

Advance Reservation of Lightpaths in Optical-Network Based Grids

Silvia Figueira^{*}, Neena Kaushik^{*}, Sumit Naiksatam^{*},
Stephen A. Chiappari⁺, and Nirdosh Bhatnagar^{*}

^{*}Department of Computer Engineering, ⁺Department of Applied Mathematics
Santa Clara University
Santa Clara, CA 95053-0566
{sfigueira, nrkaushik, snaiksatam, schiappari, nbhatnagar}@scu.edu

Abstract - Advance reservation is an important part of grid computing. It guarantees that resources will be available at a pre-determined time to participate in the execution of a grid application. The advance reservation and scheduling of optical networks are key in guaranteeing that enough bandwidth, i.e., lightpaths, will be available for an application during a pre-determined interval of time. This paper studies the behavior of two algorithms developed for scheduling multiple lightpaths requested by advance reservation. To assess the benefit of each algorithm, we compare the blocking probability introduced by each of them. The blocking probability is obtained by simulating their behavior on different topologies. This simulation is based on traces of requests generated by FONTS, our Flexible Optical Network Traffic Simulator, which provides on-demand and advance-reservation requests with different characteristics.

I. INTRODUCTION

Grid computing involves the simultaneous usage of different resources connected by network links. This characteristic of grid computing mandates that all the resources involved in the execution of a grid application be available at pre-determined times. For example, if a surgery is to be performed remotely by a surgeon located in a different place, it is crucial that the network be available, with enough bandwidth, at the time for which the surgery is scheduled. To enable this, it is important to have an advance reservation scheduler, which guarantees the availability of the network. If not, there is a chance that the network may not be available at the right time.

In optical networks, an important technology used currently in data-intensive grid computing [5], the advance reservation of lightpaths is key to guaranteeing that enough wavelengths will be available when needed. Generally, a lightpath is defined as a wavelength data channel, possibly linking multiple optical segments, with either the same wavelength throughout the path or a concatenation of different wavelengths after undergoing wavelength conversion at intermediate optical nodes.

The importance of advance reservation for grid computing has been introduced by the Globus group in the definition of GARA [6] and then in [7]. Recently, it has been discussed in [8], which focuses on data-intensive collaboration, and in [17], which aims at scheduling data placement activities. The

scheduling of resources is basically a resource management problem, which has been discussed thoroughly in [22].

This paper focuses on the advance reservation of multiple lightpaths. The usage of multiple lightpaths for communication is important when dealing with the massive amounts of data used and/or produced by data-intensive grid applications. This fact is recognized, for example, in [18]. When multiple lightpaths are requested, the wavelengths assigned may be in the same path or not. When more than one path is available between the source and destination, the strategy used to select the wavelengths to form requested lightpaths may either balance the wavelength assignment throughout the available edge-disjoint paths or concentrate them in the shortest path first followed by the alternate edge-disjoint paths.

It is a well known fact that the strategy used in the assignment of wavelengths to paths directly affects the blocking probability of the network. For on-demand requests for a single wavelength, the blocking probability of the underlying optical network, as a theoretical model, has been studied in [31, 32]. In addition, it has been shown in [21, 28, 29] that allocating the wavelengths by packing them, compared to using the first-fit, random, or spread approaches, reduces the blocking probability of the optical network. In [9], the authors show that, in networks in which the usage of alternate paths is allowed, the blocking probability is reduced, in particular if an alternate path with a lighter load is used. Note that alternate path routing has also been studied in voice-based telephone, ATM, and optical networks. In [11], the authors show that the alternate path should be a lightly loaded one, and that better results are obtained when weights, which are proportional to the loads, are assigned to the alternate paths and used to select the alternate path. In [10], the authors study the behavior of algorithms with first-policy iteration, such as *basic*, *pack*, *spread*, *porder*, *pcolor*, and *lpcolor*, in a network in Finland.

This paper studies the behavior of two algorithms developed for scheduling multiple lightpaths requested by advance reservation. One of the algorithms is based on balancing the wavelengths, while the other is based on concentrating them. These algorithms, upon request of a number of lightpaths, assign a one-wavelength channel to each lightpath. To assess the benefit of each algorithm, we compare the blocking probability imposed by each of them. The blocking probability measures

the number of requests denied due to unavailability of resources. We obtain their blocking probability by simulating their behavior with a sequence of advance reservation requests.

While WDM technology is being increasingly adopted by the grid community, most of it is still being deployed or is in the experimental stage, as in [3, 27]. To validate the efficacy of the scheduling algorithms developed, it is necessary to exercise them with a large number of user requests. However, the pattern of user traffic in the grid is significantly different from that of an on-demand network. Reservation requests are not only a function of the time at which they arrive, but also of the time in the future for which the reservation is requested. Given the novelty of our approach, there is no history of user request traces that we can use for our experiments. In the absence of a real usable network and users, it becomes necessary to simulate the user requests. We have developed FONTS [23], a Flexible Optical Network Traffic Simulator, which can be used to generate a variety of user request patterns. The requests generated by FONTS can be for on-demand and/or advance reservation.

This paper is organized as follows: Section II discusses the modeling of advance reservation requests, Section III presents FONTS, Section IV presents the two algorithms for wavelength assignment, Section V presents the experiments performed to compare the strategies presented, and Section VI concludes and discusses future work.

