

Analyzing the Advance Reservation of Lightpaths in Lambda-Grids

Sumit Naiksatam¹, Silvia Figueira¹, Stephen A. Chiappari², Nirdosh Bhatnagar¹
¹Department of Computer Engineering, ²Department of Applied Mathematics,
 Santa Clara University, Santa Clara, CA 95053-0566
 {snaiksatam, sfigueira, schiappari, nbhatnagar}@scu.edu

Abstract

The scheme of advance reservations in dynamically provisioned optical networks is novel, and there are no grid-based applications designed to utilize this scheme. We formally define and analyze this scheme and present a constrained mathematical model for advance reservations. We also introduce FONTS - the Flexible Optical Network Traffic Simulator, a tool for simulating advance reservation, on-demand, and periodic data transfer requests. FONTS is based on a stochastic model and incorporates a variety of variables, which have been identified to accurately model advance reservation requests. FONTS validates the mathematical model and also helps to analyze complex scenarios beyond the scope of this paper.

1. Introduction

Lambda-Grids are virtual aggregations of geographically distributed computational elements tightly coupled with dynamic lambda circuits which provide dedicated multi-gigabit communication channels [18]. They support sophisticated e-science applications, which require a high performance computing environment including guaranteed and sustained high bandwidth [3]. One way to ensure bandwidth guarantees is to offer the service of reserved end-to-end (e2e) dedicated data channels. This has been recognized by the grid community. Older definitions of advance reservation have been updated to incorporate network bandwidth along with computational and storage elements [1].

More specifically, the Grid High Performance Networking draft on Network Services [2] proposes an independent Network Advance Reservation Service in its service definitions. This service would allow the user to negotiate and claim specific network resources for a requested period. Interestingly, lightpaths [6], which are the communication channels in wavelength division multiplexing (WDM) optical networks, can be reserved in advance. Specialized signaling protocol abstractions like UCLP [6], GridJIT [19]

and ODIN [14], and multi-domain policy-based mechanisms for access control like the Bandwidth on Demand service [12], have been devised to dynamically set up application controlled e2e lightpaths. In fact, e2e lightpath management has already been put into practice in the CA*net4 national research network [6], the National Lambda Rail [15], and has been proposed for SURFnet 6 [16].

Reservation of lightpaths has been discussed [10], but its wide scale utilization has not yet been attempted. Consequently, there is a pressing need to not only provide a tool which can be used to study the interplay of advance reservations with other strategies used for network resource optimization, but also to define the advance reservation scheme itself in the context of dynamically provisioned optical networks. The pattern of user requests in an advance reservation scheme is significantly different from the pattern in an on-demand network environment. Advance 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 reservations are requested. Our goal is to model these requests for systems that support advance reservations.

Recently, the work in reservations has focused on reservation protocols [8], and simulation analysis has been performed in the context of packet switched networks [17]. However, the basis for analysis is different for e2e lightpaths, which is essentially a connection-oriented setup. The mathematical analysis of advance reservations presented in [7] is close to this context, and we have chosen to extend and adapt it to the optical domain. Though not specific to traffic modeling, both the discussion on the relevance of advance reservations to routing and wavelength assignment in [23] and the general concepts on advance reservations in [21] are useful. It should also be noted that the reservation scheme presented here is targeted at bulk data transfer and complements resource management strategies in evolving technologies such as optical burst switching.

In [20], Vazhkudai et al. use log data from past transfers to predict future transfer times. Such log data collected from experimental testbeds is difficult to obtain, and the authors have no control over the data collection frequency,

which limits the scope of a comprehensive analysis.

This paper presents an advance reservation model and FONTS, the Flexible Optical Network Traffic Simulator, which can be used to generate traces of a range of user request patterns. Instead of having to rely on log data, one may use FONTS to generate traces for various scenarios. Multiple traces can be combined, with possibly different reservation modes, thus making the trace generation process extremely versatile, and adaptable to most traffic conditions.

