WDM Optical Network Reconfiguration Using Automated Regression-Based Parameter Value Selection

Wojciech Golab and Raouf Boutaba

University of Waterloo
School of Computer Science
200 University Avenue West
Waterloo, Ontario, Canada
N2L 3G1

E-mail: {wgolab, rboutaba}@bbcr.uwaterloo.ca

***Index Terms*— optical networks, wavelength division multiplexing, reconfiguration, performance management, automated management**

***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 switched architecture it additionally involves the routing of bandwidth-guaranteed circuits known as lightpaths. This paper is concerned with the problem of automatically updating the configuration of an optical network in response to changes in traffic demand, which entails making a reconfiguration policy decision, selecting a new configuration, and migrating from the current to the new configuration. A technique is proposed that automatically selects values for parameters inherent in reconfiguration algorithms and reconfiguration policies, with the goal of maximizing the long-term performance gain due to reconfiguration. The effectiveness of the technique is evaluated in the context of a threshold-based reconfiguration policy. In the best case, the technique is shown to perform only 4% worse than a carefully chosen combination of static parameter values. However, a simple random assignment of parameter values is shown to perform equally well.**

I. INTRODUCTION

In the last decade, the growing popularity of the Internet has created unprecedented demand for transport of data. Optical fibre has emerged as a promising transmission medium in the struggle to accommodate this demand. Various technological advances have made it possible to achieve single-channel data rates of 40 Gb/s, and all-optical transmission distances of thousands of kilometers, by improving the quality of optical signals. Developments have also taken place in the area of wavelength division multiplexing (WDM), which involves the concurrent transmission of optical signals at multiple wavelengths through a common fibre. WDM systems carrying more than 100 densely spaced channels at rates of 10 Gb/s are currently possible, offering aggregate capacities in excess of 1 Tb/s. The combination of fast electronic interfaces and WDM technology promises to go a long way toward exploiting the enormous transmission capacity of optical fibre, which is theoretically in the range of tens of Tb/s.

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

Although the transmission of laser light over fibre plays a key role in today's high-speed optical networks, a wide variety of other components is used in order to transmit and receive data, as well as to condition and split optical signals. The particular combination of components used determines the characteristics of a network such as its traffic handling capacity, traffic grooming ability, fault tolerance, maximum size, and cost. The two most popular such combinations give rise to the passive broadcast architecture and the switched optical architecture [1], [2].

The passive broadcast architecture achieves low cost by avoiding the use of optical amplifiers, repeaters, and electronic switching equipment. Here segments of fibre are combined into a shared broadcast medium using passive optical components such as star couplers. In order to avoid feedback, the access nodes (e.g. high-end workstations and routers) and couplers must form an acyclic topology such as a bus, star, or tree. Consequently, the failure of any one coupler disconnects the network. Another drawback of this architecture is that the loss of signal power due to the splitting of optical signals restricts the network to a small geographical area.

Passive broadcast optical networks rely on WDM technology in order to allow multiple pairs of nodes to communicate at the same time over the shared medium. In the interest of reducing cost, each node is equipped with a small number of opto-electronic receivers and electro-optical transmitters. The simplest such devices are tuned to a fixed wavelength, and as a result each node can communicate directly with a particular subset of the other nodes [2]. Thus, despite there being a broadcast medium, a logical topology is defined over the network, that depends on the configuration of electronic devices. This logical topology can be optimized for a particular pattern of traffic demand, either during the planning stage or after deployment, depending on the tuning capabilities of the hardware [3]. For example, the traffic-weighted average hop count can be minimized by allowing node pairs with the greatest traffic demand to communicate directly, and routing the remaining traffic over multi-hop routes.

A different type of passive broadcast network can be constructed using fast-tunable transmitters and slow-tunable receivers. In this case, the receiver wavelengths are assigned on a long-term basis as in the multi-hop case, but the transmitters are tuned for each burst of packets destined to the same node [4]. This way, full single-hop connectivity can be established irrespective of the number of available wavelengths. The assignment of wavelengths to receivers can still be optimized for a particular traffic demand. That is to say, groups of receivers sharing the same wavelength can be chosen so as to balance the expected traffic load across all the available wavelengths.

