Policy-Driven Automated Reconfiguration for Performance Management in WDM Optical Networks *

Wojciech Golab and Raouf Boutaba
University of Waterloo
{wgolab, rboutaba}@bbcr.uwaterloo.ca

January 5, 2004

Abstract

A key feature of optical networks based on wavelength division multiplexing (WDM) technology is the ability to optimize the configuration of optical resources, i.e. wavelengths, with respect to a particular traffic demand. In the broadcast architecture, this involves the assignment of wavelengths to logical links, while in the optically switched architecture it additionally involves the routing of all-optical data paths known as lightpaths. This survey paper is concerned with the problem of automatically updating the configuration of an optical network to accommodate changes in traffic demand, which entails making a reconfiguration policy decision, selecting a new configuration, and migrating from the current to the new configuration. Existing solutions are classified according to their algorithmic properties, and compared on the basis of performance, computational cost, and flexibility. Finally, open problems and research directions are discussed.

1 Introduction

In the last decade, the growing popularity of the Internet has created an unprecedented demand for worldwide data transport. The optical fibre medium has risen to the challenge of

*This research is partially supported by the Natural Sciences and Engineering Research Council (NSERC) of Canada.

accommodating this demand thanks to its superior transmission capacity, signal attenuation, and bit error rate [17, 15]. Various innovations aiming to reduce and counteract dispersion effects have made it possible to achieve single-channel data rates of 40 Gb/s and all-optical transmission distances in excess of 2000 km [17]. Developments have also taken place in the area of wavelength division multiplexing (WDM), which involves the concurrent transmission of optical signals on different wavelengths over a common fibre. WDM systems carrying more than 100 channels at rates of 10 Gb/s are currently possible, offering aggregate capacities in excess of 1 Tb/s [12]. The combination of fast electronic interfaces and WDM technology promises to go a long way towards harvesting the enormous transmission capacity of optical fibre, which is theoretically in the range of tens of Tb/s [15].

Two types of optical network architectures have received considerable attention in literature: the passive broadcast network and the optically switched network [5, 10, 14, 21]. The passive broadcast architecture aims at simplicity and low cost. It does not use optical amplifiers, and is suitable as a distribution network that serves tens or hundreds of geographically localized nodes, for example high-end workstations or high-speed routers. Segments of fibre are combined into a shared broadcast medium using passive optical components such as star couplers. In order to avoid feedback, the network nodes and couplers must form an acyclic topology such as a bus, star, or tree. A consequence of this is that the failure of any one coupler disconnects the network.

A cost-effective way to achieve full single-hop logical connectivity in a broadcast WDM network is to equip each node with a single tunable transmitter and tunable receiver. This allows each node to communicate using multiple wavelengths, albeit not concurrently. In order to achieve high throughput despite non-negligible tuning times, it suffices to make the transmitter rapidly tunable, i.e. tunable in time comparable to packet transmission time, while using a less costly, slowly tunable receiver [19]. In that case, the transmitter is retuned between bursts of traffic destined to the same receiver, while wavelengths are assigned to the receiver on a long-term basis. Each wavelength can be shared by a set receivers, provided that the total incoming traffic demand for that set of receivers does not exceed the capacity of a wavelength. An example of such a configuration is shown in Figure 1.

Multi-hop traffic routing can be used in broadcast optical networks as an alternative to

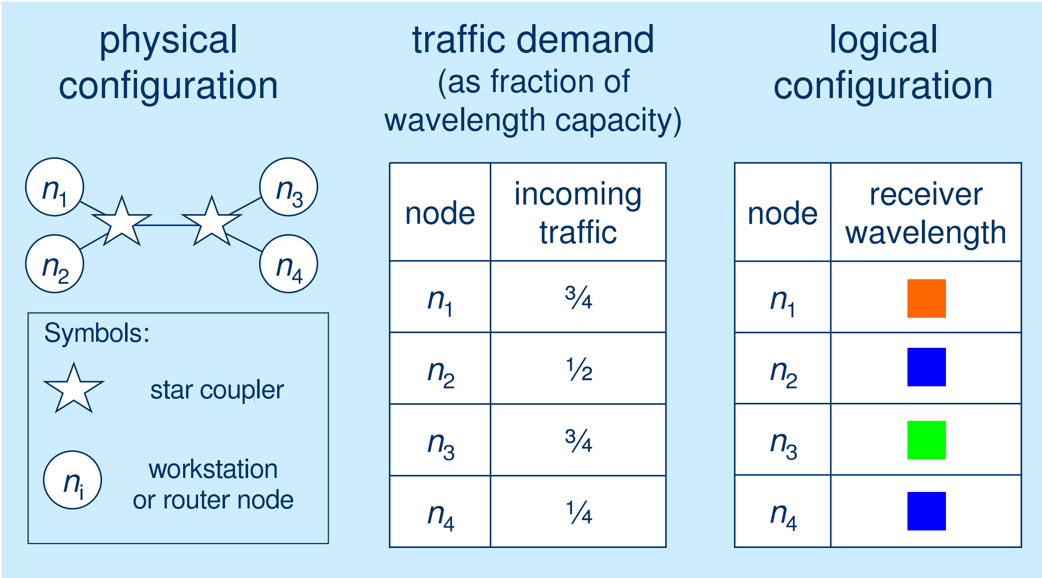

Figure 1: Example single-hop broadcast WDM network and configuration.

fast-tunable transmitters [10]. In this case, each unidirectional logical link uses a dedicated wavelength. The set of logical links is configurable through long-term assignment of transmitter and receiver wavelengths, and gives rise to a logical topology that is independent of the physical topology. Since the number of available wavelengths may not permit a fully connected logical topology, nodes must be able to not only receive but also forward traffic. A survey of topology optimization issues and solution techniques for the multi-hop broadcast architecture is presented in [9].