The paper is organized as follows: Section 2 provides a mathematical analysis of advance reservations, Section 3 introduces FONTS, Section 4 dwells on experiments performed with FONTS, and Section 5 concludes.

2. Analytical Model

A mathematical model for the advance reservation scheme, which extends the analysis provided in [7], is proposed to facilitate the understanding of the different variables at play. It provides us with two useful metrics, *Rejection Probability* and *Reservation Probability*. The rejection probability is a measure of the capacity that the underlying network can handle. The reservation probability is a measure of the utilization of the network based on the pattern of user requests and can be used to define the appropriate number of network resources, which would keep the rejection probability down to an acceptable level. While the advance reservation scheme is not constrained to fixed network topologies, it is much easier to construct a tractable model for a single link connecting two network nodes. Table 1 defines the variables used in this model. The model is based on the following assumptions:

1. A reservation interval is a time slot of any duration and is an integral multiple of minutes.
2. Reservation requests arrive in a fixed time window called the Request Arrival Window.
3. Reservations target time slots in a fixed time window called the Reservation Window.
4. The request arrival window and the reservation arrival window may overlap but the reservation window cannot start before the request arrival window. To simplify the discussion, without loss of generality, we present only the non-overlapping case.
5. Each request targets reservation for only one time slot. This is another simplifying assumption. While it is possible that the time required to transfer the data in a particular request is more than the duration of one slot, this case complicates the analysis significantly. Therefore we assume that such requests can be split into two or more requests.
6. Reservation requests targeted at time slots within the reservation window are considered as reservation candidates, while those falling outside are discarded.
7. The entire time line (the request arrivals window and the reservations window) is divided into time slots of equal duration.
8. Reservation requests arrive in a stationary Poisson stream at rate λ , such that the N_m are independent and identically distributed (i.i.d.) random variables.
9. A request targets a slot based on an Advance Reservation distribution which is a sequence of discrete i.i.d. random variables independent of the request arrival distribution.
10. The number of requests, H_m^s , are i.i.d. random variables, with mean λa_{s-m} . Thus, $H^s = \sum_{m=1}^L H_m^s = \sum_{m=1}^L \lambda a_{s-m}$. This setup is illustrated in Figure 1.
11. A reservation request is accepted if there are enough wavelengths in the time slot to accommodate the request. Otherwise, it is rejected and lost. In the preliminary discussion, only a single wavelength is assumed available per time slot. Later, we relax this assumption and extend the model to multiple wavelengths.

It is clear from the above assumptions that, when analyzing the behavior of reservations, the effects of at least two independent distributions need to be considered, request arrivals and reservation start times. Note that this differs from the commonly used metric of blocking probability based on the *Erlang-B* formula, which takes into consideration only the effects of the request arrivals distribution and subse-

Table 1. Variables in the Reservation Model.

L	Request arrivals window size in integral multiples of one slot.
W	Reservation window size in integral multiples of one slot.
K	$L + W - 1$, last time slot that can be reserved in advance.
a_i	Probability that an arrival at time t wants to reserve a time slot during $[[t] + i, [t] + i + 1)$, where $0 \leq i \leq K - 1$.
N_m	Number of reservation request arrivals in the slot $[m - 1, m)$, $m = L - i$, where $1 \leq m \leq L$ and $0 \leq i \leq K - 1$.
λ	Avg. number of reservation request arrivals per slot, such that $E(N_m) = \lambda$, where E is the expectation operator.
P_{rej}^s	Rejection Probability for a time slot s , defined as the ratio of total number of reservation requests rejected to the total number of requests for the time slot $[[s], [s] + 1)$, and is calculated when reservation requests are made until time L , where $L \leq s \leq K$.
P_{res}^s	Reservation Probability for a time slot s , defined as the probability of getting one or more reservation requests for the time slot $[[s], [s] + 1)$, and is calculated when reservation requests are made until time L , where $L \leq s \leq K$.
H_m^s	Number of requests arriving in time slot $[m - 1, m)$, which target a reservation for the time slot $[[s], [s] + 1)$.
H^s	Number of requests by time L which target a reservation for the time slot $[[s], [s] + 1)$; note that $H^s = \sum_{m=1}^L H_m^s$.
Y	Number of wavelengths available for use on each link.




