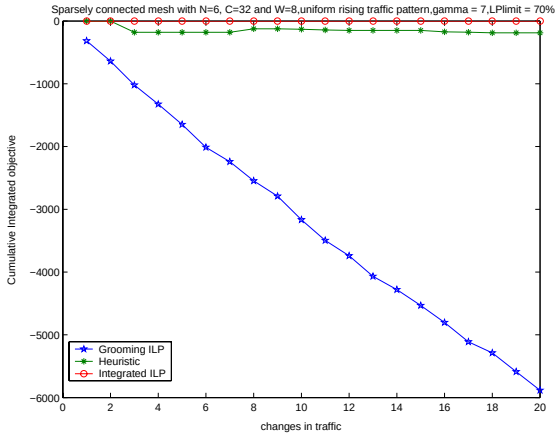

(d) Cumulative Integrated objective value for barbell topology for N=6, C=32, W=8

Figure A.17: Barbell performance under uniform rising traffic evolution with $LPlimit = 70\%$, $\gamma = 7$

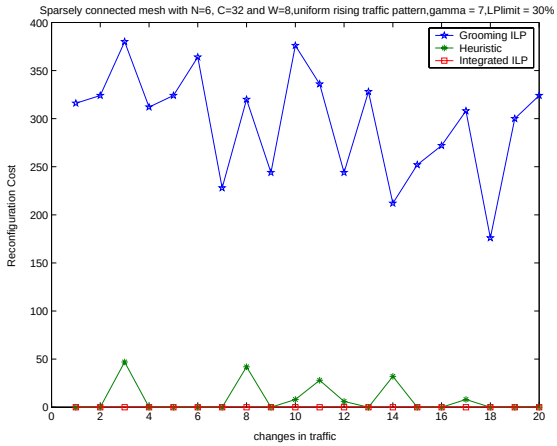

(a) Reconfiguration cost for barbell topology for N=6, C=32, W=8

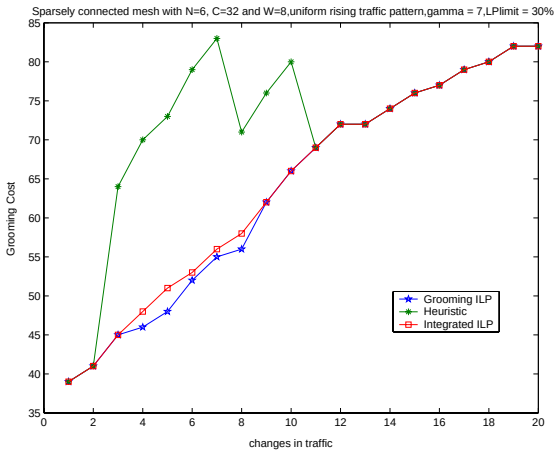

(b) Grooming cost for barbell topology for N=6, C=32, W=8

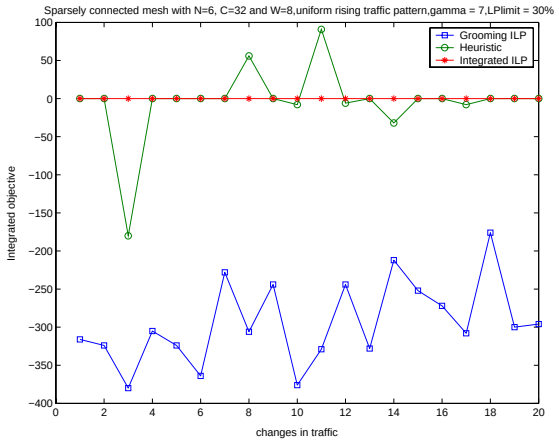

(c) Integrated objective function for barbell topology for N=6, C=32, W=8

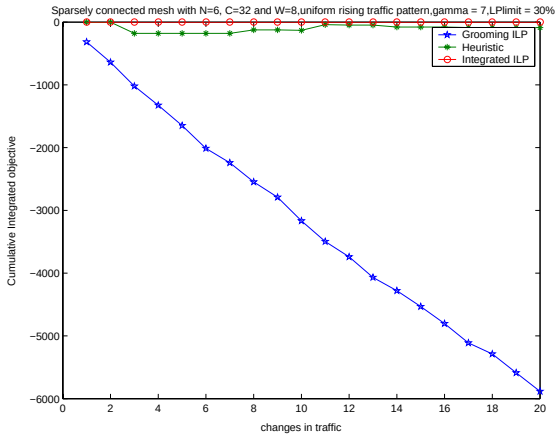

(d) Cumulative Integrated objective value for barbell topology for N=6, C=32, W=8

Figure A.18: Barbell performance under uniform rising traffic evolution with $LPlimit = 30\%$, $\gamma = 7$