The other common optical network architecture, namely the switched optical network, aims to provide greater capacity and resilience than is possible in broadcast networks. Consequently, this architecture is popular in metropolitan and wide area transport networks. The most mature and popular type of switched optical network is SONET/SDH, which is based on time division multiplexing (TDM) using synchronous timing. In typical deployments, a single fibre pair is used between adjacent nodes, each fibre carrying traffic in one direction using a single wavelength. Optical signals are converted entirely into electronic form and then back to optical form (O-E-O) as they pass through each node, which is either an add-drop multiplexer or a digital cross-connect, depending on the level of traffic grooming required. These electronic processing components are configured to provide bandwidth-guaranteed circuits, e.g. among attached IP routers. As in the case of the multi-hop broadcast network, a logical topology arises that can be optimized to reduce the load on costly high-speed routing equipment.

Switched optical networks using digital switching equipment can take advantage of WDM technology provided that the electronic components can accommodate the corresponding increase in aggregate capacity. However, a major drawback of such deployments is that each wavelength passing through a node undergoes O-E-O conversion, while all the traffic entering or leaving the node can typically be groomed onto a much smaller subset of wavelengths. Thus, in a WDM scenario, much of the costly high-speed electronic equipment simply repeats optical signals that could otherwise be passed through using less costly all-optical components. In order to address this issue, switched WDM optical networks should not only make use of digital electronics for aggregation and distribution of low rate electronic signals, but should also perform switching at wavelength or waveband granularity using all-optical versions of the add-drop multiplexer and cross-connect. This possibility makes logical topology configuration evermore important. At the same time, the topology design problem becomes more difficult as one must deal with the constraint of wavelength continuity across switching components while routing all-optical data paths known as lightpaths [5].

To summarize, a key characteristic of optical networks is the ability to optimize their configuration to suit a particular traffic demand pattern. This optimization has been studied extensively in the context of network planning and provisioning. However, it can also be considered in the context of performance management, whereby the network is reconfigured in response to a changing pattern of traffic demand. Although some variations, such as daily or seasonal oscillations in network usage, can affect all nodes in a nearly uniform manner and can be dealt with by balancing the load during the initial deployment of the network, dynamic reconfiguration can be necessary in other situations. These include the following: changes to the physical topology; changes to the set of network customers or to customer activity patterns; and even daily load variations in transport networks carrying traffic across multiple time zones, where the oscillations in traffic load can occur at different phases depending on the direction of traffic.

The problem of optical network reconfiguration is a challenging one as the reconfiguration process typically entails the temporary release of network resources, namely optical transmitters, receivers, and cross-connections, which can potentially cause loss of data in high-rate aggregated 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 in a varying traffic environment.

The remainder of this paper is organized as follows. First we give an overview of relevant literature, organized by subproblem. Next, we propose a solution to the problem of automated parameter value selection for reconfiguration algorithms and reconfiguration policies. Finally, we present simulation results and conclude with a critical evaluation and discussion of future work.

II. Related Work

This part of the paper reviews proposed approaches to the optical network reconfiguration problem. Section II-A discusses solutions to the subproblems of new configuration selection and configuration migration. Then, section II-B examines reconfiguration policies.

A. *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 approximately optimal mathematical formulations to dumb searches.

A direct approach to the reconfiguration problem is characterized by the selection of a new configuration independently of the current configuration. The new configuration is optimized with respect to the expected traffic pattern, regardless of the number of configuration changes required to reach it. Examples of this approach for various optical network architectures are discussed in [6], [7], [8]. The optimizations described in these works are approximate due to the use of simplified mathematical formulations and heuristics, as the full optimization problems are intractable. The distance between the old and new configuration (i.e. the number of configuration changes required during the configuration migration phase) is used to break ties between optimal new configurations.

The configuration migration subproblem is solved separately in a direct reconfiguration approach, once the new configuration has been selected. In single-hop broadcast networks, this amounts to retuning receivers. Rather than retuning all affected devices simultaneously, receivers can be retuned in smaller groups in order to reduce packet loss [9]. Configuration migration in multi-hop broadcast networks is more complex as wavelengths are dedicated to particular logical links, and retuning a single receiver or transmitter can destroy an entire logical link. However, the two receivers and two transmitters corresponding to a pair of links can be retuned in such a way that the number of logical links is preserved, and consequently network disruption is reduced. This is known as a branch exchange operation, and is investigated in [10], where heuristics are proposed that determine a short sequence of such operations that transforms the current configuration to the new configuration. Finally, configuration migration for switched optical networks is addressed in [11], where an algorithm is proposed that transforms the configuration of the network by creating and destroying entire lightpaths. The order of operations is chosen so that lightpath teardown is postponed as much as possible in order to minimize network disruption.