Figure 1. Calculating the number of requests targeting a reservation time slot.

quently leads to an analysis of the immediate reservation case. Also, the strategy adopted to accept a request among several competitors (call admission control) is not relevant to this model. The analysis provided here should serve as input to such an admission control policy, thus enabling optimal decisions.

2.1. Rejection Probability

In the case of a single wavelength, the rejection probability can be written as $P_{rej}^s = 1 - 1/H^s$. So,

$$P_{rej}^s = 1 - \frac{1}{\sum_{m=1}^L H_m^s} = 1 - \frac{1}{\sum_{m=1}^L \lambda a_{s-m}}. \quad (1)$$

Eq. (1) is valid for $\sum_{m=1}^L \lambda a_{s-m} \geq 1$. However, if the advance reservation distribution is assumed to be a uniform distribution, such that $a_i = 1/K$, for each i , then the expression for rejection probability can be simplified, and

$$\begin{aligned} P_{rej}^s &= 1 - \frac{1}{\sum_{m=1}^L \lambda/K} \\ &= 1 - \frac{1}{(\lambda L/K)} = 1 - \frac{L+W-1}{\lambda L}. \end{aligned} \quad (2)$$

2.2. Reservation Probability

The generating function for the random variable H^s can be written as $Ez^{H^s} = Ez \sum_{m=1}^L H_m^s$. Since H_m^s are i.i.d. variables, for fixed s , $Ez^{H^s} = \prod_{m=1}^L Ez^{H_m^s}$ and

$$\begin{aligned} Ez^{H_m^s} &= E[E[z^{H_m^s} | N_m]] \\ &= E\left[\sum_{j=0}^{N_m} \binom{N_m}{j} a_{s-m}^j (1 - a_{s-m})^{N_m-j} z^j\right] \\ &= E(1 - a_{s-m} + za_{s-m})^{N_m}. \end{aligned}$$

If $q_n = P(N_m = n)$, then

$$Ez^{N_m} = \sum_{n=0}^{\infty} (1 - a_{s-m} + za_{s-m})^n q_n, \text{ and hence}$$

$$Ez^{H^s} = \prod_{m=1}^L \sum_{n=0}^{\infty} (1 - a_{s-m} + za_{s-m})^n q_n. \quad (3)$$

Since the arrivals are assumed to be Poisson with the arrival parameter λ , we have $q_n = \exp(-\lambda)(\lambda^n/n!)$, and

$$\begin{aligned} Ez^{H^s} &= \prod_{m=1}^L \sum_{n=0}^{\infty} \frac{\exp(-\lambda)\lambda^n}{n!} (1 - a_{s-m} + za_{s-m})^n \\ &= \prod_{m=1}^L \exp(-\lambda a_{s-m}(1-z)) \\ &= \exp(-\lambda(1-z) \sum_{m=1}^L a_{s-m}). \end{aligned} \quad (4)$$

The reservation probability, P_{res}^s , of at least one reservation request, can be obtained by subtracting the first coefficient of (4), in its expansion in powers of z , from 1. So,

$$P_{res}^s = 1 - \exp(-\lambda \sum_{m=1}^L a_{s-m}). \quad (5)$$

If the advance reservation distribution is assumed to be an uniform distribution, $a_i = 1/K$, which can be used in (4) to obtain

$$Ez^{H^s} = \exp(-\lambda(1-z)L/K) = \exp(-\lambda(1-z)\frac{L}{K}). \quad (6)$$

Using (6), (5) can be rewritten as

$$P_{res}^s = 1 - \exp(-\lambda\frac{L}{K}). \quad (7)$$

It can be seen that, in the case of uniform distribution for advance reservation, the reservation probability in (7) is independent of any s terms, and is a constant for all the