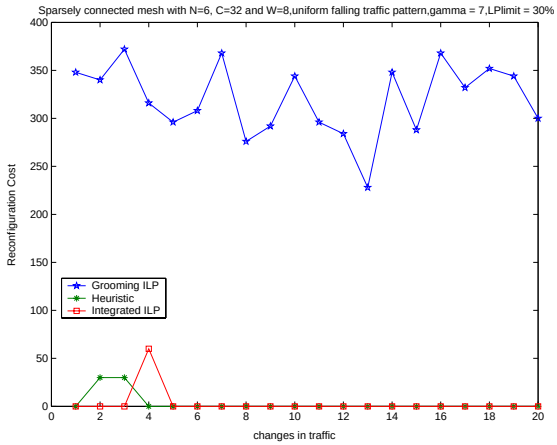

(a) Reconfiguration cost for barbell topology for N=6, C=32, W=8

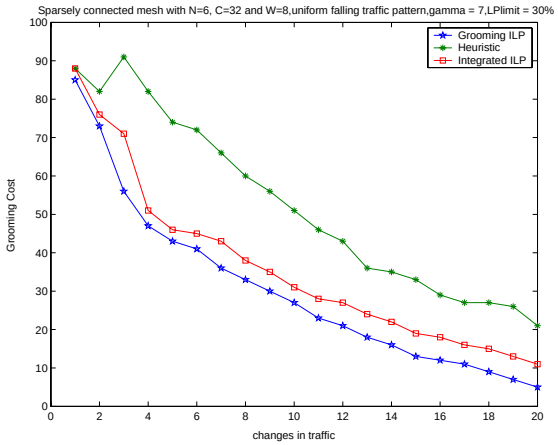

(b) Grooming cost for barbell topology for N=6, C=32, W=8

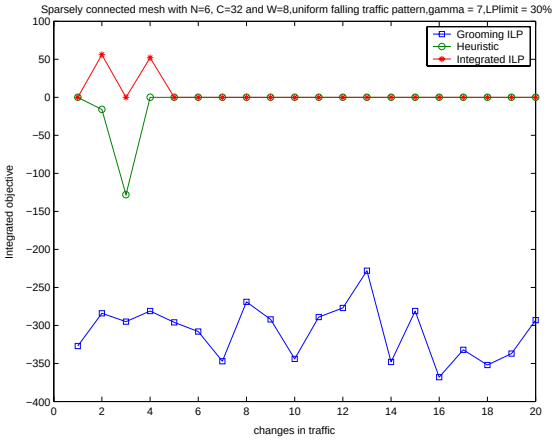

(c) Integrated objective function for barbell topology for N=6, C=32, W=8

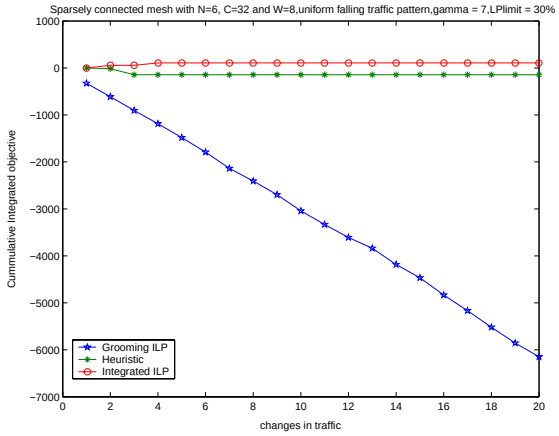

(d) Cumulative Integrated objective value for barbell topology for N=6, C=32, W=8

Figure A.19: Barbell performance under uniform falling traffic evolution with $LPlimit = 30\%$, $\gamma = 7$

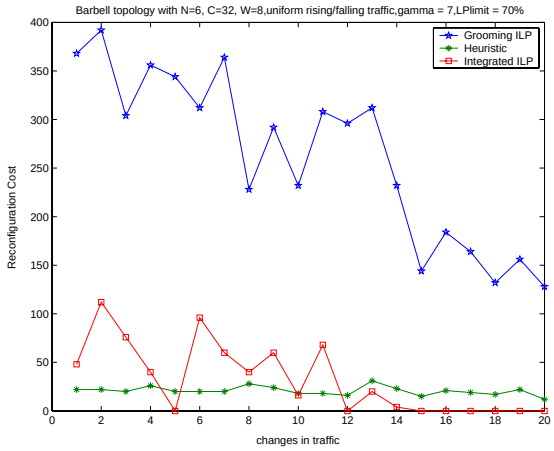

(a) Reconfiguration cost for barbell topology for N=6, C=32, W=8

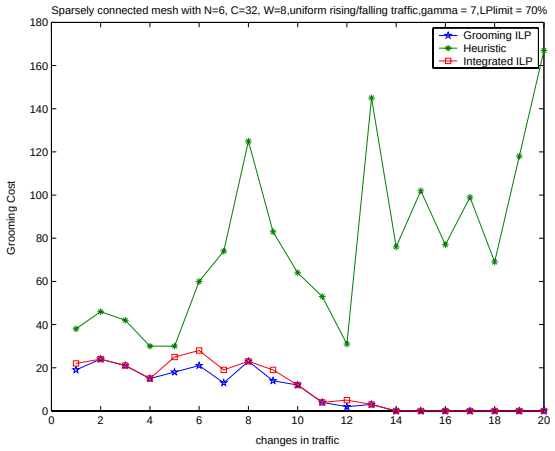

(b) Grooming cost for barbell topology for N=6, C=32, W=8

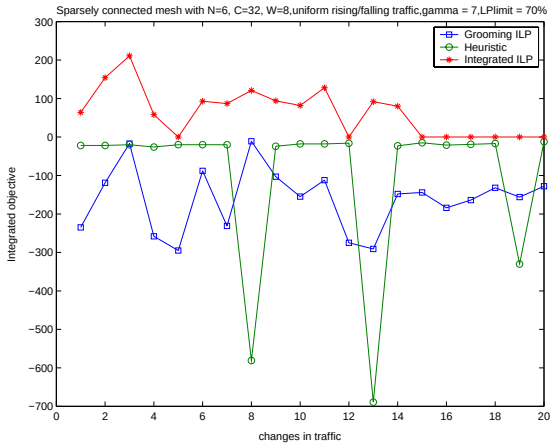

(c) Integrated objective function for barbell topology for N=6, C=32, W=8

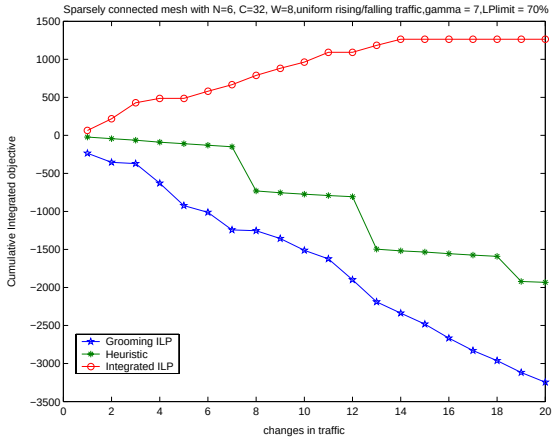

(d) Cumulative Integrated objective value for barbell topology for N=6, C=32, W=8