II. MODELING ADVANCE RESERVATIONS

While extensive research has been done on modeling the Internet traffic, there is hardly any information available on modeling of bulk data transfer within the grid infrastructure. It is being widely accepted that Internet traffic is self-similar in nature [4, 19, 33]. We identify certain fundamental differences between bulk data transfer on grids and the Internet traffic. Unlike on the Internet, where the majority of the traffic is bursty, bulk data transfer is sustained. Heavy-tailed distributions are being used to model the volume of Internet traffic, but bulk data transfer seems to be more uniform. Traffic on the Internet is believed to have long range dependence, but bulk data transfer requests are fairly independent. Based on these differences it can be hypothesized that bulk data transfer traffic in grids is closer to voice traffic in telecommunication networks than it is to the data traffic in the Internet. Moreover, voice calls require the setup of dedicated circuits, which is similar to setting up lightpaths in the optical networks. Also, the queuing and congestion effects of the packet-switched Internet are absent in connection-oriented optical networks. The above is a speculative analysis and, in the absence of conclusive evidence, it is difficult to determine which model will fit bulk data transfer in grids. It could well be possible that bulk data traffic in grids is, or will be, a combination of different stochastic processes.

We have developed an advance reservation request model, in which the user is allowed to reserve dedicated lightpaths for a time slot in the future. In a system with advance reservation, once a user makes a request, the system evaluates if it can be honored. The system then checks to see if the required resources are free to be allocated during the requested time slot. If they are available, the user is granted the time slot and receives a confirmation handle, which is used at the time of the actual data transfer. Once the reservation request is accepted, the user is guaranteed access to utilize the requested lightpaths in that time slot. If the resources are not available to satisfy the request, a failed status is returned to the user, in response to which the user can make a new request. In addition to the number of lightpaths, the user is expected to specify a number of other parameters.

Based on this model, the following variables are needed to accurately simulate an advance reservation request:

- Request arrival time: This is the time at which an advance reservation request arrives in the system. The request arrival time determines the number of requests arriving in any given time interval. For our current experiments, we use a Poisson distribution [20] to model the number of requests in an interval.
- Source node: The node at which the data transfer originates.
- Destination node: This is the intended target for the data transfer. In order to reach this node, the lightpaths may have to be established through other intermediate nodes.
- Size of data to transfer: This is the total number of bytes that are requested to be transferred. The cost of setting lightpaths is amortized when the size of data to be transferred is larger than certain threshold values [5]. Hence our reservation and scheduling scheme focuses on data sets of the order of Tera bytes.
- Number of lightpaths requested: The simplest mode of data transfer is to stream a file over a single dedicated channel. But at times it is also possible to stripe data sets into multiple streams [12]. In such cases, the user may request multiple lightpaths simultaneously.
- Time at which data transfer is initiated: This corresponds to the start time of the advance reservation time slot that is requested by the user. One can intuitively think of a traffic pattern in the future to be a reflection of the current pattern shifted in time. Since we speculate that the arrival of on-demand bulk data-transfer requests follow a Poisson distribution, we also model the number of advance reservation requests for future time slots as Poisson.
- Number of time slots: This is not specified by the user, and is calculated according to the size of the time slot in the system and the size of the data that needs to be transferred.

In addition to the above, we identify a set of traffic variables which are relevant to bulk data transfers in an environment which supports advance reservations. In [23], we list these variables, their associated distributions, a model to generate

distributions for simulation experiments, and their ranges. The advance reservation request variables listed above are either obtained from the user or derived from these traffic variables.

The above advance reservation model follows principles that are similar to the ones presented in [25, 26], which focused on packet switched networks. To the best of our knowledge, no system has been designed to use an advance reservation model for dynamically provisioned optical networks. To study advance reservation and its effects on network scheduling, we need to analyze the characteristics of the variables presented above. Due to the absence of real advance-reservation traffic, we have developed a traffic simulator to generate traces of on-demand and advance-reservation requests for lightpaths.

III. THE FLEXIBLE OPTICAL NETWORK TRAFFIC SIMULATOR

FONTS exercises a combination of different candidate distributions and generates plausible traces of advance reservation and on-demand user requests. FONTS is flexible in that it can be configured to use different stochastic processes for the aforementioned advance reservation request variables. FONTS is implemented in C and generates an ascii based trace as the output of each simulation run. Simulation parameters can be controlled via a command line interface. Table I enumerates the different operation modes of FONTS. Every simulation run is a combination of exactly one distribution for each variable.

TABLE I
FONTS' OPERATION MODES

Variable	Distributions
Request Arrivals	Poisson
Source Node	Constant
	Uniform
	Arbitrary Probabilities
Destination Node	Constant
	Uniform
	Arbitrary Probabilities
Number of Wavelengths	Constant
	Uniform
	Heavy-Tailed
Data Size	Constant
	Uniform
	Arbitrary Probabilities
	Heavy-Tailed
Advance Reservation Start Times	Poisson
Number of Time Slots	Constant
	Variable

A FONTS trace consists of a sequence of advance reservation user requests for lightpaths to transfer a specified amount of data to a chosen destination at a desired time. Fig. 1 shows

a snapshot of a FONTS trace generated for Poisson arrivals, uniform source and destination models, constant number of wavelengths, and heavy-tailed distribution for data size.

Each request in the trace shown in Fig. 1 can be treated as a discrete event and is a combination of instantaneous values of the variables defined in Section II, based on the stochastic process chosen to model that variable. To generate random variables with specific distributions, we follow the Inverse Transformation technique as shown in [23] and explained in [14].

For each request, both the arrival start time and the reservation start time are modeled using exponential inter-arrival distributions. Each request sees an advance reservation window advertised by the system and can be made for any time slot within this window. Since all requests are for time slots in the future, we make an implicit assumption that there is no correlation between the time at which a request is made and the time for which the request is made. In order to simulate this request pattern we first generate the advance reservation start times using the Poisson distribution. Fig. 2 shows this process.