Figure 2. Surface plot for number of slots in the reservation window for varying values of reservation probability and request arrivals.

time slots. This desired characteristic ensures that the resources are optimally utilized, which may not always be the case. Depending on the nature of the applications requesting advance reservations, certain periods of the day may be requested more than others (for instance, data backup operations are often done after midnight). In this case, the advance reservation distribution is going to be non-uniform and bursty. In general, the reservation probability depends on the distribution of the advance reservations. With a bursty distribution (or a non-stationary Poisson arrival process), the reservation probability has a large variance over different time slots. This leads to a fragmented utilization with alternating peak and empty usage intervals.

To calculate W , we rewrite Eq. (7) as $L = (K/\lambda) \ln[1/(1 - P_{res}^s)]$. Putting $X = (1/\lambda) \ln[1/(1 - P_{res}^s)]$ and rewriting K as $L + W - 1$ yields $L = (L + W - 1)X$, and

$$W = \frac{L(1 - X)}{X} + 1. \quad (8)$$

Since $W > 0$ and $L > 0$, $(1 - X)/X > 0$. Therefore (7) and (8) are valid for $0 < X < 1$. Figure 2 shows a surface plot for the value of W when plotted against λ , for different values of desired reservation probabilities. The surface indicates the upper bound on the values of W , such that values below the surface are all legal. If we have a reasonable estimate of the expected number of request arrivals per time slot, then we may use (8) to calculate the reservation window size for a desired level of utilization.

2.3. Multiple wavelengths

Using (7) and (2), we can express P_{rej}^s in terms of P_{res}^s :

$$P_{rej}^s = 1 - \frac{1}{\ln(1/(1 - P_{res}^s))}. \quad (9)$$




Figure 3. Plot of reservation probability vs rejection probability for a uniform distribution.

To reduce the rejection probability for higher values of reservation probability, it is necessary to provision more wavelengths. If Y wavelengths are available, then $P_{rej}^s = 1 - Y / (\sum_{m=1}^L \lambda a_{s-m})$, which, in the case of uniform distribution can be written as

$$P_{rej}^s = 1 - (YK/\lambda L). \quad (10)$$

To obtain the reservation probability for Y wavelengths, the first Y terms in the expansion of (4) have to be subtracted from 1 to obtain

$$P_{res}^s = 1 - \sum_{i=0}^{Y-1} \frac{(\lambda \sum_{m=1}^L a_{s-m})^i}{i!} \exp(-\lambda \sum_{m=1}^L a_{s-m}).$$

In the case of uniform distribution for advance reservations,

$$P_{res}^s = 1 - \sum_{i=0}^{Y-1} \frac{(\lambda L/K)^i}{i!} \exp(-\lambda L/K). \quad (11)$$

Using (10) and (11), the relation between the rejection and the reservation probability in the case of uniform distribution for advance reservations can be written as

$$P_{rej}^s = 1 - Y / \ln\left\{\left[\sum_{i=0}^{Y-1} \frac{(\lambda L/K)^i}{i!}\right] / (1 - P_{res}^s)\right\}. \quad (12)$$

Figure 3 shows the relationship defined in (12). As the number of available wavelengths increases, the rejection probability drops more rapidly as the reservation probability decreases. If the utilization (reservation probability) is known, then an appropriate value of Y can be picked to achieve an acceptable level of rejection probability.

The discussion above suggests that the reservation probability can be used as a pre-provisioning metric to plan the network capacity, and that the rejection probability can be

used as a post-provisioning metric to track the throughput of the network in terms of data transfer requests accepted.

As noted earlier, a bursty advance reservation distribution leads to non-uniform utilization. In such a case, a decision has to be made on the number of wavelengths that need to be provisioned. Choosing a number to satisfy the upper bound on the reservation probability will yield a lower rejection probability, but also a decreased utilization. Conversely, choosing a number to satisfy the lower bound on the reservation probability will yield a higher rejection probability but possibly an increased utilization. Thus, without accurate knowledge of either request arrival distribution and or advance reservation distribution, it is difficult to analytically predict this sweet spot for the number of wavelengths. Hence, for more practical considerations, performing simulations seems to be a more promising alternative. Nevertheless, the analysis of the rejection probability still turns out to be useful in specific cases, providing a precise lower bound. This is discussed in detail in section 4.