Figure A.20: Barbell performance under uniform rising and falling traffic evolution with $LPlimit = 70\%$, $\gamma = 7$

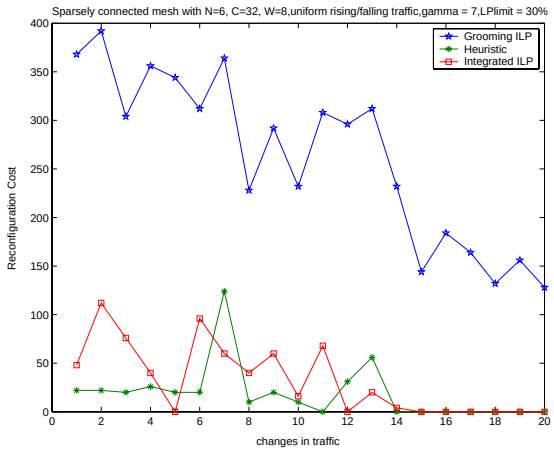

(a) Reconfiguration cost for barbell topology for N=6, C=32, W=8

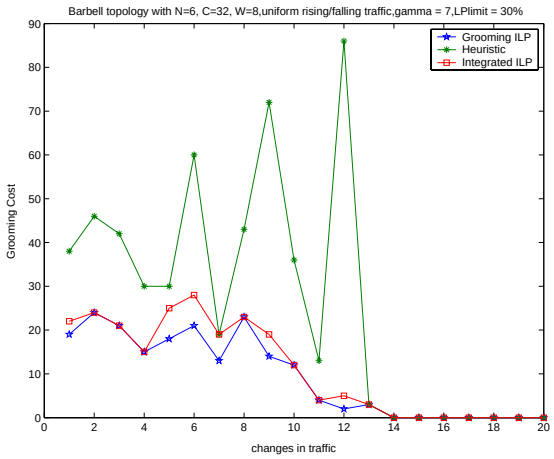

(b) Grooming cost for barbell topology for N=6, C=32, W=8

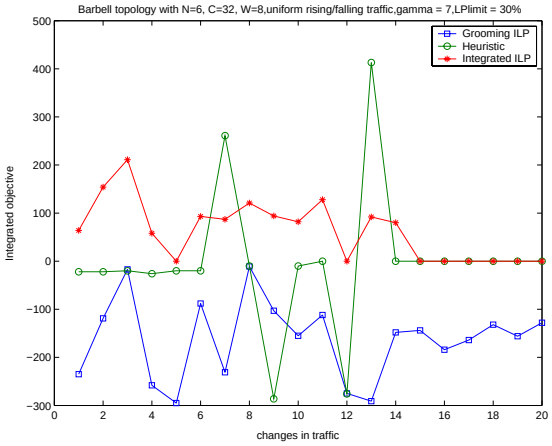

(c) Integrated objective function for barbell topology for N=6, C=32, W=8

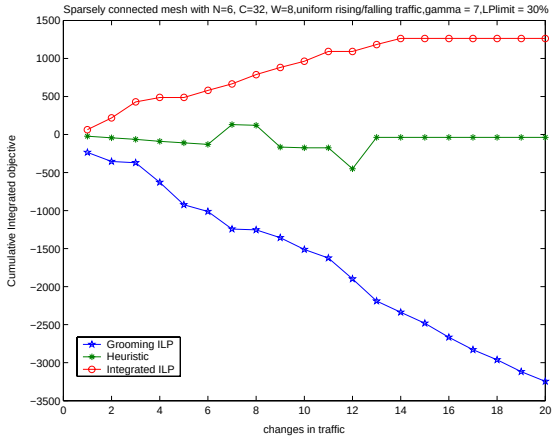

(d) Cumulative Integrated objective value for barbell topology for N=6, C=32, W=8

Figure A.21: Barbell performance under uniform rising and falling traffic evolution with $LPlimit = 30\%$, $\gamma = 7$