Another category of reconfiguration approaches, subsequently referred to as partial reconfiguration approaches, is based on the tradeoff between the optimality of the new configuration and the degree of network disruption incurred during the migration phase. The general strategy here is to identify the most loaded logical links and perform operations that redistribute the load in these links. A simple example of this in the context of single-hop broadcast networks is presented in [12], where the wavelength assignment of a receiver tuned to the most loaded wavelength is swapped with the assignment of a receiver tuned to the least loaded wavelength. The number of iterations can be varied to control the extent of reconfiguration. Another example, applicable to switched optical networks, is investigated in [13], and is presented here as Algorithm 1 since it is used in the subsequent simulation study.

Algorithm 1: Heuristic Adaptation Algorithm

Input: old configuration, traffic demand, thresholds W_L and W_H
Output: new configuration

1. if at least one pair of nodes with nonzero traffic demand is disconnected in the logical topology, attempt to establish a lightpath joining the pair with the greatest traffic demand
2. else if the load in some lightpath exceeds W_H then attempt to create a lightpath between the endpoints of the greatest contributing multi-hop traffic flow, or alternately the endpoints of the lightpath itself
3. else if no lightpath can be added and some lightpath load is below W_L, then delete the least loaded such lightpath whose deletion does not disconnect two nodes in the logical topology with nonzero traffic demand

The novel idea in this algorithm is to maintain free resources in the network by deleting underutilized lightpaths.

A more complicated algorithm for the switched architecture is discussed in [14]. Like Algorithm 1, it attempts to reduce the traffic-weighted average hop count by establishing lightpaths between the node pairs with the greatest traffic demand. However, rather than proactively freeing resources, subsets of existing lightpaths are considered for deletion only when this is necessary in order to create new lightpaths. The degree of network disruption is controlled by limiting the number of node pairs and lightpath subsets considered, as well as the number of configuration changes performed. A similar algorithm applicable only to the popular ring topology is presented in [15]. The configuration migration phase there is carried out using lightpath merge and split operations, where cross-connections are created or torn down in nodes on the segment between the particular high-traffic node pair under consideration.

A third category of reconfiguration approaches is based on a local search heuristic, whereby a set of neighbouring configurations is exhaustively explored and the best such configuration is adopted. Reference [16] discusses such a technique for switched ring networks, where the configuration neighbourhood is defined as the set of configurations reachable by a single 3-branch exchange operation. This operation entails cyclically permuting the destination nodes of three lightpaths, and is guaranteed to preserve the ring topology. The number of possible 3-branch exchanges is cubic in the number of lightpaths in a ring, and determines the time complexity of this algorithm.

Of all the reconfiguration techniques discussed so far, partial reconfiguration approaches are the most flexible in terms of the ability to control the running time and degree of network disruption through input parameters. In comparison, direct approaches do not achieve a good tradeoff between performance and cost. Local search approaches are also among the less flexible, by nature of their simplicity, and have limited applicability among various optical network architectures and topologies.

B. Reconfiguration Policies

Reconfiguration policies are concerned with the process of deciding when configuration updates should be performed. Of those works that do not ignore this important subproblem, most indicate one of three types of rudimentary approaches. The simplest of all techniques is to execute the reconfiguration algorithm at regular intervals of time [12], [16], [17]. 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 reconfiguration algorithm. Other approaches take the current traffic demand into consideration. Reconfiguration can be performed at every traffic change [14], or when an important change is detected, for example on the basis of load thresholds [13].

More sophisticated policies can be obtained using techniques that take into account long-term effects of policy decisions rather than merely the immediate reward due to improvement in some performance metric. For example, a Markov decision process formulation is proposed in [18] in the context of a single-hop broadcast WDM network. Here the state of the network is represented in terms of the current configuration and traffic pattern, and a traffic model is assumed such that the Markovian state transition property is satisfied. The reward of each state is a function of the degree of load balancing, and the cost of each transition is based on the number or receiver retunings required. The policy iteration algorithm can be used to obtain the optimal policy in this model after applying some approximations to reduce the size of the state space. A performance comparison shows that much simpler policies based on performance metric thresholds sometimes perform just as well as the optimal policies obtained using the Markovian model. This is a consequence of the assumptions of the Markovian formulation and the simplifications used to solve it.

Another formulation of the reconfiguration policy problem, similar to the Markov decision process formulation, is presented in [19]. However, rather than precomputing an optimal policy, the optimal decision in a given state is computed dynamically using traffic prediction data. Knowing the traffic pattern during a fixed number of future state transitions makes it possible to evaluate the reward incurred by all possible combinations of decisions that can be made during that period of time. The optimal decision can then be selected on the basis of this approximate long-term reward, without using the policy iteration algorithm.

