



Dynamic Lightpath Scheduling in WDM Optical Networks: A Two-Phase Approach

Dr K Venkataramana^{1*}, Dr. Ramakanta Bhoi²^{1*} Associate Professor Dept. Of Computer Science and Engineering, NIT , BBSR² Associate Professor, Dept. Of Computer Science and Engineering, NIT , BBSR

kvenkata@thenalanda.com* ramakantabhoi@thenalanda.com

Abstract — Next-generation wavelength-division multiplexing (WDM) optical networks have the important capability of lightpath scheduling, which allows for the provisioning of end-to-end lightpaths while reserving resources in advance for a predetermined time period. Dynamic scheduled lightpath demands (D-SLDs) from end users must be met in a dynamic context without awareness of pending requests. Even though the request's starting time may be hours or days in the future, the end user still wants to know whether it can be fulfilled right away. . To schedule and furnish D-SLDs in a dynamic manner, we suggest a two-phase method. The network control plane prepares a path with assured resources in the first phase, also known as the deterministic lightpath scheduling phase, in response to the arrival of a lightpath request so that the user can receive a prompt response with a deterministic lightpath schedule. We re-provision some already planned lightpaths to re-optimize for enhancing network performance in the second phase, which is known as the lightpath re-optimization phase. To redistribute network resources while preserving the current lightpath schedules, we investigate two re-optimization scenarios. Our suggested two-phase dynamic lightpath scheduling strategy can significantly lessen network blockage, according to experimental data.

I. INTRODUCTION

The next-generation wavelength-division multiplexing (WDM) optical networks will be a key enabler for many high-end applications, including those using Grid technologies, by provisioning end-to-end lightpaths in an on-demand manner. Unlike previous optical bandwidth consumers, end users largely control these new applications and thus the bandwidth demands come directly from the end users' requests. Such demands are usually dynamic, which implies that the network operation based on the assumption of static or predictable demands will be considerably inefficient. In addition, many end users need guaranteed lightpath connections during a specified period of time in the future. They usually prefer to make advance reservations for end-to-end lightpaths with predefined service durations where the starting time of the lightpath demand can be days to weeks in the future. Such an advance reservation of a lightpath is called scheduling of a lightpath and correspondingly the lightpath itself is termed as a scheduled lightpath demand (SLD) [1]. Many SLDs arrive in a dynamic manner. We distinguish between these dynamic scheduled lightpath demands (D-SLDs) as opposed to the concept of static scheduled lightpath demands (S-SLDs), where the whole set of lightpath demands are known before any actual provisioning happens in the network. Therefore, the time schedule of every S-SLD is known in advance for the network control plane. We study the problem of bandwidth allocation for D-SLDs in this paper. For the purpose of scheduling, we assume that the network time is slotted. The duration of a scheduled lightpath is measured in number of time slots. Each time slot has equal length.

In practical network operation, many end users require deterministic services. By a deterministic service, we mean that after submitting an advance reservation request for a lightpath, a user expects a quick deterministic answer as to whether the request can be accommodated and if so, the precise

schedule information for the request. Although the deterministic answer returned for a D-SLD can be optimal at the current time, the resource allocation in the network may still become sub-optimal with the arrivals of future D-SLDs. Unlike S-SLDs, arrivals of D-SLDs cannot be precisely predicted and thus resource allocation for D-SLDs is difficult to be optimized. However, before a D-SLD is physically provisioned in the network, any reallocation of resources reserved for this D-SLD without changing the schedule, e.g., rerouting and reassignment of wavelength, are transparent to the end user. Therefore, we have the opportunity to perform re-optimization for all D-SLDs scheduled to be set up in the future. In this paper, we propose that resources reserved for scheduled lightpaths be re-optimized before they are physically provisioned to improve network performance. Emerging Grid applications being deployed or planned on the emerging national and regional optical network facilities, such as the Internet2 Hybrid Optical Packet Infrastructure (HOPI), National Lambda Rail (NLR), UltraLight and TeraGrid, have similar requirements and also fit into this class of D-SLD requests. We use two example applications to characterize the two types of D-SLDs. The first application is to schedule a real-time, collaborative scientific experiment. A lightpath is scheduled between the data gathering and data processing facilities at a fixed start time and a fixed end time (i.e., for a fixed duration), e.g., [9:00 A.M., 10:00 A.M.] on Monday. We denote such a demand to be of the *time-fixed type*. As the second example, a financial institution demands a scheduled lightpath on a weekly basis to backup the huge amount of data to a data repository in a different city. The transfer takes up to one hour. The demand is specified with a loose starting time window, say [1:00 A.M., 5:00 A.M.] on Saturdays, during