The advance reservation time window is divided into discrete time slots. The granularity of the advance reservation time slot is pre-defined and is fixed for a run of the simulation. A request with start time $x + y$, where x represents the beginning of a slot, and y is less than the length of the time slot, will be placed in the slot starting at x to allow the network to be available at time $x + y$. For example, if the granularity is set to 60 min., and an advance reservation request is generated with a start time 01:45 (request 3, in Fig. 2), then this reservation request is slotted in the interval between 01:00 and 02:00. Its start time is set to the lower edge of the interval at 01:00 (request 3, in Fig. 3), and the network will be available before 01:45. Thus, all requests falling in the interval from 01:00 to 02:00 are considered to be advance reservation requests with start time 01:00. Fig. 3 shows this re-alignment process.

A pool of such advance reservation start times is generated until we reach the end of the advance reservation window. Having done this, the system must attribute each advance reservation start time to a user request. User requests are also assumed to follow an exponential distribution for the inter-arrival time, and are generated from a different stream of uniform random numbers. As each user request is generated, we randomly pick an advance reservation start time from the pool of start times generated in Step 2. This process continues until all the start times are exhausted from the pool, at which point the advance reservation request traffic generation ends. Fig. 4 shows this process.

The net effect of this exercise is that it simulates independent and identically distributed user requests for advance reservation of time slots. These requested time slots for advance reservation are again independent and identically distributed within a time window in the future.

Seq#	ReqArrivalTime	SourceNode	DestNode	Wavelengths	Size(TB)	Res.Start	Res.End
0	Sun Jun 13 04:18:18 2004	192.26.85.52	192.26.85.36	1	1.32TB	Sat Jun 26 09:00:00 2004	Sat Jun 26 09:20:00 2004
1	Sun Jun 13 06:15:18 2004	192.26.85.5	192.26.85.36	1	2.41TB	Sat Jun 26 11:00:00 2004	Sat Jun 26 11:36:00 2004
2	Sun Jun 13 06:35:18 2004	192.26.85.20	192.26.85.5	1	1.08TB	Sat Jun 26 16:00:00 2004	Sat Jun 26 16:16:00 2004
3	Sun Jun 13 06:37:18 2004	192.26.85.52	192.26.85.20	1	1.09TB	Sun Jun 20 11:00:00 2004	Sun Jun 20 11:17:00 2004
4	Sun Jun 13 06:44:18 2004	192.26.85.5	192.26.85.52	1	10.54TB	Mon Jun 21 13:00:00 2004	Mon Jun 21 15:37:00 2004
5	Sun Jun 13 06:50:18 2004	192.26.85.52	192.26.85.36	1	1.82TB	Wed Jun 23 18:00:00 2004	Wed Jun 23 18:27:00 2004
6	Sun Jun 13 07:05:18 2004	192.26.85.5	192.26.85.20	1	4.34TB	Sat Jun 26 01:00:00 2004	Sat Jun 26 02:05:00 2004
7	Sun Jun 13 07:35:18 2004	192.26.85.52	192.26.85.36	1	4.66TB	Wed Jun 23 14:00:00 2004	Wed Jun 23 15:10:00 2004
8	Sun Jun 13 07:38:18 2004	192.26.85.5	192.26.85.52	1	1.07TB	Thu Jun 24 22:00:00 2004	Thu Jun 24 22:16:00 2004
9	Sun Jun 13 07:41:18 2004	192.26.85.20	192.26.85.5	1	1.34TB	Wed Jun 23 23:00:00 2004	Wed Jun 23 23:20:00 2004
10	Sun Jun 13 07:41:18 2004	192.26.85.5	192.26.85.52	1	1.16TB	Tue Jun 22 13:00:00 2004	Tue Jun 22 13:18:00 2004
11							
12							

Fig. 1. Snapshot of a FONTS trace.

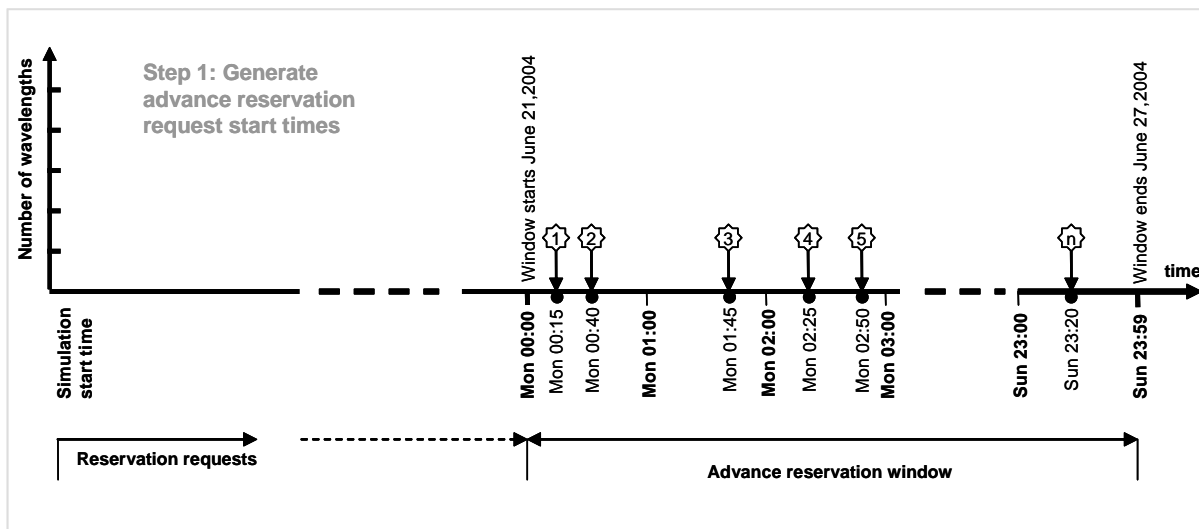


Fig. 2. Generate start times for the advance reservation requests.