3. FONTS

In the previous section, a mathematical model for analyzing advance reservation requests was presented. This model is helpful for preliminary analysis, but is limited by the constraints imposed. Extending the model to relax any of the constraints and making it general enough to cover every possible traffic scenario has proved to be difficult. To obtain a more precise analysis in a generic setup, simulation is a natural choice. While many optical network simulators are available, none of them is equipped to handle advance reservation requests. More importantly, there are hardly any applications designed to reserve lightpaths in optical networks. Hence no information about the pattern of lightpath utilization is available. FONTS was born out of this need for synthetic traces.

FONTS can be operated to generate traces in one of three reservation modes: advance, on-demand (immediate) and periodic. The advance reservation mode is discussed in section 3.1. The on-demand mode generates requests for immediate reservations and is useful in the study of the effect of coexistence of on-demand and advance reservation modes. The periodic mode generates advance reservation requests and attributes a periodic cycle to each of them. The periodic mode is specifically designed to target certain use cases in the grid computing environment, such as the periodically scheduled movement of data within a hierarchy of storage levels as in the High-Energy Physics Applications [5]. Both the on-demand and periodic modes are highly configurable, and a detailed description can be found at [11].

Each FONTS trace consists of a sequence of requests for the reservation of lightpath(s) to transfer a specified amount

of data between two nodes at a desired time. A reservation request, discussed below, is a combination of instantaneous values of the variables, based on the stochastic process chosen to model that variable. Section 4 explains how FONTS helps in providing a more generic analysis of advance reservations where the mathematical model falls short.

3.1. Advance Reservation Requests

The following variables have been identified to accurately model an advance reservation request:

Request arrival time: This is the time at which an advance reservation request arrives in the system. FONTS employs a Poisson distribution to model the number of requests arriving in every time interval.

Source and Destination Nodes: The node selection follows a constant, uniform, or non-uniform (known probabilities) distribution. In order to reach the destination node, the lightpath may have to be established through intermediate nodes. In the non-uniform model, probabilities can be assigned to nodes. This is particularly useful in the case where the topology and behavior of a network are known. For example, in power-law networks, certain nodes that handle heavy volumes of traffic are identified as super nodes. Most of the traffic passes through these nodes. These nodes can then be assigned significantly higher probabilities than the others. Consequently, when user requests are simulated, more requests will originate and terminate at these nodes.

Amount of data to transfer: The data size follows a constant, uniform, heavy-tailed, or non-uniform (known probabilities) distribution. If most of the data sizes are small and a few are large (similar to the Internet traffic), then a heavy-tailed (Pareto) distribution can be used. If the frequently occurring data sizes and their probability of occurrences are known for a system, then the known-probabilities model can be used.

Number of lightpaths requested: This variable follows a constant, uniform or heavy-tailed (Generalized Zipf's) distribution. The simplest mode of data transfer is to stream a file over a single dedicated channel. It may be also possible to stripe data sets into multiple streams. In such cases, the user may request multiple lightpaths simultaneously.

Time at which data transfer is initiated: This variable corresponds to the start time of the advance reservation time slot, which is requested by the user. The reservation start times can be modeled using a uniform or a Poisson distribution. The uniform distribution provides the simulated behavior of the uniform distribution discussed in the analytical model before. When using the uniform distribution model, the advance reservation start times are calculated according to the request arrival time. When the explicit modeling of advance reservation start times is required, the Poisson model can be used. Another argument for using the Pois-




Figure 4. Generating advance reservation start times with Uniform distribution.

son model is that one can represent the traffic pattern in the future as the current pattern translated by time. Since the request arrivals follow a Poisson distribution, the number of advance reservation requests for future time slots can be modeled as Poisson. FONTS also supports multiple inter-arrivals time rates to model the heterogeneity in traffic at different times of the day.