Intuitively, sophisticated approaches to the reconfiguration policy 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. However, policies in sophisticated approaches are difficult to optimize without either a rigid mathematical formulation or knowledge of the future traffic pattern. Threshold-based policies are attractive on account of their simplicity, but their performance is dependent on judicious selection of threshold values.

III. Automated Parameter Value Selection

The remainder of the paper discusses a solution to the problem of automated parameter value selection. We investigate a general technique that can be used to optimize the values of numerical parameters, such as inputs to reconfiguration algorithms that control the number of changes made, or thresholds in reconfiguration policies. For comparison, two additional techniques are considered that assign the values of parameters randomly. The effectiveness of all three techniques is evaluated in the context of the reconfiguration scheme proposed in [13], where the reconfiguration policy is based on two load threshold parameters. Subsequent sections describe in detail the parameter value selection techniques, the simulation environment, and simulation results.

A. Regression-Based Technique

Our approach to selection of parameter values is based on a mathematical formulation, where a performance metric such as the rate of traffic loss is modelled as a differentiable function of multiple variables, including the parameters under consideration. The optimization technique applied to this function is based on the steepest descent method. Specifically, let T denote the objective function to be minimized and let $p_1, p_2, \ldots, p_n$ denote the parameters. The gradient ∇T is computed at time intervals of length Δ_t. At the end of each interval the values of the parameters are updated according to the formula

$$p_i \longleftarrow p_i - \operatorname{sign}\left(\frac{\partial T}{\partial p_i}\right) \cdot \Delta_{p_i}$$

where Δ_{p_i} is a small positive constant.

Because the closed form of T is generally not known, an approximate form based on a partial power series expansion is used:

$$T(p_1, p_2, \ldots, p_n, \bullet) \approx c_0 + c_1 p_1 + c_2 p_2 + \ldots + c_n p_n$$

Here $\bullet$ represents additional independent variables such as time. The values of the constants $c_0, \ldots, c_n$, and consequently the values of the partial derivatives with respect to $p_1, \ldots, p_n$, can be computed by perturbing $p_1, \ldots, p_n$ and measuring the effect on the value of T. Specifically, this can be performed as follows. Let $p'_1, \ldots, p'_n$ be the parameter values set at the end of the last time interval of length Δ_t. The parameters are randomly perturbed at N equally spaced times during the next interval according to the formula

$$p_i \longleftarrow p'_i + \phi \Delta_{p_i} / 2$$

where ϕ is chosen uniformly at random to be either $+1$ or -1. The value of T is observed over a period of Δ_t/N time units following each perturbation. Thus, N tuples of data of the form $[T, p_1, p_2, \ldots, p_n]$ are collected for each interval of length Δ_t. Applying linear least squares regression to these data, we obtain the constants $c_0, \ldots, c_n$. The gradient ∇T can then be approximated according to $\frac{\partial T}{\partial p_i} \approx c_i$. The approximation is accurate only if the constants Δ_{p_i} are sufficiently small.

In the context of the reconfiguration scheme proposed in [13], the two parameters of interest are the high and low load thresholds that are used as inputs to the Heuristic Adaptation Algorithm (Algorithm 1). Since the range of legal values of these parameters is known in advance (i.e. the interval $[0, 1]$), updates to parameter values at the end of each time interval of length Δ_t are appropriately constrained. Similarly, because the two parameters have a common domain and similar semantics, it is reasonable to define $\Delta_{p_1} = \Delta_{p_2} = \Delta_p$.

B. Experimental Controls

The following two techniques are used as experimental controls in the simulation study described below. The first one, subsequently referred to as the Brownian technique, is the same as the above technique except that the constants $c_0, \ldots, c_n$ are assigned uniformly at random in the interval $[-0.5, 0.5]$. This way, the signs of the estimated partial derivatives $\frac{\partial T}{\partial p_i}$ are equally likely to be $+1$ and -1, and the values of the parameters p_i exhibit variation analogous to Brownian motion. This control makes it possible to evaluate the accuracy of the proposed gradient approximation technique. The second technique, subsequently referred to as the uniformly random technique, uniformly randomly assigns the values of the parameters in their respective domains. This occurs N times in each interval of length Δ_t, i.e. at the same times as the perturbations in the regression-based technique. Both controls restrict the values of the parameters to the domain $[0, 1]$.

C. Simulation Environment

The simulation environment used to produce the results presented in the next section is based on the one described in [13]. The physical network topology is shown in Figure 1. The network consists of 16 nodes and 26 bidirectional links. Each node represents a cross-connect device that performs all-optical switching at wavelength granularity, as well as an electronic processing engine equipped with 8 receiver/transmitter pairs. No wavelength conversion or traffic buffering capability is assumed in the cross-connect devices. For simplicity, lightpath establishment and teardown are assumed to occur instantaneously. Each link carries up to 16 wavelengths. Each wavelength carries a unidirectional optical signal with a data rate of 10 Gb/s. Lightpaths are created over shortest routes, using the first available wavelength.