Fig. 3. Align advance reservation requests.

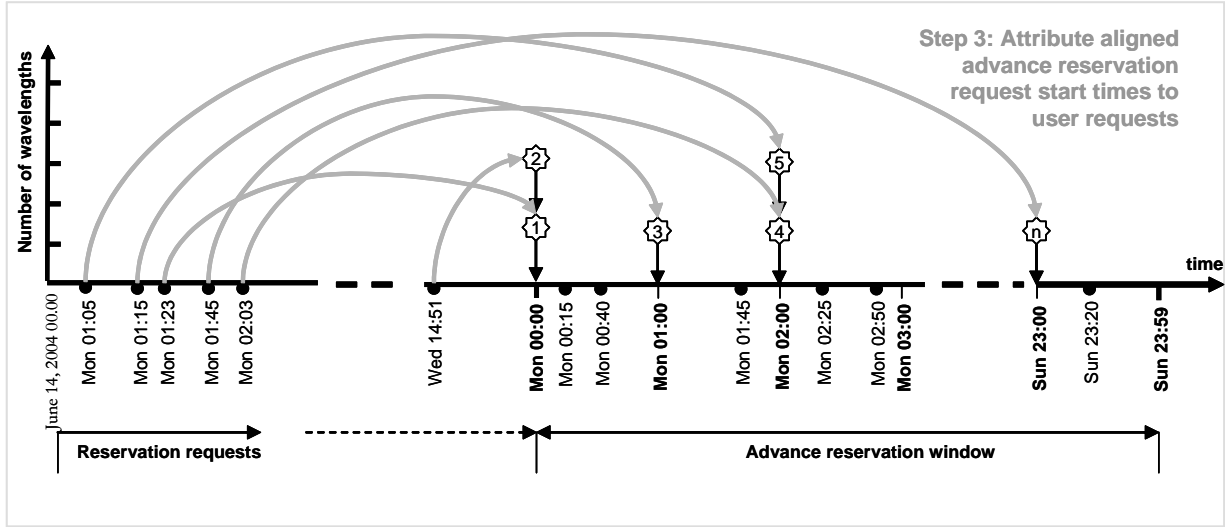


Fig. 4. Attribute advance reservation request start times to user requests.

It must also be noted that, on-demand user requests are just a special case of advance reservation requests. Based on its time of arrival, the on-demand request contends for the upcoming time slot. The difference is that, instead of aligning the request with the lower edge of the time slot, we align it with the upper edge of the time slot. For example, if an on-demand request arrives at 01:20 in Fig. 2, then it will be aligned with the time slot starting at 02:00 and ending at 03:00.

To validate the traces, the observed mean of the trace is compared with the theoretical mean of the Poisson distribution. In addition, the distribution of the advance reservation requests needs to be tested with the Poisson distribution for goodness-of-fit. This is achieved by using both the Chi-Square and the Kolmogorov-Smirnov tests as described in [2]. The tests show that the theoretical mean of the traces generated by FONTS is quite close to the observed mean, i.e., the traces passed both goodness-of-fit tests.

IV. ASSIGNING WAVELENGTHS TO LIGHTPATHS

The advance reservation and scheduling of lightpaths involves assigning wavelengths to requested lightpaths. The strategy used for this assignment should both minimize the blocking probability of the network and maximize the network utilization.

It is important to note that optical networks may have no wavelength converters, some converters, or all converters. The upper bound on the number of wavelengths needed in an optical network with no wavelength converters, with some wavelength converters, and with all wavelength converters is described in detail in [24]. Networks with no converters, such as the OMNInet [13], have the highest blocking probability due to the wavelength continuity constraint. In this case, to optimize the optical network throughput, i.e., to minimize the network blocking probability, it may be important to use an

efficient strategy to assign wavelengths to lightpaths. In [10], the authors deal with the allocation of one lightpath per request and show that, for these networks, *packing* by iterating over alternative paths leads to higher network utilization than *spreading* the wavelengths over the shortest path(s).

This paper addresses networks with no converters and requests for multi-wavelength lightpaths. When multiple lightpaths are requested, the wavelengths assigned may be in the same path or not. When multiple paths are available between the source and destination, the strategy used to select the wavelengths to form requested lightpaths may balance the wavelength assignment throughout the available paths (i.e., pack) or concentrate them in the shortest path(s) (i.e., spread). In this paper, we study the effects of balancing wavelengths by comparing two strategies for assigning wavelengths to lightpaths. One of the strategies balances the wavelengths, while the other concentrates them. Both algorithms are based on edge-disjoint paths. Note that, due to the wavelength continuity constraint, since we are assuming no wavelength converters, each of the requested lightpaths between the source and destination is formed using the same wavelength.