Compared to the broadcast architecture, the switched optical network provides transport over longer distances, with greater capacity and resilience. It relies on more sophisticated hardware, including a combination of the following: optical amplifiers, which allow all-optical transmission over long distances; optical add-drop multiplexers (OADMs), which allow a selected wavelength to be extracted from a fibre and electronically processed while optically forwarding others; and optical cross-connects (OXCs), which allow multiple wavelengths to be added, dropped, or optically switched across multiple fibres. The topology of a switched optical network is typically a ring, commonly used in metropolitan areas, or a mesh, more typical of backbone networks.

Optically switched networks are generally multi-hop in nature, where each logical link is an all-optical data path known as a lightpath, possibly passing through multiple OADMs or OXCs. A lightpath uses a dedicated wavelength in each fibre segment along its route, but wavelength sharing is possible among lightpaths that follow fibre-disjoint routes. The wavelength used by a lightpath must be continuous from source to destination, unless wavelength conversion is available at intermediate nodes. Although the latter feature is costly, sparse deployment of wavelength converters has been shown to achieve benefits comparable to full deployment [15]. Consequently, some works assume full wavelength conversion capability, or equivalently ignore the wavelength continuity constraint.

Similarly to a multi-hop broadcast network, an optically switched network has an associated logical topology that is defined by the set of established logical links, namely lightpaths. Realizing a particular logical topology entails routing lightpaths over the physical topology, and assigning wavelengths to lightpaths. Figure 2 illustrates this in a simple ring network. References [6] and [18] discuss optimization techniques that can be applied to arbitrary physical topologies given the expected traffic demand. In the event that the future traffic demand is not known or is difficult to predict, lightpaths can also be provisioned dynamically.

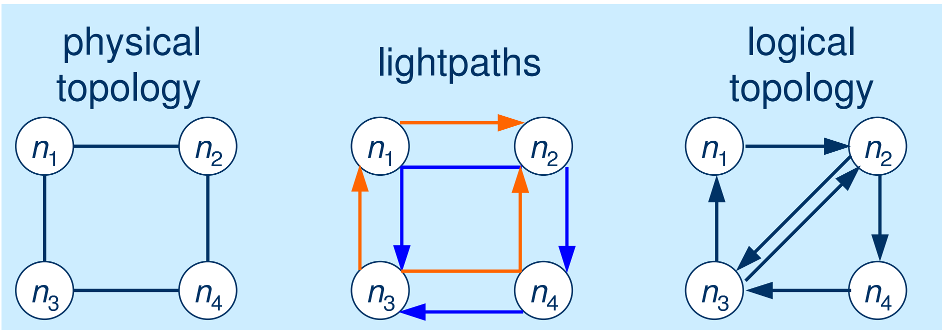

Figure 2: Example optically switched ring network and configuration. Each node is equipped with two transmitter/receiver pairs. Each fibre is bidirectional and carries up to two wavelengths. The logical topology shown has been optimized for bidirectional traffic between nodes n_2 and n_3.

Both broadcast and optically switched WDM optical networks are configurable in a manner that affects their ability to accommodate a particular traffic demand. Consider a traffic demand matrix $\mathbf{T}$ defined as

$$\mathbf{T}_{i,j} = \begin{cases} \text{directed traffic demand from node } i \text{ to node } j & \text{if } \; i \neq j \\ 0 & \text{if } \; i = j \end{cases}$$

If the demand is static, one can optimize the configuration of a multi-hop broadcast or switched optical network by minimizing the traffic-weighted average hop count. However, this can give rise to heavily loaded links as a result of packing traffic flows along shortest routes. Small deviations in the actual traffic demand from the expected demand may cause congestion. Avoiding this is especially critical in optical networks due to the transmission rates involved. A more conservative optimization strategy is to maximize the real number γ such that the traffic demand $\gamma\mathbf{T}$ can be satisfied, which is equivalent to minimizing the maximum link load [6]. This technique allows the traffic demand to scale up maximally. However, the traffic matrix can also change in ways that are not so uniform. Large bursts in the volume of traffic offered between a particular pair of nodes can occur during the transfer of large files or high-quality streaming video. On a larger time scale, daily oscillations in traffic load can occur at various phases in various parts of the network if users are distributed across multiple time zones.

The difficulty in selecting an optimal static configuration, as well as the limited ability of such a configuration to accommodate changing traffic patterns, motivates the process of reconfiguring the network on an ongoing basis. This way, the network can not only adapt to changing traffic patterns, but can also respond to unforeseen changes in the physical topology, such as equipment failures and network expansion. However, the reconfiguration procedure is a delicate one since it typically involves both the creation and teardown of logical links, and consequently the disruption of certain traffic flows. The decision to reconfigure, the new configuration, and the configuration migration strategy must all be carefully considered in order for the network to adapt efficiently and effectively.

This paper is organized into two main parts. First, reconfiguration techniques and reconfiguration policies are discussed in Section 2. Then, open problems and research directions are discussed in Section 3.