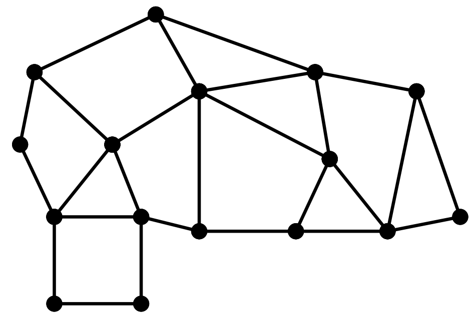

Fig. 1. Physical network topology used during simulation.

The traffic model aims to represent random short-term variations, as well as regular oscillations due to time of day, occurring at various phases as in a wide area network spanning multiple time zones. The directed traffic demand in Mb/s between each pair of nodes is of the following form:

$$D(t) = \gamma \cdot \alpha \cdot \beta(t) \cdot \left[1 + \frac{\sin(2\pi t + \delta)}{2}\right]$$

Time t is measured in units of days. The constants γ, α, and δ are set per run of the simulation. The constant γ, in units of Mb/s, is global and controls the overall magnitude of the traffic demand matrix. The constant α determines the relative demand between a particular pair of nodes. It is randomly set for each node pair with the following distribution: zero with a probability of 0.4, and uniformly random on the interval $[0.5, 1.5]$ with a probability of 0.6. Thus, on average 60% of the node pairs exhibit a nonzero traffic demand in one run of the simulation. The phase shift δ is selected uniformly at random for each node pair from the interval $[-0.3, 0.3]$. The function $\beta(t)$ adds a bursty character to the traffic demands. Its value is randomly chosen for each node pair at times $t = k \cdot 0.05$ days, $k \in \mathbb{Z}$, and is determined by linear interpolation at times between. Each random value is chosen from the interval $[1-\rho, 1+\rho]$, where ρ is a global constant that determines the overall degree of burstiness, for example $\rho = 0.5$. A sample traffic trace is presented in Figure 2.

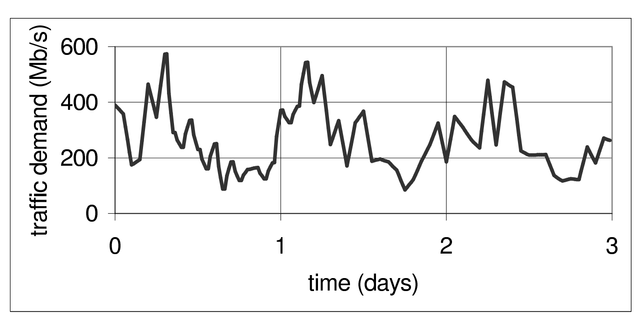

Fig. 2. Example traffic trace generated using $\gamma = 500$ Mb/s and $\rho = 0.5$.