which any lightpath starting time is acceptable. We denote such a demand to be of the *time-window type*.

In this paper, we propose a two-phase dynamic lightpath scheduling approach. In the first phase, we employ a deterministic lightpath scheduling algorithm to determine if the current request can be accommodated into the network along with a schedule if the D-SLD is a time-window type. This is called the *deterministic lightpath scheduling phase*, in which each request is assigned a route and wavelength with a fixed time schedule or is denied. In the second phase, a re-optimization procedure is carried out to re-provision those lightpaths that have been scheduled, but are not in service yet (i.e., they have not been physically provisioned). This is called the *lightpath re-optimization phase*. The lightpath re-optimization phase is triggered when a blocking of D-SLD occurs in the phase 1 or during the beginning of a time-slot where one of more scheduled D-SLDs start. Note that any re-provisioned lightpath in phase 2 must be guaranteed a set of resources (path and wavelength) and its time schedule should be unaltered.

The scheduling or advance reservation issue is common in communication networks that allow for reserving and dedicating resources to specific traffic demands. Below, we briefly review the related work in this area. Advance reservation in general networks was studied in [3]. An approach was proposed in [4] for providing advance resource reservation in ATM networks. In [5], the authors studied the complexity of path computation for advance reservation. For the case of WDM wavelength-routed networks, in [1], [2], [6],

[7] and [8] the authors studied the static lightpath scheduling problem. Both integer linear programming (ILP) and heuristic and meta heuristics approaches were used to accommodate the static schedulable lightpath demands (S-SLD). None of the above research studies considered dynamic SLD requests and most of them except [1], [2] and [8] considered only time-fixed demands. In [9], the authors proposed a dynamic traffic model for advance reservation. Several simple RWA algorithms were proposed to schedule the D-SLDs.

The remaining portion of this paper is organized as follows. Section II presents the network model and problem statement. Section III describes our deterministic lightpath scheduling algorithm for both time-fixed and time-window D-SLDs in the deterministic lightpath scheduling phase. Section IV describes a re-optimization procedure and several rescheduling schemes for the lightpath re-optimization phase. Section V presents our experimental results. Section VI concludes this paper.

II. NETWORK MODEL AND PROBLEM STATEMENT

A. Network Model

In this paper, we consider WDM wavelength-routed mesh networks. Such a network consists of a set of reconfigurable optical cross-connects (OXC) interconnected by optical fiber links. Each link has two fibers in opposite directions, while each fiber has a fixed number of wavelengths. We assume that the OXCs have no wavelength conversion capability; thus all lightpaths are subject to wavelength continuity constraint. We make the following operational assumptions and use these notations in our network model.

- Network time is slotted. A time slot is the minimum time unit in the network, each having equal, fixed length. We denote the time slots starting from a given time 0 by using a sequence of indices $\{ T_0, T_1, \dots \}$.
- The duration of any D-SLD is an integral multiple of a time slot. A lightpath can only be scheduled at the beginning of a time slot. Note that in the rest of this paper, we use the term lightpath interchangeably with D-SLD.
- Each D-SLD occupies the whole capacity of a single wavelength. The i^{th} D-SLD in the network is denoted by $(s^i, d^i, t^i, \square^i, l^i)$, where s^i is the source node, d^i is the destination node, t^i and \square^i are the starting time and duration (in time slots) for the scheduled lightpath, and l^i is the maximum path length in kilometers. For a time-fixed D-SLD, t^i is a fixed value. For a time-window D-SLD, t^i represents a range of contiguous time slots.
- For any wavelength on any link, its availability in one time slot is independent of its availability in other time slots. However, the wavelength assigned for a specific lightpath must be the same when that lightpath spans multiple time slots.