The wavelength-balancing algorithm allocates lightpaths in the following order: the first wavelength along the first edge-disjoint path, then the first wavelength along the second edge-disjoint path, and so on until the last edge-disjoint path. After that, it allocates the second wavelength along each edge-disjoint path in order, then the third wavelength along each edge-disjoint path in order, and so on until all the lightpaths are reserved. Note that, due to the order in which the allocation takes place, wavelengths are packed as much as possible. The algorithm is described below:

Wavelength-Balancing Algorithm:

```

Begin
  For i = 1 to number of wavelengths
    For j = 1 to number of edge-disjoint paths
      If wavelength[i] is available for all segments in
        edge-disjoint path[j]
        Allocate wavelength[i] for all segments in
          edge-disjoint path[j]
    If all requests of user are satisfied
      Print success
    Else
      Print the number of requests satisfied and number denied
  End

```

The wavelength-concentrating algorithm allocates the wavelengths along the paths in order, thereby utilizing all the wavelengths needed along the shortest path first followed by the alternate edge-disjoint paths. The algorithm is described below.

Wavelength-Concentrating Algorithm:

```

Begin
  For i = 1 to number of edge-disjoint paths
    For j = 1 to number of wavelengths
      If wavelength[j] is available for all segments in
        edge-disjoint path[i]
        Allocate wavelength[j] for all segments in
          edge-disjoint path[i]
    If all requests of user are satisfied
      Print success
    Else
      Print the number of requests satisfied and number denied
  End

```

Intuitively, balancing wavelengths should lead to schedules with better network usage, since it distributes the lightpaths along different paths, not blocking segments. This contributes to decreasing the probability of blocking due to wavelength continuity constraints. To illustrate, consider the OMNInet topology in Fig. 2 and the following requests:

- 4 lambdas from node 1 to node 3,
- 4 lambdas from node 2 to node 4,
- 1 lambda from node 1 to node 2, and
- 2 lambdas from node 2 to node 4.




Fig. 5. The OMNInet Topology

As shown in Table II and Table III, the wavelength-concentrating algorithm does a worse job and two requests are denied. The algorithm with wavelength-balancing is able to

accommodate all the requests. Note that OMNInet has a partial-mesh topology with four lambdas in each segment.

TABLE II
ALLOCATION WITH WAVELENGTH CONCENTRATING

Request path-xy from node x to node y	Number of wavelengths	Action
path-13	4	Allocate 4 wavelengths along <1,2,3>
path-24	4	Allocate 4 wavelengths along <2,4>
path-12	1	Not possible to allocate
path-24	2	Not possible to allocate

TABLE III
ALLOCATION WITH WAVELENGTH BALANCING

Request path-xy from node x to node y	Number of wavelengths	Action
path-13	4	Allocate 2 wavelengths (1st, 2nd) along <1,2,3> and 2 wavelengths (1st, 2nd) along <1,4,3>
path-24	4	Allocate 3 wavelengths (1st, 2nd, and 3rd) along <2,4> and 1 wave- length (3rd) along <2,1,4>
path-12	1	Allocate 1 wavelength (4th) along <1,2>
path-24	2	Allocate 1 wavelength (4th) along <2,4> and 1 wavelength (3rd) along <2,3,4>

V. EXPERIMENTAL ANALYSIS

We have written a scheduling simulator for calculating the blocking probabilities for the algorithms presented above. For each algorithm, the simulator accepts a FONTS trace, as shown above. It deals with one request for one time slot at a time, allocating lightpaths for each request, if possible. At the end of a one-week window, it calculates the blocking probability for each time slot, and then reports the average blocking probability.

We have run the simulator for several scenarios. In each case, requests are made for one week, which is split into 168 one-hour slots, unless otherwise specified. One week is the interval used, because the week is the smallest significant time interval in the calendar (hours and days are too small). We have experimented with three categories of traffic loads and three characteristics of wavelength requests, as shown in Table IV. We have also experimented with three different topologies. For each of these topologies, the edge-disjoint paths were pre-determined.

A. Topologies Studied

The behavior of the wavelength-balancing and wavelength-concentrating algorithms were analyzed in the topologies shown in Fig. 6. The 4-node partial-mesh is the topology used in the OMNInet. The 3-node-ring-with-a-spike topology introduces a concept of asymmetry in the network, which is useful in comparing the wavelength-balancing and wavelength-concentrating algorithms. The 4-node ring is basically a ring, a common configuration used in optical metro networks, which has been the focus of experiments with different algorithms [16]. These topologies comprehensively cover the 4-node based vertex graph for a variety of different features as mentioned above, and also resemble a practical network, the OMNInet, which motivates the use of a four-node graph in our first set of experiments.

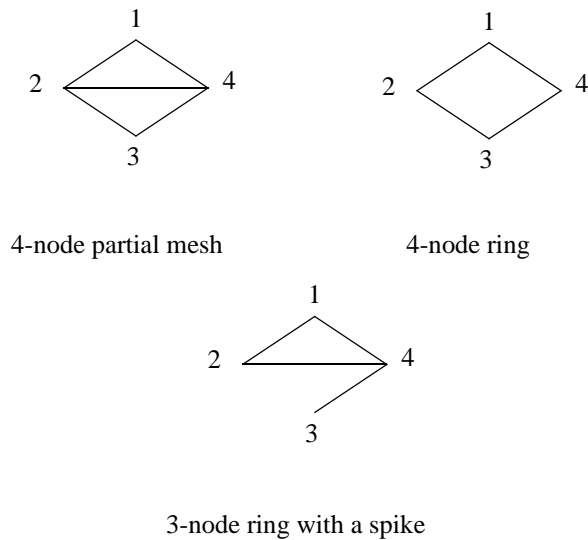


Fig. 6. Topologies Analyzed

For every link in the OMNInet, there are four wavelengths, which can be allocated in any direction. These links are therefore bi-directional with a capacity of four paths. The edge routers in case of OMNInet do have the capability of controlling in which direction, the communication will take place. For the simulations performed and the results presented in this paper, it has been assumed without loss of generality, that if the lightpath is allocated along one source-destination pair, it cannot be allocated in reverse.