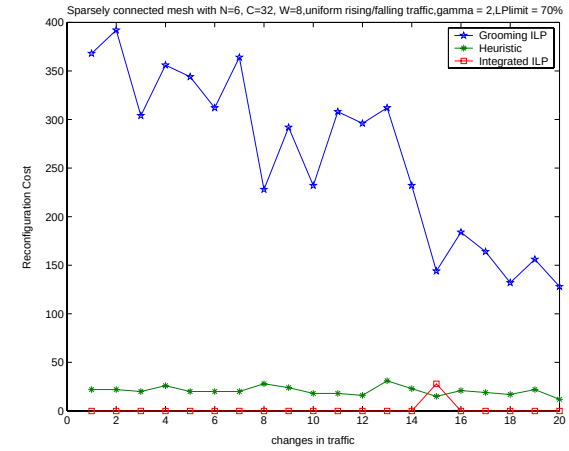

(a) Reconfiguration cost for barbell topology for N=6, C=32, W=8

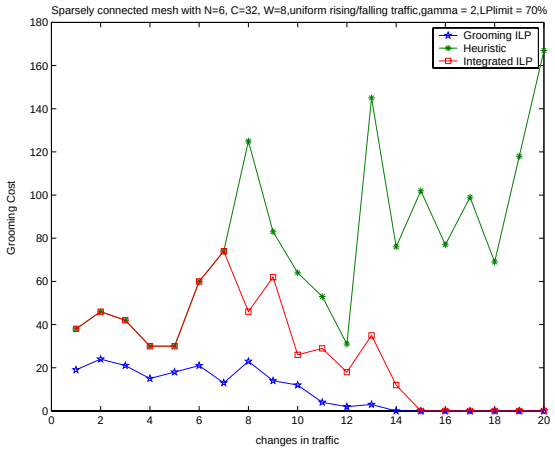

(b) Grooming cost for barbell topology for N=6, C=32, W=8

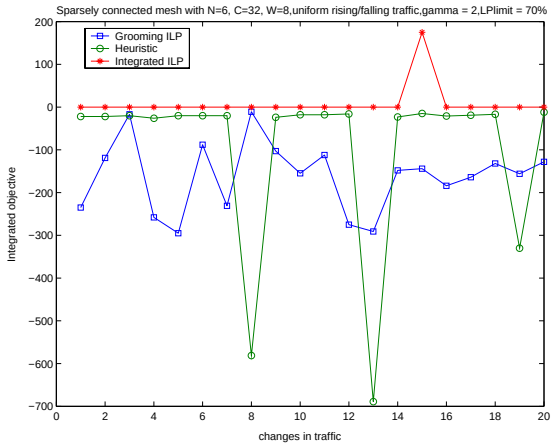

(c) Integrated objective function for barbell topology for N=6, C=32, W=8

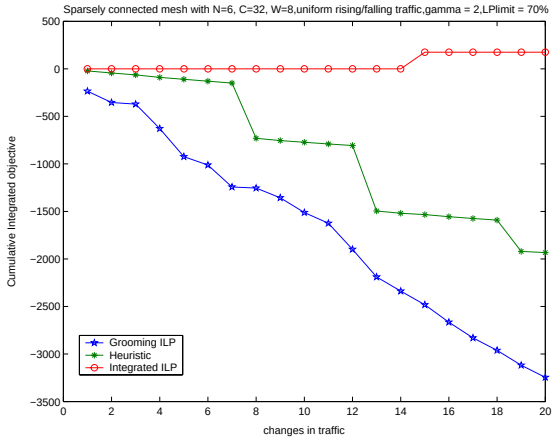

(d) Cumulative Integrated objective value for barbell topology for N=6, C=32, W=8

Figure A.22: Barbell performance under uniform rising and falling traffic evolution with $LPlimit = 70\%$, $\gamma = 2$

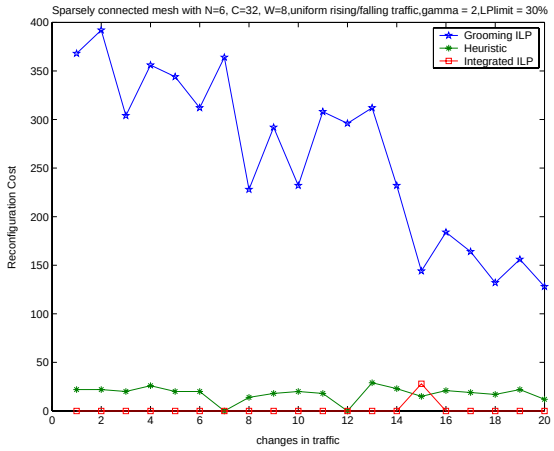

(a) Reconfiguration cost for barbell topology for N=6, C=32, W=8

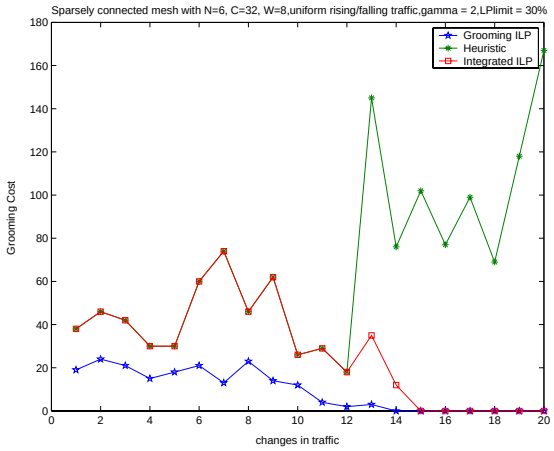

(b) Grooming cost for barbell topology for N=6, C=32, W=8

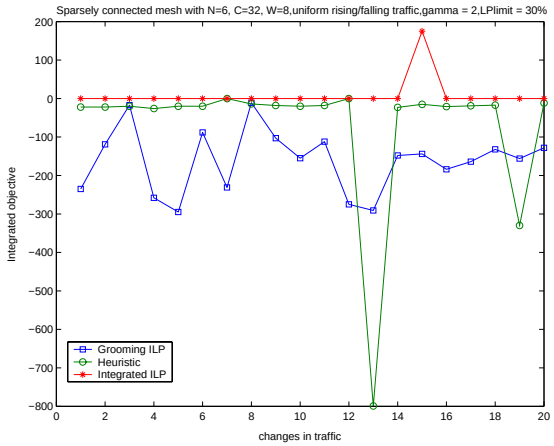

(c) Integrated objective function for barbell topology for N=6, C=32, W=8

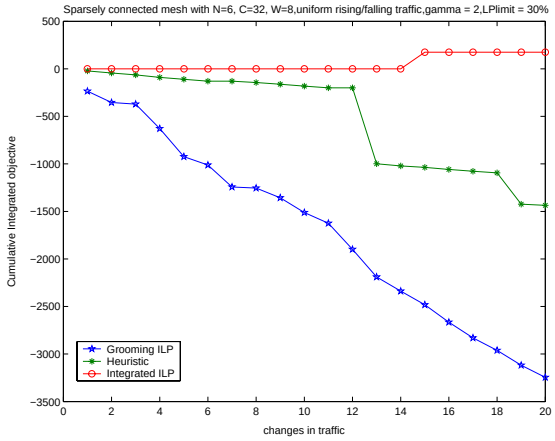

(d) Cumulative Integrated objective value for barbell topology for N=6, C=32, W=8