A lightpath in the network is in one of the following two states:

- (a) *scheduled*: In this state, the lightpath has been scheduled and its starting time cannot be changed; however, its route and wavelength may be re-assigned.
- (b) *in-service*: In this state, the lightpath has been physically provisioned and its starting time, routing and wavelength assignment cannot be changed.

As discussed in Section I, once a lightpath is scheduled, its starting time cannot be changed because the request has been guaranteed resources. However, a lightpath can be rerouted for the re-optimization purpose at any time before its starting time. To avoid interrupting ongoing services, the lightpath cannot be altered when it is in service (after its scheduled starting time). Each lightpath is also subject to a maximum path length constraint, which is usually considered in practice to address physical impairments and end-to-end delay.

B. Problem Statement

We propose a solution to the scheduling and provisioning of D-SLD requests using a two phase approach. Phase I is the deterministic lightpath scheduling phase and phase II is the lightpath re-optimization phase. When a D-SLD request arrives, the system enters phase I, where the system

quickly checks if the D-SLD can be scheduled with a route and wavelength being reserved at its starting time and for the requested duration. Phase II re-provisions those scheduled lightpaths to improve network performance. In theory, we can perform re-optimization after the arrival of each D-SLD, whose overhead, however, would be overwhelming. To reduce the overhead, we perform lightpath re-optimization only during the following two scenarios, namely, *re-optimization at blocking* and *re-optimization at kick-off*. In re-optimization at blocking, the system enters phase II when a D-SLD is blocked in the phase I. In re-optimization at kick-off the system enters phase II at the beginning of each time slot, to re-optimize the resource allocation of some scheduled lightpaths.

In the following part of this section, we present the problem statement for the deterministic lightpath scheduling and lightpath re-optimization phases respectively. We first define some common notations. By using a graph $G(V, E)$, we represent an optical network topology, where V is a set of nodes and E is a set of links in the network. By $U[t]$ we denote the number of used wavelength-links during time slot t , and by $u_e[t]$ the number of used wavelengths during time slot t on the link $e \in E$. Therefore, $U[t] = \sum_{e \in E} u_e[t]$.

Phase I – Deterministic Lightpath Scheduling

Given a D-SLD $(s^i, d^i, t^i, \lambda^i, l^i)$ and a network topology $G(V, E)$ with wavelength availability information, the problem is to provision the lightpath with two possible objectives as follows:

Objective 1 (MWL): Minimizing the number of used $t_c \in \mathbb{1}$. The objective is either MWL or LB as defined in phase I.

III. DETERMINISTIC LIGHTPATH SCHEDULING ALGORITHM

We develop a deterministic lightpath scheduling algorithm for solving the problem in phase I. Given a D-SLD, we need to compute a path, a wavelength, and a starting time slot for the request. The algorithm should find a solution quickly so that the requestor is provided an answer in a short time.

We use fixed-alternate routing as the path selection scheme. A set of pre-computed k-shortest paths between the source and destination are used as the candidate routes for the demand. They are denoted by $\{P_1, P_2, \dots, P_K\}$. These k-shortest paths should satisfy the path length constraint on the demand.

We develop a wavelength assignment scheme called Slotted First-Fit (SFF) in our deterministic lightpath scheduling algorithm. A wavelength is said to be available for a set of contiguous time slots if and only if it is available in each of these time slots. Given the starting time slot, the duration and a pre-computed routing path for a D-SLD, SFF picks up the first wavelength that is available on every link of the path in all the time slots from its starting time through its ending time. The complete design of our deterministic lightpath scheduling algorithm is described in Figure 2. The inputs are a topology graph $G(V, E)$ and a D-SLD represented by (s, d, t, λ, l) , where the value of t is an integer in a range $[t_{lower}, t_{upper}]$.