Number of time slots: The available advance reservation window is divided into discrete time slots. Each advance reservation is aligned with the beginning of one slot. The user can then request one or more time slots depending both on the length of the time slot and on the size of the data that needs to be transferred.

3.2. Simulating Advance Reservation Requests

Figure 4 illustrates the simulation of advance reservation requests according to a uniform distribution. Every user request has an equal probability of selecting any time slot for reservation in the advance reservation window. To achieve this, the total time on the entire time line (request arrival window + advance reservation window) is calculated. Request arrivals are generated using a Poisson distribution with the specified arrival rate. For each request, a time is selected at random from the total time calculated before and added to the current request arrival time. If the resultant reservation time falls outside the advance reservation window, the request is discarded. Otherwise, the reservation time is aligned with the start of the time slot in which it occurs (step 1). This is regarded as the reservation start time and attributed to the request being processed (step 2).

Figure 5 shows how advance reservation requests are generated with a Poisson distribution. A reservation start time is selected to follow a Poisson distribution with the specified inter-arrival rate. As before, this reservation time is aligned with the start of the time slot in which it occurs. A pool of such advance reservation start times is generated until the end of the advance reservation window is reached. Then each advance reservation start time is attributed to a




Figure 5. Generating advance reservation start times with Poisson distribution.

user request. User requests are also assumed to follow an exponential distribution for the inter-arrival time. As each user request is generated, an advance reservation start time is randomly selected from the pool of start times generated in step 1. This process continues until all the start times are exhausted from the pool, at which point the generation of advance reservation request ends (step 2).

4. Simulation experiments using FONTS

To illustrate the power and flexibility of FONTS, we show how the rejection probability changes according to the traffic distribution in a 4-node network and how FONTS can help to analyze that effect. Interesting results are also observed with respect to the interplay of file size distribution and reservation probability. They are presented in a future publication due to the space constraints here.

This section extends the analysis of the rejection probability from section 2.3 by relaxing the two-node constraint imposed in the mathematical model. A fully interconnected 4-node network is considered here, but this analysis can be extended to any arbitrary topology (a fully interconnected network has the same number of paths between all the pairs of nodes which simplifies the analysis). The nodes are labeled 1, 2, 3 and 4, and the objective of the experiment is to find the rejection probability for the requests between a given pair of nodes, say 1 and 2. It is assumed that there is one wavelength available on each link. There are multiple paths between nodes 1 and 2. The availability of a particular path (in a time slot) depends not only on the request traffic between nodes 1 and 2, but also on the traffic between other combinations of source-destination pairs. For example, the path 1-3-2 can be used to reach from node 1 to 2, but if a previous request has already used the link from node 3 to 2, then the path 1-3-2 can no longer be used. As such, it will be interesting to obtain the rejection probability between nodes 1 and 2, with varying traffic loads between the other nodes. This requires traces of requests such that the traffic load between nodes 1 and 2 remains constant, and

the contribution of traffic load between the other pairs is progressively increased. These types of experiments can be effectively modeled using traces generated by FONTS. For example, we can choose five different test cases in which the traffic load between nodes 1 and 2 is respectively 100%, 80%, 60%, 40% and 20% of the traffic load between nodes in all other pairs.