Figure A.23: Barbell performance under uniform rising and falling traffic evolution with $LPlimit = 30\%$, $\gamma = 2$

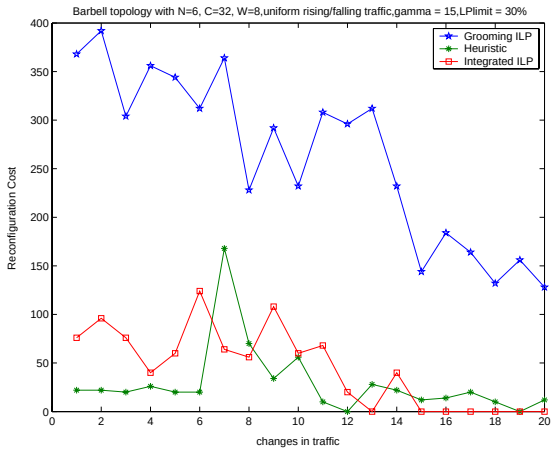

(a) Reconfiguration cost for barbell topology for N=6, C=32, W=8

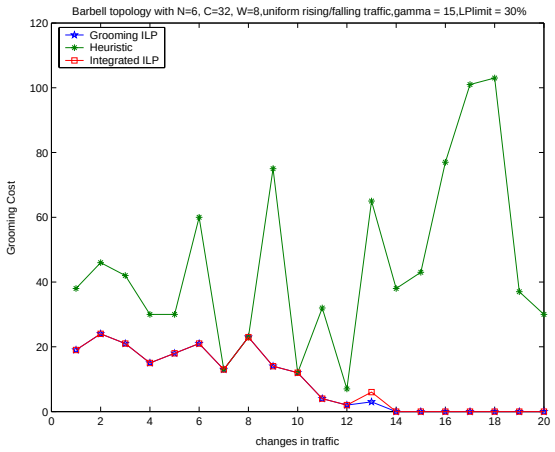

(b) Grooming cost for barbell topology for N=6, C=32, W=8

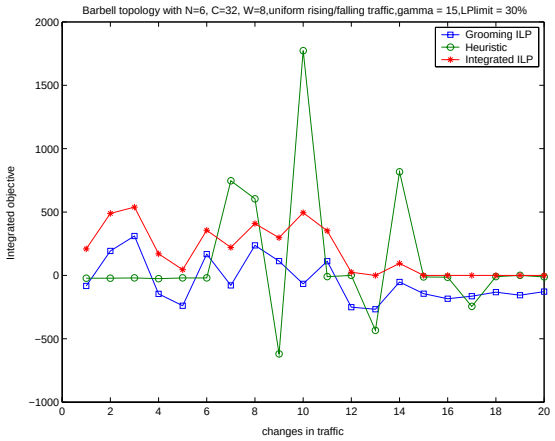

(c) Integrated objective function for barbell topology for N=6, C=32, W=8

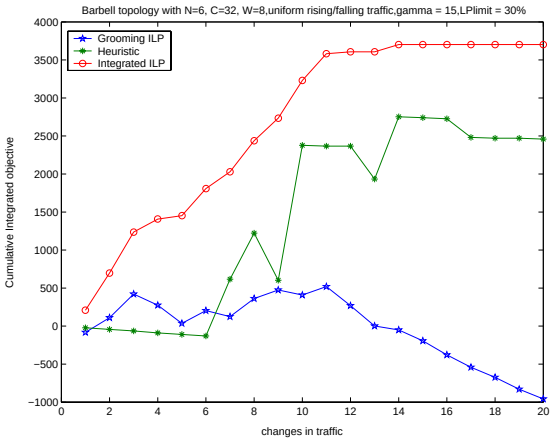

(d) Cumulative Integrated objective value for barbell topology for N=6, C=32, W=8

Figure A.24: Barbell performance under uniform rising and falling traffic evolution with $LPlimit = 30\%$, $\gamma = 15$