2 Reconfiguration Techniques and Policies

This part reviews proposed approaches to the optical network reconfiguration problem. Section 2.1 discusses solutions to the subproblems of new configuration selection and configuration migration. Then, section 2.2 examines reconfiguration policies.

2.1 Reconfiguration Techniques

Reconfiguration techniques can be categorized according to the nature of the algorithms used. The subproblems of selecting a new configuration and migrating to that configuration can be solved jointly or separately. Furthermore, the configuration selection subproblem lends itself to a variety of solution strategies, ranging from complex mathematical formulations to dumb searches. The organization of this section is based on these characteristics rather than the network architecture under consideration, since they determine the properties of a reconfiguration scheme such as the degree of network disruption, the optimality of the new configuration, and the computational cost.

2.1.1 Direct Approaches

A direct approach to the reconfiguration problem is characterized by the selection of a new configuration independently of the current configuration. That is to say, the configuration selection subproblem is solved separately from the configuration migration subproblem, using the same techniques as in the static configuration case. These techniques have been widely discussed in literature, and are beyond the scope of this survey. Instead, this section focuses strictly on the configuration migration subproblem.

Configuration migration for single-hop broadcast networks using rapidly tunable transmitters and slowly tunable receivers is investigated in [3]. The study explores the impact of the length of the migration phase on the performance of the network. Specifically, it considers a family of migration techniques where the relevant receivers are retuned in stages, each stage affecting a fixed-size group of receivers. On the one hand, it is desirable to make the group size as large as possible in order to shorten the migration phase and reduce the amount of time the network spends in the old, suboptimal state. On the other hand, the network tends to be more congested early in the migration phase, and retuning a receiver then is more disruptive. Consequently, there is also a reason to retune receivers more gradually by using a smaller group size. Simulation results show that a compromise between these two alternatives yields the lowest packet loss.

A configuration migration technique for multi-hop broadcast networks is explored in [11]. Because each logical link requires a dedicated wavelength in this architecture, links must be reconfigured carefully in order to avoid resource conflicts. To this end, the branch exchange operation is proposed, where the logical links between a pair of transmitters and a pair of receivers are rearranged as illustrated in Figure 3.

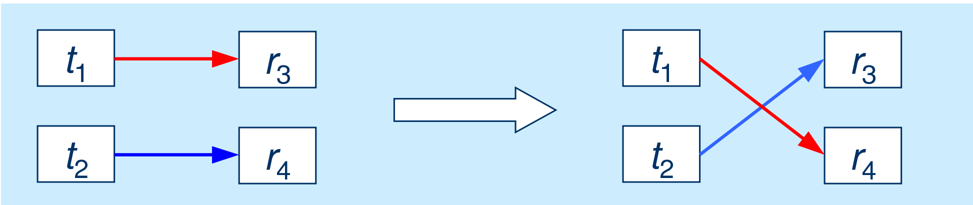

Figure 3: Example branch exchange operation in a multi-hop broadcast network among transmitters t_1, t_2 and receivers r_3, r_4. Solid arrows indicate unidirectional logical links, each using a dedicated wavelength.

It can be shown that a sequence of branch exchanges can transform an initial logical topology to any target logical topology with the same number of links. Naturally, the problem of finding the shortest such sequence is important in minimizing network disruption during the configuration migration phase. Two key properties are identified in [11] that simplify this problem. First, the shortest sequence only contains branch exchanges that decrease the number of links that remain to be established in order to reach the target configuration. For comparison, in the general case a branch exchange can also preserve or increase this number. Second, two branch exchanges that do not operate on the same logical link can be applied in any order without affecting the final result. These properties can be exploited in constructing heuristics that produce approximate solutions to the shortest branch exchange sequence problem. A simulation study of three such algorithms is presented in [11].

A direct technique for reconfiguring optically switched networks is discussed in [13], using the Incremental Reconfiguration Migration heuristic. This algorithm maintains lists of lightpaths that remain to be created and that remain to be deleted. Each iteration of the algorithm creates as many lightpaths as possible from the first list, and then tears down the least loaded lightpath from the second list. In order to minimize the impact of the latter step, only those lightpaths are considered whose deletion does not disconnect the logical topology. Thus, the heuristic reduces network disruption by postponing lightpath deletion and prioritizing lightpath creation. This is possible in an optically switched network because lightpaths occupy varying amounts of resources, depending on their length, and it is not always necessary to perform one lightpath deletion for every lightpath creation.

2.1.2 Partial Reconfiguration Approaches

This section discusses reconfiguration approaches that use heuristics to improve the configuration of an optical network, sacrificing the optimality of the new configuration in order to reduce the degree of network disruption. Specifically, these algorithms focus on a smaller number of more beneficial configuration changes rather than considering reconfiguration of the entire network.

A partial reconfiguration heuristic for single-hop broadcast networks is described in [1]. As its name suggests, the Most and Least Loaded Channel Balance Algorithm attempts to reduce the maximum channel (i.e. wavelength) load by swapping the wavelength assignments of one receiver assigned to the most loaded channel and another assigned to the least loaded channel. The candidate receiver pair that is actually reconfigured is chosen on the basis of reduction in maximum channel load. In contrast to a direct reconfiguration approach, this procedure redistributes load within the network while disrupting a minimal number of logical links.