Wavelength-Links between time slot i , minimizing $\sum_{e \in E} U[t]$.



t^i and
 $t^i \square \square i \square 1,$

For a time-fixed D-SLD, $t_{lower} \square t_{upper}$. For a time-window D-SLD, $t_{lower} \square t_{upper}$. The algorithm loops through each

$$t^i \square t \square t^i \square \square i \square 1$$

Objective 2 (LB): Load-Balancing, i.e., minimizing

possible starting time slots between t_{lower}

and t_{upper}

(in step 3)

$MAX u_e [t]$. In other words, LB strives to

$$t^i \square t \square t^i \square \square i \square 1, e \square V$$

minimize the maximum number of used wavelengths among all the links in the network.

Phase II – Lightpath Re-optimization

Different objectives should be applied under different re-optimization scenarios. Therefore, we define separate sub- problems for the re-optimization phase.

- *Re-optimization at blocking scenario*

Given a D-SLD that is blocked in phase I and a set of D- SLDs that have been scheduled, but have not been in service, the problem is to determine how to re-provision those lightpaths. The objective is to provision the blocked D-SLD without blocking any of the already scheduled D-SLDs and without changing their schedules.

- *Re-optimization at kick-off*

Given the current time slot t_c and a set of scheduled D- SLDs, the problem is to determine how to re-optimize those scheduled lightpaths whose starting time is equal to and each candidate path of the k-shortest paths (in step 4). Then, it uses SFF to find a wavelength (in step 5) and inserts the solution, i.e., the combination of a path, a wavelength and a starting time, into a solution list, L (in step 6). If the solution list is empty after looping through all the possible starting time slots and candidate paths, the demand is blocked (in step 9). Otherwise, it schedules the demand by using the solution that has the minimum objective value (in step 10). The objective values are calculated using either the MWL objective or the LB objective, whose expressions are shown below.

- *MWL*: the number of links on the path

- *LB*: $MAX u [x]$

$$t \square x \square t \square \square 1, e \square P$$

IV. LIGHTPATH RE-OPTIMIZATION

Lightpath re-optimization re-provisions the already scheduled D-SLDs to reduce blocking probability or minimizing network resource consumption. Due to the nature of the advance reservation, there exist lightpaths that have been scheduled, but are not in service. Although we cannot change the starting time of these lightpaths, their routing and wavelength assignment can be altered without interrupting their services.

We consider two scenarios of re-optimization, namely re-optimization at blocking and re-optimization at kick-off. Because the re-optimization problem is NP-Complete [10], we develop heuristic algorithms to solve the problems with respect to these two re-optimization scenarios. Our re-optimization heuristics operate on the set of time-overlapped lightpath sets. In this section, we first introduce the concept of time-overlapped lightpath set. We then present the heuristic algorithms for re-optimization at blocking and re-optimization at kick-off respectively.

A. *Time-overlapped lightpath set*

Given two lightpaths whose starting times are t_1 and t_2 and whose durations are d_1 and d_2 respectively, they are said to be

blocking in phase I, the algorithm re-provisions the scheduled lightpaths in the system with the hope of accommodating the blocked D-SLD. The algorithm operates on the time-overlapped lightpath set that overlaps with the blocked D-SLD. Except for the blocked lightpath, all the other lightpaths in the set are in the *scheduled* state. Figure 3 presents the algorithm for re-optimization at blocking for only time-fixed D-SLD. This algorithm is later used as building block of the algorithm for time-window D-SLD, which will be introduced later in this section.

The algorithm for re-optimization at blocking for a time-window D-SLD simply loops through all the possible starting times for the blocked D-SLD using the re-optimization time-overlapped if and only if

$$t_1 \leq t_2 \leq t_1 + d_1 \leq 1.$$

Deterministic Lightpath Scheduling Algorithm

Input: A D-SLD request (s, d, t, τ, l) and $G(V, E)$

Output: A schedule of the demand or refusal

1. **empty** the solution list L
2. **compute** the k -shortest path set $\{P_1, P_2, \dots, P_K\}$ between s and d with path length no greater than l .
3. **for** $t = t_{lower}$ to $t = t_{upper}$ (t_{lower} and t_{upper} are time window bounds)
4. **for each** path $P \in \{P_1, P_2, \dots, P_K\}$
5. use *SFF* as wavelength assignment scheme to find an available wavelength w on the path P
6. **if** a wavelength is found, insert (P, w, t) as a solution into the solution list L
7. **end for**
8. **end for**
9. **if** the solution list L is empty, the demand is blocked.
10. **else** compute the objective value for each solution in the list L and select the one with minimized objective value.