B. Experiments

We have compared the algorithms proposed in Section IV by calculating their blocking probability. The blocking probability is calculated as the ratio of the total number of requests rejected to the total number of requests for that experiment, i.e.,

$$BlockingProbability = \frac{TotalRequestsRejected}{TotalRequests}.$$

The reservation requests target one-hour time slots within a week, and all requests are for advance reservations. If the total number of wavelengths in a request is greater than the maximum number of wavelengths that the network can handle for that time slot, then the request is blocked.

The graphs below compare the blocking probability obtained by the two scheduling algorithms presented in Section IV for the topologies defined in Subsection A. Each graph contains the blocking probability obtained by each algorithm in each experiment described in Table IV. In all the experiments, the inter-arrival time for the requests is ~10min. The inter-arrival time for the request reservation (i.e., start time) is different for each experiment and is specified in Table IV.

TABLE IV
SCENARIOS USED IN THE EXPERIMENTS

Number	Experiment	Reservation Inter-Arrival Time	Number of Wavelengths
1	High traffic with constant wavelength requests	5min	Constant: 1
2	High traffic with heavy-tailed wavelength requests	5min	Zipf's Generalized: exponent = 3 capacity = 4
3	High traffic with uniform wavelength requests	5min	Uniform: [1-4]
4	High traffic with constant wavelength requests	5min	Constant: 4
5	Medium traffic with constant wavelength requests	15min	Constant: 1
6	Medium traffic with heavy-tailed wavelength requests	15min	Zipf's Generalized: exponent = 3 capacity = 4
7	Medium traffic with uniform wavelength requests	15min	Uniform: [1-4]
8	Medium traffic with constant wavelength requests	15min	Constant: 4
9	Low traffic with constant wavelength requests	30min	Constant: 1
10	Low traffic with heavy-tailed wavelength requests	30min	Zipf's Generalized: exponent = 3 capacity = 4
11	Low traffic with uniform wavelength requests	30min	Uniform: [1-4]
12	Low traffic with constant wavelength requests	30min	Constant: 4

The graphs in Fig. 7, Fig. 8, and Fig. 9 represent the results obtained in our first set of experiments. Fig. 7 shows the blocking probability obtained in the partial mesh topology,

Fig. 8 shows the blocking probability obtained in the ring topology, and Fig. 9 shows the blocking probability obtained in the ring-with-one-spike topology. Note that the algorithms present a very close blocking probability in the three topologies studied, in all the 12 scenarios described in Table IV.




Fig. 7. Blocking probability for the partial-mesh topology.




Fig. 8. Blocking probability for the ring topology.

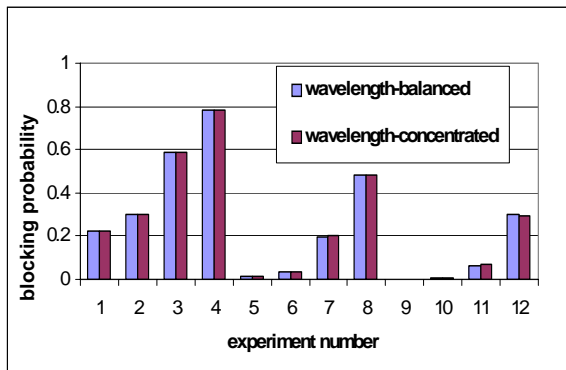


Fig. 9. Blocking probability for the ring-with-one-spike topology.

The graphs in Fig. 10, Fig. 11, and Fig. 12 represent the results obtained in our second set of experiments, which shows that the type of requests affects the efficiency of the algorithms. The graphs compare the blocking probability obtained by the two scheduling algorithms presented in Section IV for the partial mesh, when the percentage of 2-hop

requests (between nodes 1 and 3) is fixed. Fig. 10 shows the blocking probability obtained when 25% of the requests were for the 2-hop path, Fig. 11 shows the blocking probability obtained when 50% of the requests were for the 2-hop path, and Fig. 12 shows the blocking probability obtained when 75% of the requests were for the 2-hop path. In each case, the remaining requests are distributed evenly among the other paths. Note that the percentage of requests for the 2-hop light-path, between nodes 1 and 3, affects the behavior of the algorithms, since 2-hop requests use more links. The graphs indicate that decreasing the percentage of these requests favors the wavelength-balanced algorithm and that the probabilistic distribution of the distance between the end-points in the requests definitely affects the behavior of the algorithms.

The graph in Fig. 13 represents the results obtained in our third set of experiments. The graph compares the blocking probability obtained by the two scheduling algorithms presented in Section IV for the partial mesh, when there are no requests for the link between nodes 2 and 4. This simulates the existence of a redundant link, which is generally the case in network cores. Note that eliminating these requests affects the efficiency of the algorithms and favors the wavelength-balanced algorithm, since the link between nodes 2 and 4 can be used as an alternate path for other requests.




Fig. 10. 25% of the requests for the link between nodes 1 and 3.

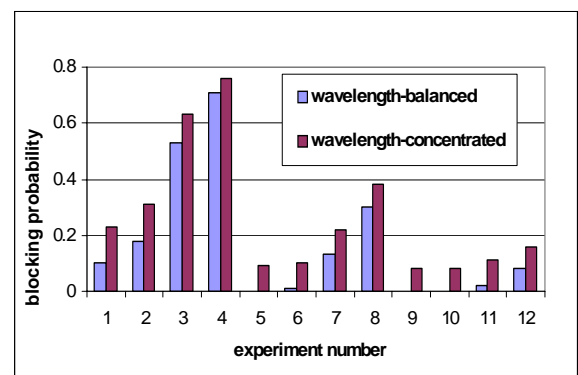


Fig. 11. 50% of the requests for the link between nodes 1 and 3.