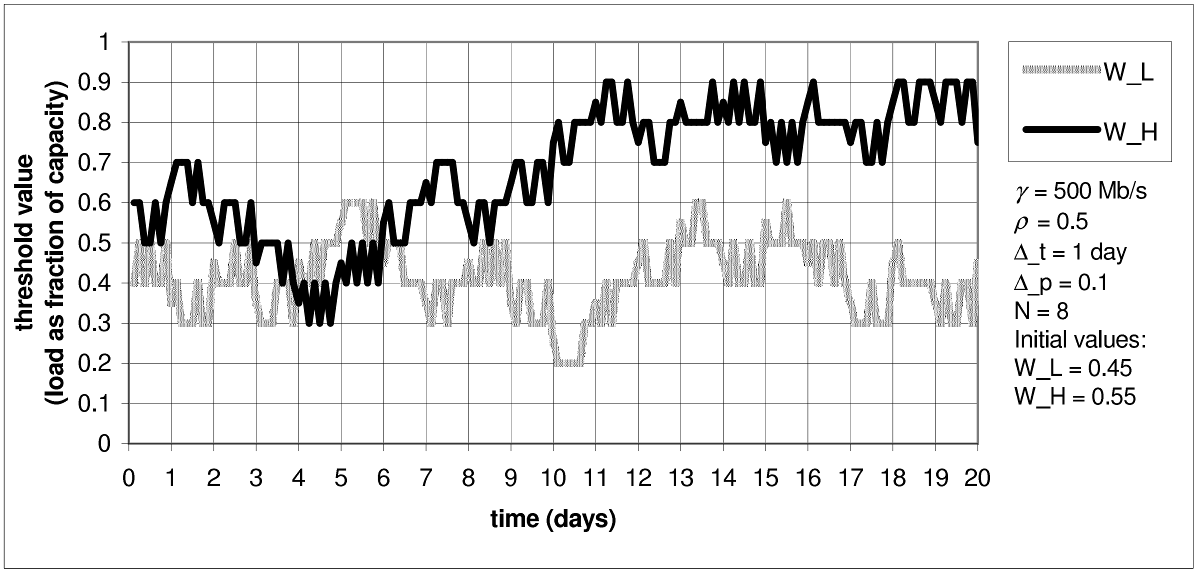

Fig. 3. Updates to threshold values over a twenty-day period using the regression-based parameter value selection technique.

D. Simulation Results

The operation of the regression-based parameter value selection technique is illustrated in Figure 3, which presents the sequence of values of W_L and W_H generated over a twenty-day simulation period. Repeated runs of the simulation show that the weak stabilization of parameter values seen in the figure is typical.

Next, the performance of the Heuristic Adaptation Algorithm (Algorithm 1) is evaluated using static combinations of threshold values in order to establish a performance baseline for the automatic parameter value selection techniques. Figure 4 presents traffic loss over a ten-day simulation period, using a regular mesh of parameter values.

There are two causes of traffic loss in the simulation: link congestion and lack of connectivity in the logical topology. The figure shows that the total traffic loss decreases as W_L increases. This is consistent with the fact that as W_L increases, the amount of spare resources increases, which facilitates reconfiguration. Similarly, as W_H decreases, spare resources are reallocated more intensively, and performance improves. However, the figure shows that the latter effect is much weaker, and that the change in traffic loss is not always monotonic. The figure also suggests that the region does not have any pronounced local minima that might interfere with the steepest descent search. The search is expected to eventually settle in the rightmost region of the figure and achieve a long-term rate of traffic loss between 8.0 and 8.4 Tb / 10 days.

Finally, Figure 5 presents a performance comparison of static threshold value combinations and automatic parameter value selection techniques. The leftmost two bars represent static thresholds. The minimum and maximum traffic loss values observed in Figure 4 are shown, the

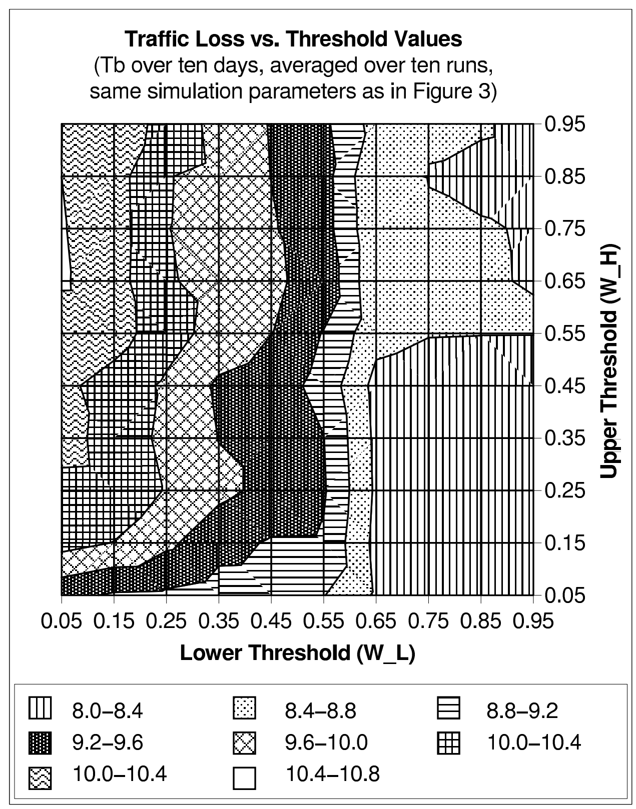

Fig. 4. Performance of Heuristic Adaptation Algorithm (Algorithm 1) using various combinations of static threshold values.