$$t_2 \leq t_1 \leq t_2 + d_2 \leq 1 \text{ or}$$

algorithm for the time-fixed D-SLD. If for any starting time the algorithm for time-fixed D-SLD succeeds, the algorithm returns success. Otherwise, it returns failure.

The ordering scheme used in step 5 of the re-optimization algorithm in Figure 3 is crucial. Based on the work in [11], the chances of provisioning all lightpaths becomes higher if the network resources are first assigned to those lightpaths that are most difficult to be provisioned, e.g., the lightpath with



the maximum number of hops on the shortest path between its source and destination. However, the introduction of time slots complicates the ordering scheme since we have an extra dimension, the time line.

Figure 1: Deterministic lightpath scheduling algorithm.

To define the time-overlapped lightpath set, we need to construct a time-overlapped graph. Given a set of lightpaths, we create the time-overlapped graph for these lightpaths as follows: (a) each lightpath is mapped to a node in the graph; (b) an edge is added between two nodes if their corresponding lightpaths are time-overlapped.

In the constructed time-overlapped graph, each connected component exactly corresponds to a time-overlapped lightpath set. The routing and wavelength assignment of lightpaths in a time-overlapped lightpath set are independent of lightpaths in another set. Given a time-overlapped graph and a lightpath, we can use the breadth-first-search (BFS) to find the set of time-overlapped lightpaths for the given lightpaths.

B. Re-optimization at blocking

We develop a heuristic algorithm for performing the re-optimization at blocking scenario. Upon occurrence of a

Figure 3: Algorithm for re-optimization at blocking for time-fixed D-SLDs.

In our approach, we use a combination of multiple keys for ordering. Given a set of scheduled lightpaths, we first sort them according to the non-decreasing order of their starting time. If two lightpaths have the same starting times, we break the tie by sorting them according to the non-decreasing order of the number of links on their minimum-hop paths. If they are still tied, we break the tie by sorting them according to the non-decreasing order of their service durations. Therefore, the lightpaths are sorted according to the rules of the earliest starting time first, maximum hop first, and longest duration first with decreasing priorities.

C. Re-optimization at kick-off

Re-optimization at kick-off provides a last chance to adjust the scheduled lightpaths before they are set up in service. At the beginning of each time slot, the re-optimization at kick-off algorithm collects the set of the lightpaths that are in the same time-overlapped lightpath set as each lightpath starting at the next time slot. It then releases all those scheduled lightpaths and re-provisions them according to the order discussed in the previous section. If any of the re-provisioning fails or the objective value is not reduced, it restores the re-provisioned lightpaths and returns failure. Otherwise, it returns success.

V. NUMERICAL RESULTS

In this section, we present numerical results for our proposed approach to dynamic and deterministic lightpath scheduling. We conduct our experiments on US-NET which has 24 nodes and 86 unidirectional links. A link in the topology represents two unidirectional links on each direction with the same length. Each link in the network has W wavelengths, where W

$\square \{8, 16, 32, 64\}$.

We simulate the arrivals of D-SLDs consisting of a mix ratio of *time-fixed* and *time-window* demands of 7:3 because we expect more time-fixed demands in the network. We set the duration of a time slot to 15 minutes. The window size of the time-window demands is set uniformly in a range of 4 to 48 time slots. We assume that the inter-arrival time between D-SLDs is exponentially distributed with mean \square . The duration between the D-SLD arrival time and the lightpath starting time is also exponentially distributed with mean of a number of 100 time slots. The duration of D-SLDs is measured in number of time slots and is a weighted uniform distribution in the range [1, 50]. For 50% of demands, the duration varies uniformly in the range [1, 10], and for 25%, 10%, 10%, and 5% of the demands the duration varies in the range of [11-20], [21, 30], [31, 40] and [41, 50] respectively. We set 10 as the value of k for computing the *k-shortest* paths. The path length constraint for each demand is set to 600 km, which is considered as a typical reach distance of all-optical signals using today's amplification and dispersion compensation technologies.