A partial reconfiguration approach suitable to optically switched networks is presented in [20]. The proposed Path-Add heuristic aims to improve the configuration of the network by reducing the average traffic-weighted hop count. The algorithm considers K node pairs where the associated traffic flows have the greatest traffic-weighted hop count, and considers joining each such pair with a single lightpath. If a lightpath cannot be established, then up to L subsets of existing lightpaths are considered for deletion. An example of this is shown in Figure 4. The

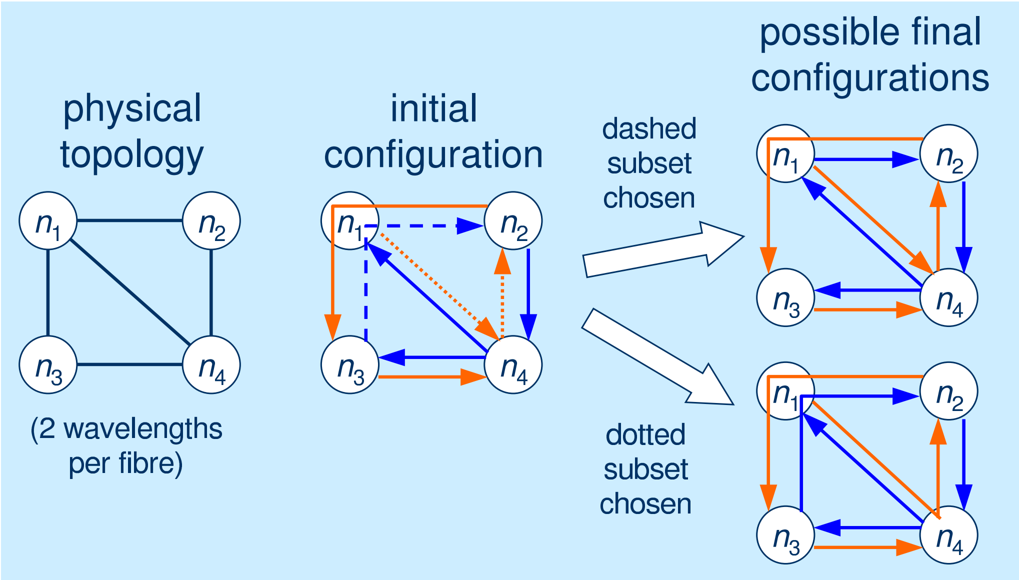

Figure 4: Example subset selection in Path-Add heuristic with $K = 1$ and $L = 2$. The dashed and dotted subsets of lightpaths in the initial configuration are considered for deletion in order to establish a lightpath from n_1 to n_2.

number of reconfiguration changes allowed here is bounded by another parameter, D. Of all the options explored, the algorithm selects the one that gives rise to a logical topology with the smallest traffic-weighted average hop count. The algorithm allows the degree of network disruption to be controlled via the parameter D, which also affects the computational cost. Similarly, computation time can be traded off with the optimality of the new configuration by varying K and L.

While the Path-Add heuristic is designed to work with optically switched networks having arbitrary physical topologies, [7] presents a technique specific to the popular ring topology. Rather than completely deleting existing lightpaths in order to liberate the receivers, transmitters, and wavelengths required to create a new lightpath, the proposed approach applies

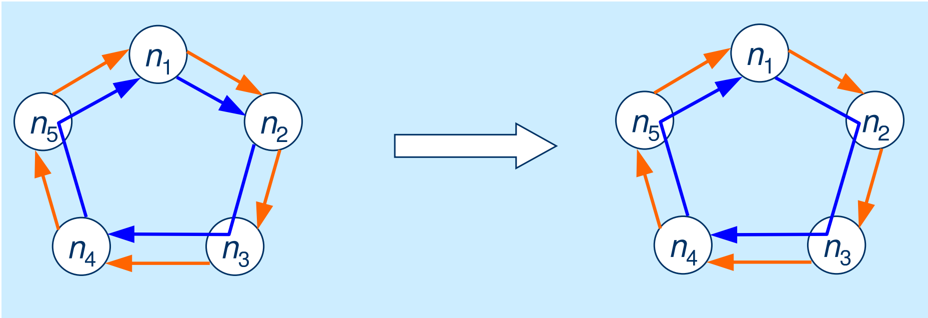

Figure 5: Example merge operation to establish a lightpath from node n_1 to node n_4 in a unidirectional ring. The inverse operation can be used to establish a lightpath from n_2 to n_4.

lightpath merge and split operations. A lightpath merge operation takes two lightpaths with compatible directions and endpoints, and cross-connects them to form a single lightpath. A split operation is equivalent to replacing a cross-connection with an add/drop at a node that has an available receiver/transmitter pair. Examples of these operations are presented in Figure 5. A sequence of such operations is decided using the Merge-Split Reconfiguration (MSR) Algorithm, which attempts to create lightpaths between high-demand node pairs. Split operations are used at the source and destination, while merge operations are applied on the path between. The algorithm considers various wavelengths for the new lightpath, choosing one on the basis of reduction in maximum lightpath load.

The performance of the MSR algorithm is compared in [7] against a direct approach that uses a logical topology design heuristic. Tests done on unidirectional and bidirectional ring topologies show that the maximum lightpath loads obtained using the two approaches are within 15% of each other, while MSR performs eight to ten times fewer configuration changes.