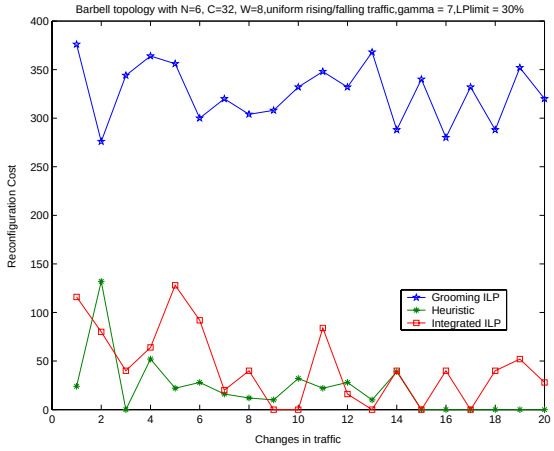

(a) Reconfiguration cost for barbell topology for N=6, C=32, W=8

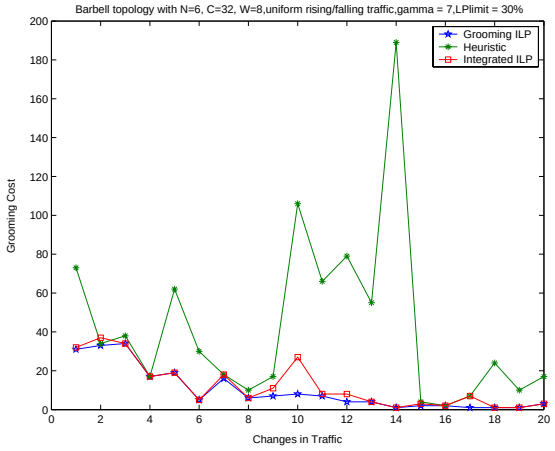

(b) Grooming cost for barbell topology for N=6, C=32, W=8

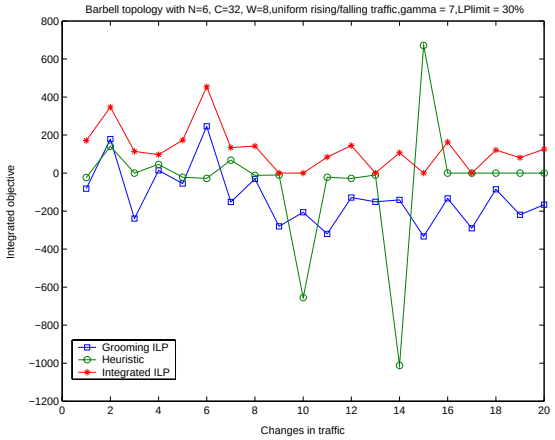

(c) Integrated objective function for barbell topology for N=6, C=32, W=8

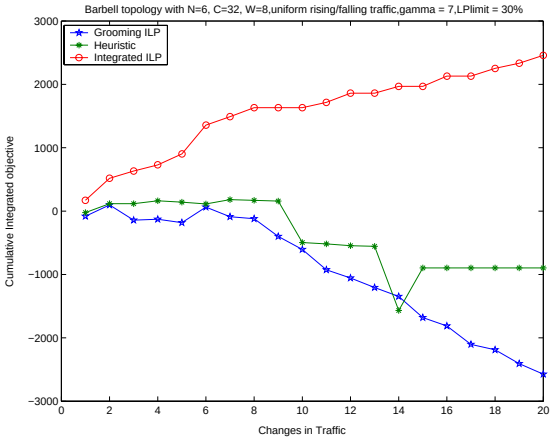

(d) Cumulative Integrated objective value for barbell topology for N=6, C=32, W=8

Figure A.25: Barbell performance under uniform rising and falling traffic evolution with $LPlimit = 30\%$, $\gamma = 7$

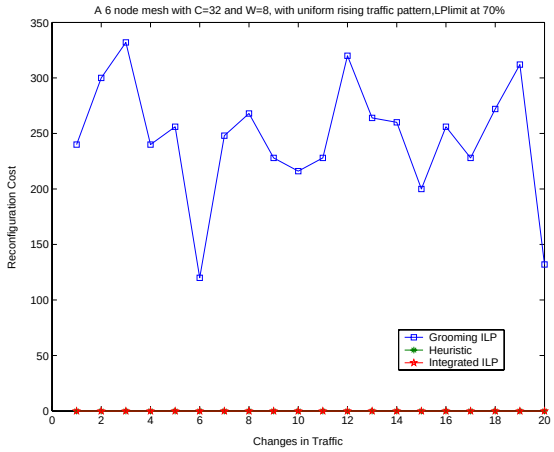

(a) Reconfiguration cost for mesh topology for N=6, C=32, W=8

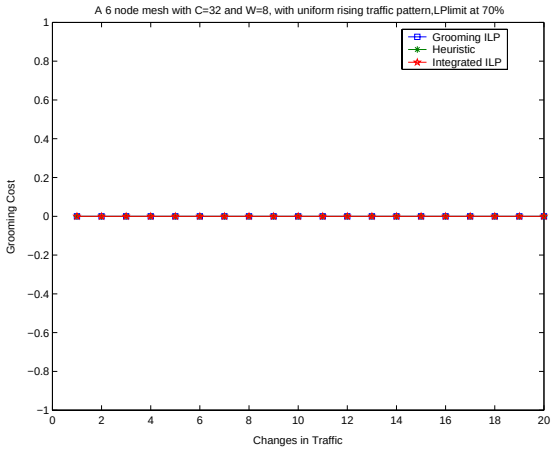

(b) Grooming cost for mesh topology for N=6, C=32, W=8

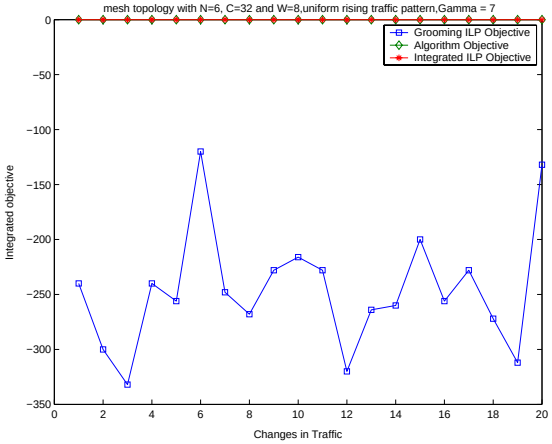

(c) Integrated objective function for mesh topology for N=6, C=32, W=8

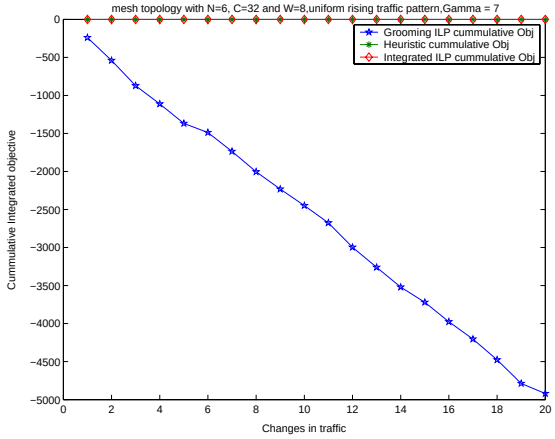

(d) Cumulative Integrated objective value for mesh topology for N=6, C=32, W=8

Figure A.26: Mesh topology performance under rising traffic evolution with $LPlimit = 70\%$, $\gamma = 7$