We use two metrics for comparison, blocking probability (*bp*) and service blocking probability (*sbp*). The blocking probability is calculated as the ratio of the number of blocked D-SLDs to the total number of D-SLDs. The service blocking probability is measured as the ratio of the sum of the durations of blocked D-SLDs to the sum of the durations of all the D-SLDs. Because different D-SLDs may have different durations, service blocking probability provides a fair measurement on the network performance.

A. *With and without re-optimization at blocking*

In this subsection, we compare the network performance with and without performing re-optimization at blocking. In the experiment, we dynamically establish D-SLDs one by one using the deterministic lightpath scheduling algorithm with load-balancing (LB) as its objective. The algorithm for re-optimization at blocking is invoked after each blocking. The simulation stops when the total number of D-SLDs exceeds a hundred thousand.

Figure 5 and Figure 6 plot the blocking probability (*bp*) and the service blocking probability (*sbp*) vs. the demand inter-arrival mean in networks with 32 and 64 wavelengths respectively. Each figure shows results obtained with and without performing the re-optimization (*reopt*) at blocking scenario. As shown in the figures, the re-optimization at blocking scenario reduces both the blocking probability and the service blocking probability significantly. We can also observe that this reduction is consistent as the network load changes. However, as the network load increase, the improvement obtained by the re-optimization degrades slightly. This is due to the fact that there is less room left for running re-optimization when the network becomes congested.

Table 1 summarizes the average performance of re-optimization at blocking in our experiments on US-NET. We observe that on average 49%, 58%, 58%, and 54% of blocked D-SLDs is eliminated by running re-optimization at blocking. We also observe that the performance improvement in terms of service blocking probability remains at the same level as blocking probability. The average number of time-overlapped lightpaths involved in each re-optimization is one of the factors affecting the scale of the problem. As the number of lightpaths per run increases, the running time becomes longer. However, the data in the last column show that our re-optimization algorithm can always find solutions in a short time.

B. With and without re-optimization at kickoff

In this subsection, we compare the network resource consumption values obtained with and without performing the re-optimization at kickoff. In the experiment, each D-SLD is scheduled with the objective of minimizing the number of used wavelength-links (MWL) in the deterministic lightpath scheduling phase. The re-optimization at kickoff algorithm is invoked at the beginning of each time slot, also with the MWL objective. We stop the experiment after generating ten thousand demands. We measure the number of used wavelength-links saved after each re-optimization invocation.

Table 2 shows the performance of the algorithm for re-optimization at kickoff in terms of the average saved wavelength-links per run. The first column (W) of the table is the number of wavelengths on each link in the experiment. The second column presents the average number of used wavelength-links that are saved after each invocation of re-optimization at kickoff. The third column is the percentage of saved wavelength links over the total number of wavelength links. It is followed by the average number of D-SLDs operated in each re-optimization invocation and the average running time of re-optimization. On average, we observe 5% to 7% of the total number of wavelengths is saved by running re-optimization at kickoff. Those saved wavelengths may be used to provision more connections in the future. Additional results from our simulation experiments are available in [12].




Figure 5 Results with and without re-optimization at blocking in US-NET with 32 wavelengths. scheduling phase, we employed heuristic algorithms to provision both time-fixed and time-window D-SLDs. In the second phase, lightpath re-optimization phase, we gave those scheduled D-SLDs that have not been placed in service a second chance to be re-provisioned based on the latest network state. We trigger the re-optimization phase based on two scenarios, namely re-optimization at blocking and re-optimization at kickoff. Simulation experiments showed that in the first scenario about 30%-60% of blocking can be eliminated and that in the second scenario a number of used wavelength-links can be reduced. This is a great performance improvement for telecom carriers that accommodate network users who require both dynamic and deterministic scheduled lightpath services.