Another partial reconfiguration approach is introduced in [8], that allows an optically switched network to evolve gracefully in response to a gradually changing traffic demand. This is achieved by not only creating lightpaths to relieve congestion, but also proactively deleting underutilized lightpaths. The latter procedure liberates network resources in a less disruptive manner than the lightpath deletion in the Path-Add heuristic. The proposed Heuristic Adaptation Algorithm is executed at regular time intervals and attempts to create or delete a single lightpath at each execution. First, lightpath creation is considered between pairs of nodes that have positive traffic demand but are disconnected in the logical topology. If no such node pair exists, the algorithm identifies a congested link, and attempts to create a lightpath between the endpoints of the greatest contributing multi-hop traffic flow. In the absence of congested links, the algorithm attempts to delete an underutilized lightpath without disconnecting the logical topology. Congested and underutilized links are defined in terms of high and low load thresholds (watermarks), W_H and W_L, which are specified as parameters to the algorithm. In fact, these thresholds can be thought of as specifying part of the reconfiguration policy, the other part being the rate at which the algorithm is invoked.

A simulation study based on a lightly loaded mesh network is presented in [8], showing that the Heuristic Adaptation Algorithm maintains lightpath loads between W_L and W_H with high probability. Traffic loss can still occur in the event that an underutilized lightpath is deleted and the corresponding traffic flows cannot be supported on alternate routes due to congestion. In that case, the next iteration is likely to address this congestion condition by creating an additional lightpath, which will help to restore the disrupted traffic flows. Such oscillations can be beneficial in the long term by rearranging lightpaths into a more optimal configuration, in a similar manner as done by the Path-Add heuristic.

2.1.3 Local Search Approaches

This section discusses a family of techniques where a set of neighbouring configurations is exhaustively explored, and the most optimal of these configurations is adopted. The search neighbourhood is defined in terms of reachability from the initial configuration via a single application of some simple reconfiguration operation. The initial configuration itself can be included in the neighbourhood, in which case the procedure is analogous to the steepest descent optimization technique.

The branch exchange operation, introduced in Section 2.1.1 in the context of multi-hop broadcast networks, serves as a good basis for a local search as it only disrupts a small portion of the network. This operation can be extended to optically switched networks, where rerouting of lightpaths is used rather than retuning of transmitters and receivers. It can also be generalized in terms of the number of links exchanged. Whereas the 2-branch exchange has been discussed previously, [16] considers a branch exchange operation the acts on groups of three links. Namely, a 3-branch exchange on links from a_1 to a_2, from b_1 to b_2, and from c_1 to c_2, results in links from a_1 to b_2, from b_1 to c_2, and from c_1 to a_2. An example of this is shown in Figure 6. An

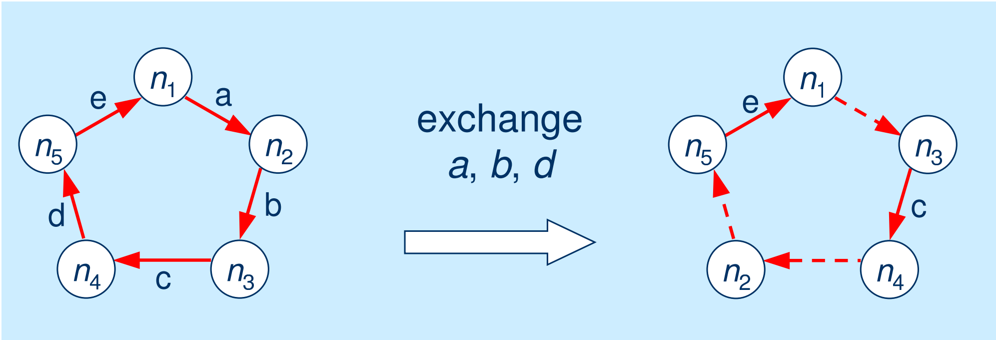

Figure 6: Example 3-branch exchange operation in a unidirectional ring. The operation changes the order of nodes in the logical topology.

advantage of the 3-branch exchange is that it never disconnects a ring network, whereas this can occur with the 2-branch exchange. In general, a k-branch exchange operation on a network with L logical links gives rise to a search neighbourhood of size $\binom{L}{k}$. However, the ability to perform a particular branch exchange operation in an optically switched network is conditional upon the availability of wavelength resources on affected routes.

The performance of a local search based on the 3-branch exchange in ring networks is evaluated in [16]. Simulations on ten-node unidirectional rings with one or two receiver/transmitter pairs per node show that ten iterations of the local search after each change in traffic demand yield a reduction in maximum lightpath load within 10% of what can be achieved using a direct approach.

2.1.4 Comparison

Direct reconfiguration approaches yield configurations that are closest to optimal. However, because the underlying optimization problems are NP-hard, approximate problem formulations are used based on various simplifications such as ignoring the wavelength continuity constraint, or solving subproblems separately rather than jointly. Solving these simplified problems optimally can still be very costly, necessitating the use of heuristics for networks of nontrivial size. In comparison, partial reconfiguration and local search approaches, both of which are also based on heuristics, achieve slightly worse results. The solutions discussed in this paper typically achieve objective function values 5-15% worse than what is possible with direct approaches. This is because partial reconfiguration approaches purposely sacrifice optimality, while local searches do not always reach an optimal configuration, either because of slow convergence or convergence to a locally but not globally optimal configuration.