Fig. 12. 75% of the requests for the link between nodes 1 and 3.

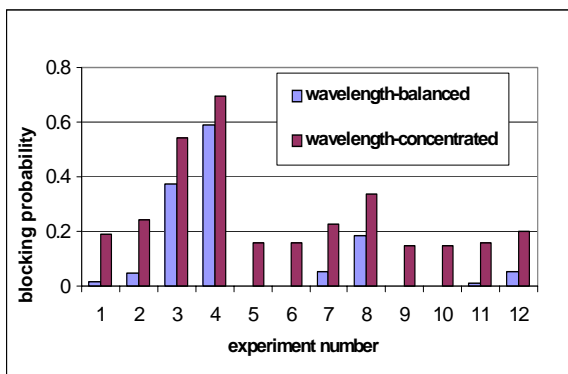


Fig. 13. Blocking probability when there are no requests for the 2-hop path.

The graph in Fig. 14 represents the results obtained in our fourth set of experiments. The graph compares the blocking probability obtained by the two scheduling algorithms presented in Section IV for the partial mesh, when the time slot used by the simulator is 30 minutes, instead of one hour. Note that decreasing the duration of the time slot affects the efficiency of the algorithms. With a shorter time slot, both algorithms obtain a lower blocking probability, since there are fewer requests per time slot, especially when the requests do not require the network for more than 30 minutes. In addition, a short time slot favors the wavelength-balanced algorithm, because fewer requests per time slot results in proportionally fewer requests both for the link between nodes 2 and 4 which, according to the numbers shown in Fig. 13, favors the wavelength-balanced algorithm, and for the link between 1 and 3 which, according to our second set of experiments, also favors the wavelength-balanced algorithm.

This leads to the conclusion that balancing wavelengths may be a better option in networks in which some of the links are not directly requested by the users, because these redundant links can be utilized in the alternate paths. Therefore, in networks in which all the links are requested uniformly, concentrating the wavelengths may be a better option, because there are no extra links to be used in the alternate paths, and using

alternate paths for one request may conflict with the request for another.




Fig. 14. Blocking probability with a 30-minute time slot.

For these experiments, the edge-disjoint paths were pre-determined. In [1], it has been shown that finding two edge disjoint paths for a fault-tolerant network is an NP-complete problem and, consequently, so is the problem of finding k -edge disjoint paths, where the paths are considered in a pair. As this project continues, and more sophisticated topologies are studied, it will be necessary to use techniques and algorithms, such as the ones presented in [15, 30] to find edge-disjoint paths.

VI. CONCLUSION

Advance reservation and scheduling of lightpaths in optical networks are essential ingredients in the recipe for a successful grid infrastructure. This paper explores the semantics of advance reservation requests and provides an analytical treatment for incorporating such requests in a user model. Although in simulation, this is to the best of our knowledge, the first attempt to study the behavior of advance reservation requests and their effects on network resource allocation in the context of dynamically provisioned optical networks. The authors have effectively laid down some criteria for further research work in evolving the advance reservation paradigm and quantifying the subsequent impact of optimization operations, such as scheduling of lightpaths. The discussion and results presented in this paper are generic in nature and are aimed at complementing the ongoing wave of research in creating service oriented grid stacks for dynamic network provisioning.

ACKNOWLEDGMENT

We would like to thank the DWDM-RAM group, in particular, Tal Lavian and Franco Travostino from Nortel Networks and Joe Mambretti from iCAIR and Northwestern University, for their support.