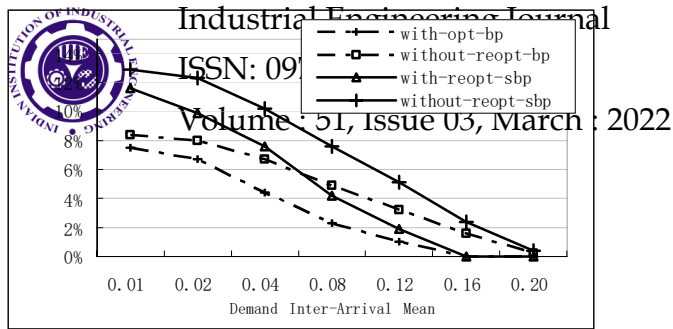


Figure 6 Results with and without re-optimization at blocking US-NET with 64 wavelengths.

Table 1 Performance of re-optimization at blocking in US-NET

# of Wavelengths	Average Improvement		Average # of Lightpaths per Run	Average Running Time (seconds)
	blocking probability	service blocking probability		
8	49.8%	51.8%	263	0.06
16	58.9%	59.9%	640	0.37
32	58.8%	59.1%	1429	1.10
64	54.7%	51.8%	3283	3.89

Table 2 Performance of re-optimization at kickoff in US-NET

VI. CONCLUSIONS

REFERENCES

- [1]. J. Kuri, N. Puech, M. Gagnaire, E. Dotaro, R. Douville, "Routing and wavelength assignment of scheduled lightpath demands", *Journal of Selected Areas in Communications (JSAC)*, vol. 21, no. 8, pp. 1231-1240, Oct. 2003.
- [2]. W. Su and G. Sasaki, "Scheduling of periodic transfers with flexibility," *41st Annual Allerton Conference on Communication, Control, and Computing*, Monticello, IL, Oct. 2003.
- [3]. D. Wischnik and A. Greenberg, "Admission control for booking ahead shared resources," in *Proc. of IEEE INFOCOM '98*, San Francisco, CA, USA, Mar. 1998.
- [4]. A. Schill, S. Kuhn and F. Breiter, "Resource reservation in advance in heterogeneous networks with partial ATM infrastructures," in *Proc. of the IEEE INFOCOM '97*, pp. 612-619, Kobe, Japan, Apr. 1997.
- [5]. R. Guerin and A. Orda, "Networks with advance reservations: The routing perspective," in *Proc. of the IEEE INFOCOM '00*, pp. 118 – 127, Tel-Aviv, Israel, Mar. 2000.
- [6]. A. Maach and A. Hafid, "Resources reservation in advance in slotted optical networks," in *Proc. of the Third International Conference on Optical Communications and Networks (ICOON 2004)*, Hong Kong, China, Nov. 2004.
- [7]. J. Kuri, N. Puech, M. Gagnaire, and E. Dotaro, "Routing foreseeable lightpath demands using a tabu search metaheuristic," in *Proc. of IEEE GLOBECOM 2002*, Taipei, Taiwan, Nov. 2002.

Wavelengths	Average # of links saved	Percentage of links saved	Average # of invocation	Average time (seconds)
8	36	5.2%	197	0.03
16	60	4.3%	381	0.08
32	166	6.0%	1025	0.36
64	398	7.2%	2194	1.42



[10]. E. Bouillet, J-F. Labourdette, R. Ramamurthy and S. Chaudhuri, "Lightpath re-optimization in mesh optical networks," *IEEE/ACM Transactions on Networking*, Apr. 2005.

In this paper, we studied the problem of scheduling and provisioning dynamic advance lightpath reservation requests known as D-SLD's in WDM networks. We consider two types of D-SLDs, namely, *time-fixed* where the starting time is fixed and *time-window* where the starting time can vary in a range. We proposed an efficient two-phase lightpath scheduling approach to schedule and provision the D-SLD requests dynamically. In the first phase, deterministic lightpath

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

[12]. L. Shen, Resource Allocation in Wavelength-Routed WDM Mesh Networks, PhD Dissertation, University of Nebraska- Lincoln, August 2006.