The computational cost of various approaches is generally related to the optimality of the solution. Direct approaches that solve a mixed integer linear program (MILP) formulation are the most costly, with running times growing rapidly as the size of the network increases. Direct approaches that use heuristics, as well as partial reconfiguration approaches, have well-behaved polynomial running times. In particular, partial reconfiguration approaches can be quite fast when parameters are used to constrain the number of permitted configuration changes. The cost of local search approaches is determined by the size of the search neighbourhood, which is proportional to the cube of the number of lightpaths in the case of the 3-branch exchange technique. The number of lightpaths, in turn, is at most the number of nodes times the number of receiver/transmitter pairs per node. Local searches are generally slower than partial reconfiguration heuristics because a larger number of target configurations is considered.

The degree of network disruption incurred is highest with direct approaches since they ignore the current configuration of the network. In contrast, local search approaches produce minimal network disruption at each step by using families of very simple reconfiguration operations, although the number of steps needed must also be considered. However, to date, such families have only been defined for broadcast and optically switched ring networks, and not for optically switched mesh networks. The behaviour of partial reconfiguration approaches with respect to network disruption varies depending on the algorithm and the choice of input parameters, and is potentially much less than in direct approaches since the old configuration is taken into account. In addition, algorithms such as Merge-Split and the threshold-based Heuristic Adaptation make a special effort to avoid the removal of high-traffic data paths.

Partial reconfiguration approaches are the most flexible because the number of configuration changes performed can typically be controlled through parameters; the Path-Add heuristic is one example. A more coarse form of control can be achieved with local search approaches by varying the number of iterations performed at each reconfiguration step. Direct approaches are the least flexible due to the lack of relevant parameters, and the fact that multiple iterations are not beneficial.

The characteristics of the three categories of optical network reconfiguration techniques discussed in this paper are summarized in Table 1.

	Approach Type		
	Direct	Partial Reconfiguration	Local Search
optimality	*best*	*good*, susceptible to continual approximation	*good*, susceptible to local extrema
computational complexity	*intractable* (MILP) to *polytime* (heuristics)	*polytime*, can be parameterized	*polytime*, determined by neighbourhood size
network disruption	*highest*	potentially *low*, can be parameterized	*low*, by choice of reconfiguration operation
flexibility	*least*	*greatest*, due to appropriate parameters	*moderate*, multiple iterations possible
applicable architectures	*all*	*all*	*limited* to WDM rings and broadcast networks

Table 1: Comparison of three categories of reconfiguration approaches.

2.2 Reconfiguration Policies

Whereas reconfiguration algorithms are concerned with the task of updating the configuration of a network, reconfiguration policies are concerned with the process of deciding when such algorithms should be executed. Most works focus on the former, resorting to rudimentary reconfiguration policies in the simulation environment. This section begins with an overview of such rudimentary techniques, and goes on to discuss more sophisticated solutions that attempt to address the long-term effects of policy decisions.

2.2.1 Primitive Approaches

Of those works that do not ignore the topic of reconfiguration policies, most indicate one of three types of simple approaches. The simplest of all techniques is to execute the reconfiguration algorithm at regular intervals of time [1, 16, 4]. In that case, the interval length must be chosen to reflect the rate of change in the pattern of traffic with respect to the degree of adaptation achieved at each execution of the algorithm. Other approaches take the current traffic demand into consideration. Reconfiguration can be performed at every traffic change [20], or only when an important change is detected, for example on the basis of load thresholds [8].

2.2.2 Sophisticated Approaches

The formulation of a reconfiguration policy as a Markov decision process is proposed in [2]. This mathematical formulation is the first step in developing a smarter policy that considers the long-term benefits and costs of reconfiguration.

The state of the network in the formulation is represented by a tuple $(\mathcal{R}, \mathbf{T})$, where $\mathcal{R}$ is the current network configuration and $\mathbf{T}$ is the current traffic demand. Assuming that the next state depends only on the current state, the corresponding process embedded at times when $\mathbf{T}$ changes is a discrete-time Markov process. At each state transition, i.e. when $\mathbf{T}$ changes to $\mathbf{T}'$, a decision is made based entirely on $(\mathcal{R}, \mathbf{T})$ to either migrate from $\mathcal{R}$ to a new configuration $\mathcal{R}'$ (which is unique as a result of fixing the parameters of the reconfiguration algorithm) or to maintain $\mathcal{R}' = \mathcal{R}$. Associated with the next state are an immediate reward and a reconfiguration cost. The immediate reward is modelled as $\alpha[\phi(\mathcal{R}', \mathbf{T}')]$, where α is a nonincreasing function and $\phi(\mathcal{R}', \mathbf{T}')$ is the degree of load balancing achieved by $\mathcal{R}'$ under traffic $\mathbf{T}'$, defined as (maximum link load $-$ average link load) / average link load. The reconfiguration cost is $\beta[\mathcal{D}(\mathcal{R}, \mathcal{R}')]$, where β is a nondecreasing function and $\mathcal{D}(\mathcal{R}, \mathcal{R}')$ is the number of configuration changes needed to migrate from $\mathcal{R}$ to $\mathcal{R}'$.

Given particular cost and reward functions, as well as the set of state transition probabilities, one can define an optimal policy as one that maximizes the expected (long-term) reward, defined as

$$\lim_{k\to\infty} \frac{1}{k} E\left\{ \sum_{l=1}^{k} \alpha[\phi(\mathcal{R}^{(l)}, \mathbf{T}^{(l)}] - \beta[\mathcal{D}(\mathcal{R}^{(l-1)}, \mathcal{R}^{(l)})] \right\}$$

where $\left(\mathcal{R}^{(k)}, \mathbf{T}^{(k)}\right)$ is the state after the k'th transition. In general, the optimal policy cannot be computed exactly and must be approximated. The assumption that the next state is determined solely by the current state simplifies this task, allowing solutions to be obtained using well-known iterative techniques. However, due to high computational cost, these techniques are useful only if the state space is finite and fairly small, which motivates an alternate state representation in terms of the degree of load balancing ϕ and the distance D to the new configuration determined by the reconfiguration algorithm, i.e. as a tuple $(\phi, \mathcal{D})$. Techniques are proposed in [2] of discretizing ϕ, and estimating the state transition probabilities based on a near-neighbour traffic model where ϕ makes independent random transitions to the next-highest, next-lowest, or same discrete state, with fixed probabilities.