REFERENCES

- [1] J. Armitage, O. Crochat and J.-Y. Le Boudec, "Design of a Survivable WDM Photonic Network," in *Proceedings of the INFOCOM'97*, pp. 244-252, April 1997.
- [2] J. Banks, J. S. Carson, B. N. Nelson, "Discrete-Event System Simulation," Prentice Hall, pp. 375 - 378, 2nd edition, 1996.
- [3] DeFanti, T., Brown, M., Leigh, J., Yu, O., He, E., Mambretti, J., Lillethun, D., and Weinberger, J. "Optical switching middleware for the OptIPuter," in a *Special Issue on Photonic IP Network Technologies for Next-Generation Broadband Access*, *IEICE Transactions in Communications*, E86-B, 8 (Aug. 2003), 2263--2272.
- [4] S. Dixit, Editor, "IP over WDM: Building the Next Generation Optical Internet," John Wiley & Sons, Ltd., New Jersey, 2003.
- [5] S. Figueira, S. Naiksatam, H. Cohen, D. Cutrell, P. Dasput, D. Gutierrez, D. B. Hoang, T. Lavian, J. Mambretti, S. Merrill, F. Travostino, "DWDM-RAM: Enabling Grid Services with Dynamic Optical Networks," *IEEE CCGRID/GAN 2004, Workshop on Grid and Advanced Networks*, April 2004.
- [6] I. Foster, C. Kesselman, C. Lee, R. Lindell, K. Nahrstedt, and A. Roy, "A Distributed Resource Management Architecture that Supports Advance Reservations and Co-Allocation," *International Workshop on Quality of Service*, 1999.
- [7] I. Foster, A. Roy, and V. Sander, "A Quality of Service Architecture that Combines Resource Reservation and Application Adaptation," in *8th International Workshop on Quality of Service (IWQOS 2000)*, pp. 181-188, June 2000.
- [8] I. Foster, J. Vockler, M. Wilde, and Y. Zhao, "The Virtual Data Grid: A New Model and Architecture for Data-Intensive Collaboration," *Proceedings of the First CIDR - Biennial Conference on Innovative Data Systems Research*, January 2003.
- [9] H. Harai, M. Murata and H. Miyahara, "Performance of Alternate Routing Methods in All-Optical Switching Networks", in *Proceedings of the IEEE INFOCOM*, 1997, pp. 516-524, 1997.
- [10] E. Hyytia and J. Virtamo, "Dynamic Routing and Wavelength Assignment Using First Policy Iteration," *Fifth IEEE Symposium on Computers and Communications (ISCC'2000)*, July 2000.
- [11] C-F. Hsu and J.Y. Hui, "Load Balanced K-Shortest Path Routing for Circuit-Switched Networks", in *Proceedings of IEEE NY/NJ Regional Control Conference*, August 1994.
- [12] <http://www.globus.org/datagrid/gridftp.html>
- [13] <http://www.icair.org/omninet>
- [14] R. Jain, "The Art of Computer Systems Performance Analysis," John Wiley & Sons, pp. 496 - 498, 1991.
- [15] J. Kleinberg, "Approximation algorithms for disjoint paths problems", Ph.D. Thesis, MIT, Cambridge, USA, May 1996.
- [16] M. Knoke and H.L. Hartmann, "Dynamic Wavelength Assignment Schemes for Wavelength Routed WDM Ring Transport Networks", in *IEEE Global Telecommunications Conference*, vol. 5, 2003, pp. 2607-2611.
- [17] T. Kosar and M. Livny, "Stork: Making Data Placement a First Class Citizen in the Grid," in *Proceedings of 24th IEEE International Conference on Distributed Computing Systems (ICDCS2004)*, March 2004.
- [18] H. Lee, M. Veeraraghavan, H. Li and K.P. Chong, "Lambda scheduling algorithm for file transfers on high-speed optical circuits", *IEEE International Symposium on Cluster Computing and the Grid (CCGrid 2004)*, April 19-22, 2004, Chicago, Illinois, USA. <http://www.ccgrid.org/ccgrid2004>.
- [19] K. H. Liu, "IP over WDM," John Wiley & Sons, Ltd., England, 2002.
- [20] B. B. Lowekamp, B. Tierney, L. Cottrell, R. Hughes-Jones, T. Kielmann, and M. Swany, "A hierarchy of network measurements for grid applications and services," *draft at <http://www.wdidc.lbl.gov/NMWG>*.
- [21] A. Mokhtar and M. Azizoglu, "Adaptive Wavelength Routing in All-Optical Networks," *IEEE/ACM Transactions on Networking (TON)*, vol. 6, no. 2, pp. 197-206, April 1998.
- [22] J. Nabrzyski, J. Schopf, and J. Weglarz, Editors, "Grid Resource Management," Kluwer Publishing, 2003.
- [23] S. Naiksatam, S. Figueira, S. Chiappari, and N. Bhatnagar, "Modeling Advanced Reservation Requests in Optical Network Grids," COEN/SCU Technical Report, July 2004.
- [24] R. Ramaswami and K. N. Sivarajan, "Optical Networks - A Practical Perspective," Chapter 8, Morgan Kaufmann, 1998.
- [25] O. Schelen and S. Pink, "Sharing Resources through Advance Reservation Agents," *Proceedings of the IFIP International Workshop on Quality of Service*, May 1997.
- [26] O. Schelen and S. Pink, "An Agent-based Architecture for Advance Reservations," *Proceedings of the IEEE Conference on Computer Networks*, November 1997.
- [27] L. Smarr, A. Chien, T. DeFanti, J. Leigh, P. Papadopoulos, "The OptIPuter," *Special Issue: Blueprint for the Future of High Performance Networking*, Communications of the ACM, vol. 46, no. 11, pp. 58-67, November 2003.
- [28] H. Zang, J. Jue, L. Sahasrabudhe, R. Ramamurthy, and B. Mukherjee, "Dynamic Lightpath Establishment in Wavelength-Routed WDM Networks," *IEEE Communications Magazine*, pp. 100-108, September 2001.
- [29] J. Zhang, K. Zhu, L. Sahasrabudhe, S.J.B. Yoo, and B. Mukherjee, "On the study of routing and wavelength assignment approaches for survivable wavelength-routed WDM mesh networks," in *Optical Networks Magazine*, vol. 4, no. 6, November/December 2003, SPIE/Kluwer.
- [30] L. Zhao, "Disjoint Paths in Random Graphs", Department of Mathematical Sciences Carnegie Mellon University Pittsburgh, PA 15213, May 20, 1999.
- [31] Y. Zhu, G.N. Rouskas and H.G. Perros, "Blocking in Wavelength Routing Networks, Part I: The Single Path Case", in *Proceedings on INFOCOM*. IEEE Press, Piscataway, NJ, 321-328.
- [32] Y. Zhu, G.N. Rouskas and H.G. Perros, "Blocking in Wavelength Routing Networks, Part II: Mesh Topologies", in *Proceedings of the Sixteenth International Teletraffic Congress (ITC 16)*, Elsevier, Amsterdam, The Netherlands, pp. 1321-1330.
- [33] M. Zukerman, T. D. Neame, and R. G. Addie, "Internet Traffic Modeling and Future Technology Implications," *IEEE INFOCOM 2003*, vol. 1, pp. 587 - 596, March 2003.