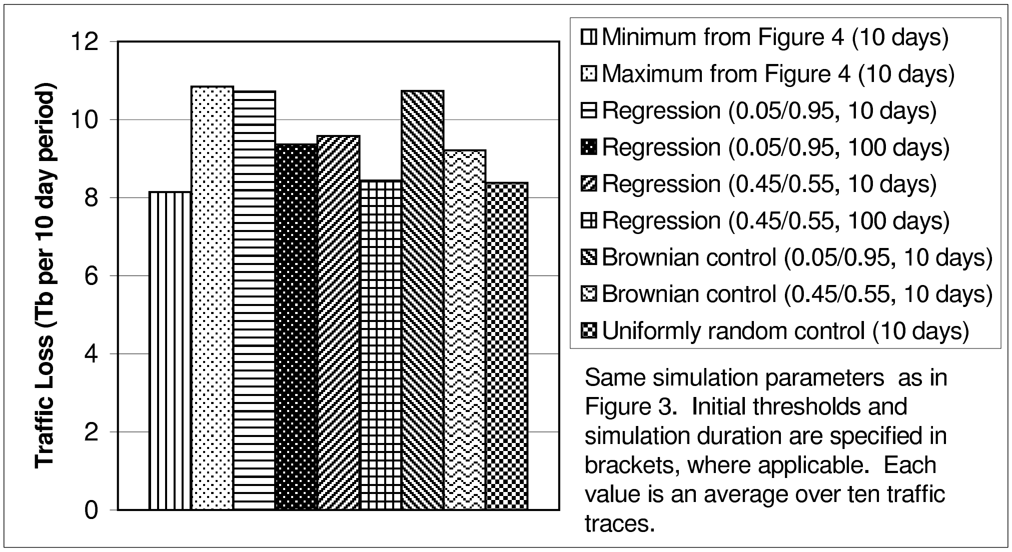

Fig. 5. Traffic loss per ten-day period incurred using fixed thresholds, the regression-based technique, as well as the randomized control techniques.

maximum being 33% greater than the minimum. The next two bars show the performance of the regression-based parameter value selection technique, using the initial values $W_L = 0.05$ and $W_H = 0.95$. Over a ten-day period the performance is 32% worse than the minimum from Figure 4. However, amortized traffic loss over a 100-day period is only 15% worse. This suggests that in the long run the technique does improve performance by selecting better parameter values. Similar behaviour occurs using initial values $W_L = 0.45$ and $W_H = 0.55$. Amortized traffic loss over a 100-day period is only 4% greater than the minimum from Figure 4. At 8.4 Tb / 10 days, the traffic loss is as predicted above.

Looking at the experimental controls, the Brownian technique performs as well as the regression-based technique over a ten-day period. Using initial values $W_L = 0.45$ and $W_H = 0.55$, the traffic loss is 13% greater than the minimum from Figure 4. The uniformly random technique performs best of all the automated techniques over a ten-day period, achieving traffic loss only 3% worse than the minimum from Figure 4.

IV. Conclusions and Future Work

The results presented in Figure 5 demonstrate that the regression-based parameter value selection technique does work over time scales of tens of Δ_t. This is seen by considering performance levels relative to the minimum and maximum values exhibited in Figure 4, and relative to the Brownian technique. In the best case, the level of traffic loss is 4% greater than the minimum value from Figure 4, and is close to the level expected based on the features of the objective function shown in that figure. Because amortized performance is seen to improve as the length of the simulation period increases, it may be possible to achieve even better performance over longer time scales. Interestingly, the very simple uniformly random technique performs as well as the regression-based technique in the best case. A possible explanation for this unexpected outcome is that in order to achieve optimal performance levels, it suffices to use optimal parameter values on occasion, as opposed to in the long term. The uniformly random scheme certainly does this by design, albeit blindly.

In terms of automation, the proposed regression-based technique can potentially simplify the task of performance management by selecting appropriate values for relevant parameters, but at the same time it introduces parameters of its own, namely Δ_t, Δ_p, and N. However, there are two benefits to using the technique. First, the number of free parameters (i.e. those that must be set by hand) is fixed at three, which in some scenarios is a reduction. Second, the remaining free parameters have simple semantics that are independent of the original parameters under consideration. For example, $\Delta_t = 1$ day was used in this study since it coincides with the period of the traffic demand functions, and is the shortest observation period during which the level of traffic loss is not dependent on the starting time of the period. During experimentation, larger values of Δ_t appeared to provide little benefit, while smaller values gave rise to weaker stabilization of the parameter values with increasing time.

The regression-based technique is general in the sense that it makes few assumptions concerning the objective function and its parameters. Specifically, it only requires the objective function T to resemble a differentiable function in the region over which the parameter values are selected. However, achieving near-optimal performance is dependent on the shape of the objective function, presence of local extrema in particular, and on the choice of values for the parameters Δ_t, Δ_p, and N. In the simulation study presented above, a good level of performance was achieved, but only over a long simulation period. Thus, another drawback of the regression-based technique is potentially slow convergence. This drawback does not apply to the uniformly random technique.