Another formulation of the reconfiguration policy problem is presented in [22]. As in the first example, states are represented as tuples $(\mathcal{R}, \mathbf{T})$. A discrete process is formed with an immediate reward related to the change in weighted average hop count, and a reconfiguration cost equal to the number of lightpaths added or deleted. At each state transition, a decision is made to either perform one of several possible reconfigurations, or to remain in the current configuration. However, rather than precomputing an optimal policy, the optimal decision in a given state is computed online using traffic prediction data. Information about the future traffic pattern makes it possible to evaluate the rewards and costs incurred by all possible combinations of policy decisions that can be made during a fixed period of time known as the prediction horizon. The optimal decision can then be selected on the basis of this approximate long-term reward.

2.2.3 Comparison

Intuitively, sophisticated approaches to the reconfiguration policy optimization problem promise to yield better-performing policies than primitive approaches, since they attempt to maximize the expected future reward of policy decisions rather than merely maximizing the immediate reward. A comparison is presented in [2] of optimal policies obtained using the Markov decision process formulation, against much simpler policies based on two thresholds. In the latter policies, reconfiguration is performed after a change in the traffic demand whenever the alternative is to enter a state $(\phi, \mathcal{D})$ such that $\phi > \phi_{max}$ or $\mathcal{D} \leq \mathcal{D}_{max}$. Under certain combinations of the thresholds ϕ_{max} and $\mathcal{D}_{max}$, this two-threshold policy achieves a greater long-term reward than an optimal policy obtained using the Markovian model. This is possible since the latter model is approximate as a result of using a simplified state space. In comparison, the traffic prediction approach uses a much simpler computation. However, its performance is dependent on the ability to predict the future traffic demand.

The ongoing cost associated with making policy decisions is smallest in primitive approaches and in the Markov decision process approach, where the policy is precomputed. The techniques given in the literature for forming the initial policy are trial-and-error in the primitive approaches, and policy iteration in the Markovian approach, both of which allow solution optimality to be traded off with computational cost. The traffic prediction approach incurs a greater ongoing cost as it entails evaluating the benefits of many possible sequences of reconfiguration decisions, in addition to performing the prediction itself. The computational cost can be controlled by varying the traffic prediction horizon as well as the set of possible reconfigurations considered in each policy decision (e.g. by restricting the number of configuration changes allowed). Because the policy is formed dynamically in the traffic prediction approach, there is no cost associated with policy precomputation.

The robustness of a reconfiguration policy can be considered in terms of the assumptions made concerning the traffic demand during policy computation. No such assumptions are made in primitive approaches, but at the same time this precludes the use of most optimization techniques and requires that parameters such as thresholds be chosen by hand. The Markovian formulation assumes that the traffic demand varies according to a near-neighbour model, which allows the use of well-known iterative techniques. This assumption is generally not valid over time scales of one day or more, where patterns such as oscillations or long-term growth trends occur in the traffic demand. In contrast, the traffic prediction approach relies on such trends in order to accurately predict traffic demand for some period of time in the future.

3 Open Problems and Research Directions

The set of works surveyed in this paper exhibits a range of approaches to the problem of automated optical network reconfiguration. Most studies are concerned with algorithms that select a target configuration. A variety of solutions has been proposed for both broadcast and optically switched network architectures. Most of these can be categorized according to their algorithmic character either as direct, partial reconfiguration, or local search approaches. Collectively, reconfiguration algorithms exhibit a tradeoff between the optimality of the new network configuration and the degree of network disruption incurred during the configuration migration phase, both of which affect performance metrics such as packet loss and delay. Reconfiguration policies and configuration migration strategies, in comparison, have received much less attention in literature, despite their importance as elements of a complete performance management solution. Many studies either ignore these subproblems altogether, or use very rudimentary solutions.

Future contributions are expected in all major areas of the automated reconfiguration problem. New target configuration selection algorithms that achieve more optimal configurations with less network disruption are one possibility, along with analysis of their properties. Further work can be done especially in the context of optically switched mesh networks, where the complexity of the topology allows the greatest flexibility in solutions. Two existing ideas that may lead to such developments are performing reconfiguration using minimally disruptive operations akin to lightpath merging and splitting [7], and maintaining spare resources to facilitate configuration migration, as in [8].

Less disruptive configuration migration techniques are another possibility for future work. For example, the techniques of ordering and grouping configuration changes to minimize network disruption, which have been introduced in [13, 3], can be explored further. Whereas the idea of using link load to prioritize configuration changes has already been proposed in [13], the delay associated with control messages during lightpath creation and teardown can also be considered in order to produce a more accurate reconfiguration cost model, yielding more optimal orderings or groupings of configuration changes.