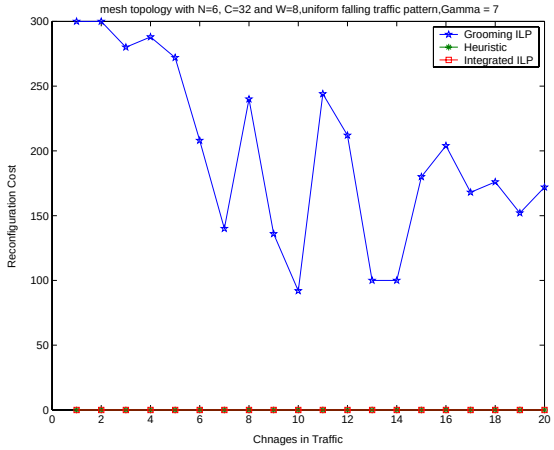

(a) Reconfiguration cost for mesh topology for N=6, C=32, W=8

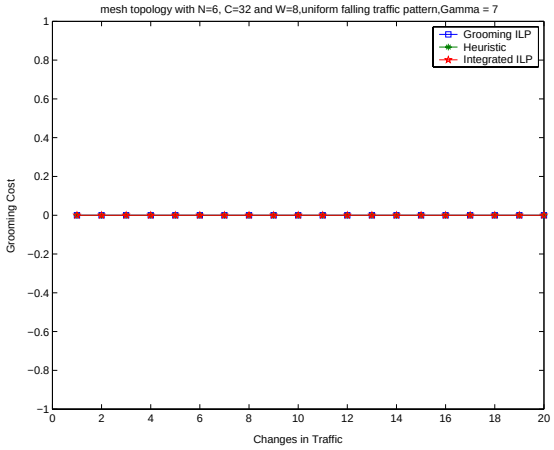

(b) Grooming cost for mesh topology for N=6, C=32, W=8

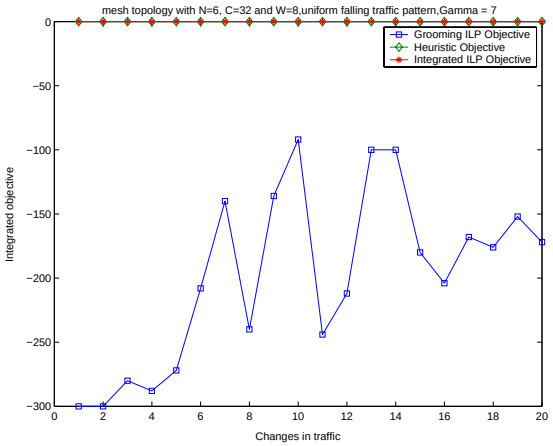

(c) Integrated objective function for mesh topology for N=6, C=32, W=8

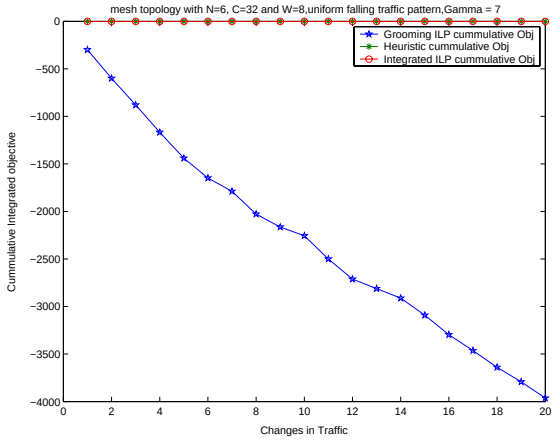

(d) Cumulative Integrated objective value for mesh topology for N=6, C=32, W=8

Figure A.27: Mesh topology performance under falling traffic evolution with $LPlimit = 70\%$, $\gamma = 7$