Further analysis is needed in order to evaluate the behaviour of the proposed regression-based technique with respect to the variables Δ_t, Δ_p, and N, the traffic model parameters α, γ, and ρ, and the physical network configuration including the presence of wavelength conversion. A more sophisticated reconfiguration cost model should also be examined, where a finite period of time is needed to establish and tear down cross-connections, during which traffic must be buffered. Also, additional application scenarios must be considered, for example optimizing the values of the parameters used in the heuristic proposed in [14], in order to evaluate the applicability of the technique. Objective functions other than traffic loss can be studied, for example maximum link load or buffer occupancy. Finally, further investigation is required in order to account for the surprisingly good performance of the uniformly random technique.

Several improvements can be made to the regression-based technique in an attempt to improve its performance. In order to speed up initial convergence, a systematic search phase can be added that considers various parameter value combinations, for example over a regular mesh in the parameter space, as in Figure 4. Another modification is to vary the magnitude of the perturbations to the values of the parameters, which would help to deal with the presence of local extrema. Parameter value updates occurring at time intervals of length Δ_t can also be enhanced by applying a statistical test to the regression coefficients in order to determine whether their sign is statistically significant. This information could be used to produce more intelligent parameter value update decisions. Finally, greater use of historical data can be made in the approximation of the gradient, which currently only uses measurements from the last interval of length Δ_t.

REFERENCES

[1] V. Chan, K. Hall, E. Modiano, and K. Rauschenbach, "Architectures and technologies for high-speed optical data networks," *Journal of Lightwave Technology*, vol. 16, no. 12, pp. 2146–2168, Dec. 1998.

[2] J.-F. Labourdette and A. S. Acampora, "Logically rearrangeable multihop lightwave networks," *IEEE Transactions on Communications*, vol. 39, no. 8, pp. 1223–1230, Aug. 1991.

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

[4] V. Sivaraman and G. Rouskas, "A High Performance Reservation Protocol with look-ahead for broadcast WDM networks," in *Proc. of IEEE INFOCOM*, 1997, pp. 1272–1279.

[5] R. Dutta and G. N. Rouskas, "A survey of virtual topology design algorithms for wavelength routed optical networks," *Optical Networks Magazine*, vol. 1, no. 1, pp. 73–89, Jan. 2000.

[6] I. Baldine and G. Rouskas, "Reconfiguration and dynamic load balancing in broadcast WDM networks," *Photonic Network Communications Journal*, vol. 1, no. 1, pp. 49–64, June 1999.

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

[8] B. Ramamurthy and A. Ramakrishnan, "Virtual topology reconfiguration of wavelength routed optical WDM networks," in *Proc. GLOBECOM*, Nov. 2000, pp. 1269–1275.

[9] I. Baldine and G. N. Rouskas, "Traffic adaptive WDM networks: A study of reconfiguration issues," *Journal of Lightwave Technology*, vol. 19, no. 4, pp. 433–455, Apr. 2001.

[10] J.-F. Labourdette, G. W. Hart, and A. S. Acampora, "Branch-exchange sequences for reconfiguration of lightwave networks," *IEEE Transactions on Communications*, vol. 42, no. 10, pp. 2822–2832, Oct. 1994.

[11] 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*, vol. 50, no. 10, pp. 1662–1679, Oct. 2002.

[12] I. Alfouzan and A. Jayasumana, "Dynamic reconfiguration of wavelength-routed WDM networks," in *Proc. IEEE LCN*, Nov. 2001, pp. 477–485.

[13] A. Gençata and B. Mukherjee, "Virtual-topology adaptation for WDM mesh networks under dynamic traffic," in *Proc. IEEE INFOCOM*, vol. 1, 2002, pp. 48–56.

[14] N. Sreenath, B. Gurucharan, G. Mohan, and C. S. R. Murthy, "A two-stage approach for virtual topology reconfiguration of WDM optical networks," *Optical Networks Magazine*, vol. 2, no. 3, pp. 58–71, May/Jun 2001.

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

[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*, vol. 18, no. 10, pp. 1972–1979, Oct. 2000.

[17] M. Brunato, R. Battiti, and E. Salvadori, "Load balancing in WDM networks through adaptive routing table changes," in *Proc. NETWORKING*, ser. Lecture Notes in Computer Science, vol. 2345. Springer, 2002, pp. 289–300.

[18] I. Baldine and G. N. Rouskas, "Dynamic reconfiguration policies for WDM networks," in *Proc. of IEEE INFOCOM*, 1999, pp. 313–320.

[19] 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*, 2002, pp. 127–131.