Contributions can also be made in the area of smarter reconfiguration policies. These might take the form of simple approaches based on multiple thresholds, or more sophisticated formulations that maximize long-term reward without making strong assumptions about the nature of changes in the traffic demand. In both cases, techniques are needed for computing the optimal policy, which becomes more complicated as the number of thresholds increases, and as the assumptions made in sophisticated formulations are relaxed.

Finally, there are some general open problems that apply to multiple areas of automated reconfiguration. These include the selection of optimal values for parameters of reconfiguration algorithms and thresholds in reconfiguration policies, accurate traffic modeling, and parallelizing computation-intensive tasks using multiple network nodes. In addition, certain specific scenarios can be considered, such as backbone networks that span multiple management domains, and reconfiguration of survivable networks that use protection paths. The first scenario lends itself to a solution where each domain performs reconfiguration on its own subset of resources. In that case, coordination is required between domains as reconfiguration changes in one domain can affect the traffic demand offered to a neighbouring domain. In the second scenario, one faces the problem that working and protection paths must be reconfigured jointly, and survivability must be preserved at every step during the migration phase.

To summarize, the topic of automated reconfiguration for performance management in optical networks is expected to be the subject of many contributions in the near future. While past studies have considered the basic aspects of the problem, most have focused on one key area while largely ignoring others. Further work is needed, especially in the areas of configuration migration and configuration policies, in order to make the reconfiguration process less disruptive, more adaptive, and completely automated.

References

[1] I. Alfouzan and A. Jayasumana. Dynamic reconfiguration of wavelength-routed WDM networks. In *Proc. IEEE LCN*, pages 477–485, November 2001.

[2] I. Baldine and G. N. Rouskas. Dynamic reconfiguration policies for WDM networks. In *Proc. of IEEE INFOCOM*, pages 313–320, 1999.

[3] I. Baldine and G. N. Rouskas. Traffic adaptive WDM networks: A study of reconfiguration issues. *Journal of Lightwave Technology*, 19(4):433–455, April 2001.

[4] M. Brunato, R. Battiti, and E. Salvadori. Load balancing in WDM networks through adaptive routing table changes. In *Proc. NETWORKING*, volume 2345 of *Lecture Notes in Computer Science*, pages 289–300. Springer, 2002.

[5] V. Chan, K. Hall, E. Modiano, and K. Rauschenbach. Architectures and technologies for high-speed optical data networks. *Journal of Lightwave Technology*, 16(12):2146–2168, December 1998.

[6] 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, January 2000.

[7] P. Ernest, G. Mohan, and V. Bharadwaj. An efficient algorithm for virtual topology reconfiguration in WDM optical ring networks. In *Proc. IEEE CCN*, pages 55–60, 2001.

[8] A. Gençata and B. Mukherjee. Virtual-topology adaptation for WDM mesh networks under dynamic traffic. In *Proc. IEEE INFOCOM*, volume 1, pages 48–56, 2002.

[9] J.-F. Labourdette. Traffic optimization and reconfiguration management of multiwavelength multihop broadcast lightwave networks. *Computer Networks and ISDN Systems*, 30:981–998, 1998.

[10] J.-F. Labourdette and A. S. Acampora. Logically rearrangeable multihop lightwave networks. *IEEE Transactions on Communications*, 39(8):1223–1230, August 1991.

[11] J.-F. Labourdette, G. W. Hart, and A. S. Acampora. Branch-exchange sequences for reconfiguration of lightwave networks. *IEEE Transactions on Communications*, 42(10):2822–2832, October 1994.

[12] K. Liu and J. Ryan. All the animals in the zoo: The expanding menagerie of optical components. *IEEE Communications Magazine*, 39(7):110–115, July 2001.

[13] K. H. Liu, L. Changdong, J. L. Pastor, A. Roy, and J. Y. Wei. Performance and testbed study of topology reconfiguration in IP over optical networks. *IEEE Transactions on Communications*, 50(10):1662–1679, October 2002.

[14] M. Maeda. Management and control of transparent optical networks. *IEEE Journal on Selected Areas in Communications*, 16(7):1008–1023, September 1998.

[15] C. Siva Ram Murthy and M. Gurusamy. *WDM Optical Networks: Concepts, Design, and Algorithms*. Prentice Hall, 2002.

[16] 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):1972–1979, October 2000.

[17] R. Ramaswami. Optical fiber communication: From transmission to networking. *IEEE Communications Magazine*, 40(6):138–147, May 2002.

[18] R. Ramaswami and K. N. Sivarajan. Design of logical topologies for wavelength-routed optical networks. *IEEE Journal on Selected Areas in Communications*, 14(5):840–551, June 1996.

[19] V. Sivaraman and G. Rouskas. A High Performance Reservation Protocol with $l$ook-ahead for broadcast WDM networks. In *Proc. of IEEE INFOCOM*, pages 1272–1279, 1997.

[20] N. Sreenath, B.H. Gurucharan, G. Mohan, and C. Siva Ram Murthy. A two-stage approach for virtual topology reconfiguration of WDM optical networks. *Optical Networks Magazine*, 2(3):58–71, May/Jun 2001.

[21] J. Y. Wei. Advances in the management and control of optical Internet. *IEEE Journal on Selected Areas in Communications*, 20(5):768–785, May 2002.

[22] L. Zhang, K.-H. Lee, C.-H. Youn, and H.-G. Yeo. Adaptive virtual topology reconfiguration policy employing multi-stage traffic prediction in optical Internet. In *Proc. IEEE High Performance Switching and Routing*, pages 127–131, 2002.