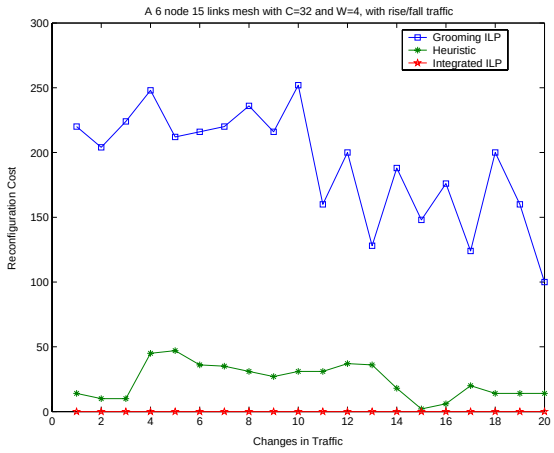

(a) Reconfiguration cost for mesh topology for N=6, C=32, W=4

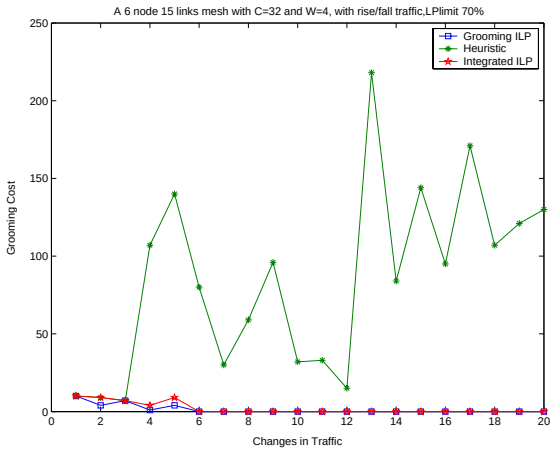

(b) Grooming cost for mesh topology for N=6, C=32, W=4

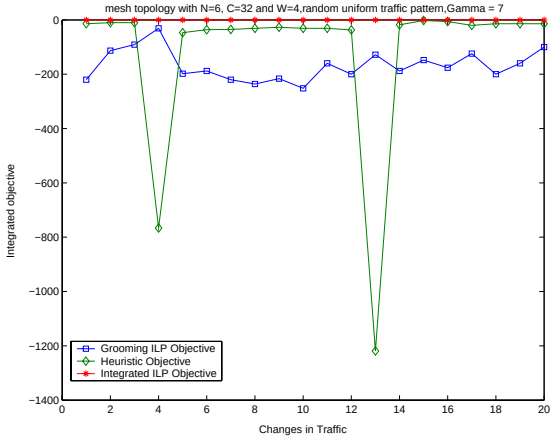

(c) Integrated objective function for mesh topology for N=6, C=32, W=4

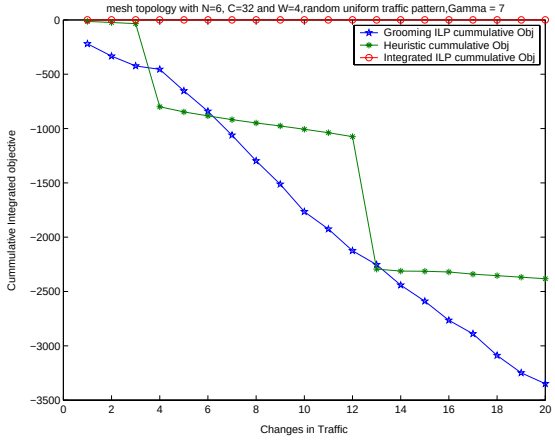

(d) Cumulative Integrated objective value for mesh topology for N=6, C=32, W=4

Figure A.28: Mesh topology performance under rising and falling traffic evolution with $LPlimit = 70\%$, $\gamma = 7$

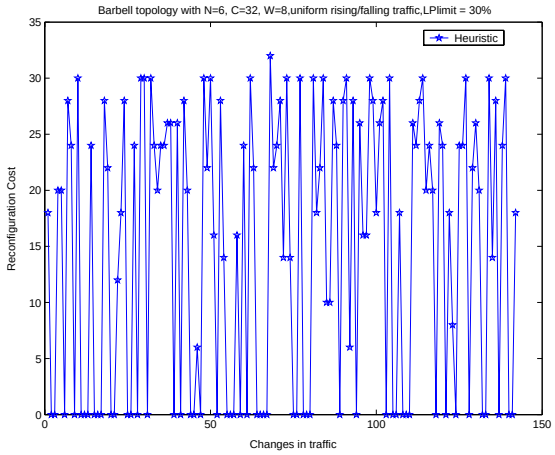

(a) Reconfiguration cost for barbell topology for N=6, C=32, W=8

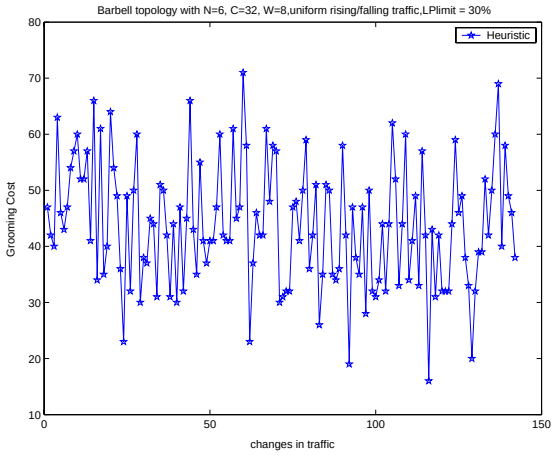

(b) Grooming cost for barbell topology for N=6, C=32, W=8

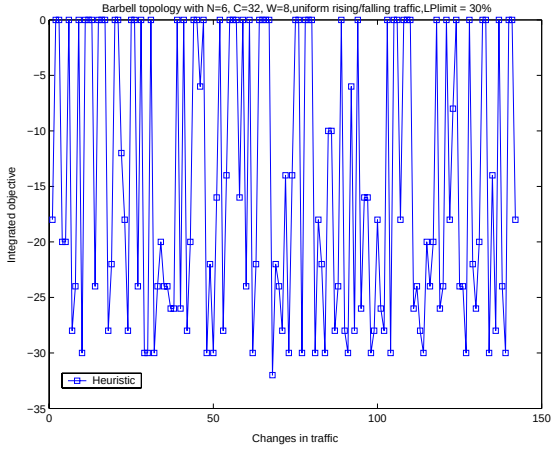

(c) Integrated objective function for barbell topology for N=6, C=32, W=8

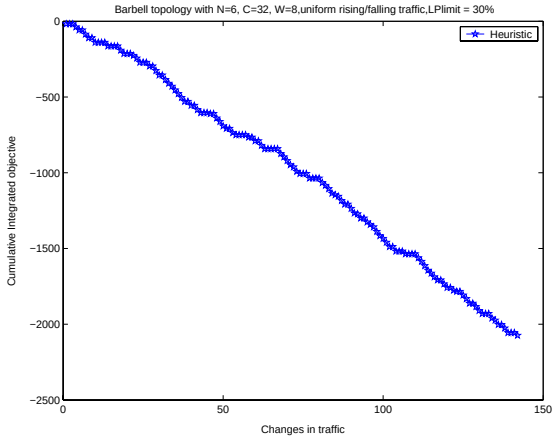

(d) Cumulative Integrated objective value for barbell topology for N=6, C=32, W=8

Figure A.29: Barbell topology performance for only the heuristic under rising and falling traffic evolution with $LPlimit = 30\%$