Assume 20 reservation-request arrivals per 5-min time slot for data transfer between nodes 1 and 2. If the 80% case is considered, these 20 requests per 5 min need to be exactly 80% of all requests entering the network. Therefore, there are 25 reservation-request arrivals per 5 min, and 20 requests out of these are for transferring data between nodes 1 and 2. This can be achieved by generating two traces and then combining the two. The first trace is generated such that it has a uniform distribution of all source-destination node pairs. There are 6 possible source-destination pairs in a 4 node network (ignoring the order of nodes in a pair). Out of these, the 1-2 node pair is the one under investigation. So the request arrival rate between the other 5 pairs of nodes is $25 - 20 = 5$ requests per 5 min or 1 request per 5 min for each pair. However, this trace is generated with uniform distribution of source and destination nodes, and some of the requests are going to be for the 1-2 node pair. The effective request arrival rate for the 4 node trace is calculated as $5 + 1 = 6$ requests per 5 min. The second trace is generated exclusively for the 1-2 node pair. Since 1 request per 5 min of traffic load of the 1-2 node pair has already been generated in the 4 node trace, the effective arrival rate for this 2 node trace is $20 - 1 = 19$ requests per 5 min. The result of combining the two traces is that we obtain an average arrival rate of 25 requests per 5 min with approximately 20 requests per 5 min for the 1-2 node pair.

Given this scheme of generating traces of traffic, we use (10) to find a lower bound on the rejection probability. In the case of an n -node fully interconnected network, there are $n - 1$ shortest edge disjoint paths between any two pairs of nodes. If only one wavelength is available per link, at most a maximum of $n - 1$ reservation requests can be satisfied between any two nodes for any time slot. With respect to (10), these $n - 1$ paths can be treated as Y wavelengths in a two node network. Using this data, the value of P_{rej}^s in (10) can be calculated. For the n -node network, if all the traffic is only between one pair of nodes (say 1-2) then the observed rejection probability will be the same as P_{rej}^s in (10), calculated by making $Y = n - 1$. As the traffic between the other pairs of nodes increases, some of the paths between the 1-2 node pair are going to get used up for the requests between the other pairs of nodes. Hence a smaller number of requests for the 1-2 node pair will be satisfied and, subsequently, the rejection probability will increase.

Figure 6 shows a plot of the number of nodes versus the rejection probability. The curve generated from (10) gives

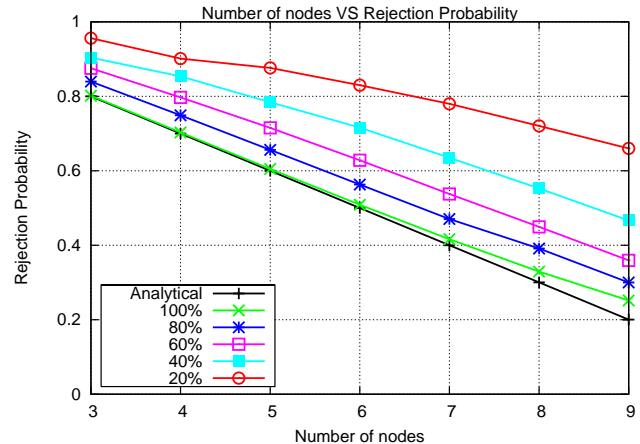


Figure 6. Analytical rejection probability as the lower bound.

the lower bound, and the remaining five curves were obtained with simulation. Simulations were performed for fully interconnected networks ranging from 3 to 9 nodes (no wavelength converters involved), and traces were generated for varying traffic loads as described above. These traces were then fed to a custom-built First Come First Serve (FCFS) optical network scheduling simulator. The FCFS simulator uses the First-Fit algorithm for wavelength assignment and Shortest-Path-First algorithm for routing [22]. Note that if we extend the simulation beyond 9 nodes, the graph will follow the observed trend.

As expected, the curve with 100% of the traffic between just one pair of nodes approaches the lower bound set by the curve obtained analytically. As the percentage of traffic between these two nodes reduces in comparison to the overall network traffic, the rejection probability increases. The analytically calculated lower bound serves as a tight lower bound on the rejection probability. Therefore, if it is known in advance that the traffic between a particular pair of nodes dominates, the rejection probability can be directly approximated from (10). In cases where the number of nodes in the network is very large and the traffic between the pair of nodes under consideration amounts to a small fraction of the total network traffic, the discrepancy between the observed rejection probability and the lower bound is large. This method of approximation should not be used in such cases.

5. Conclusion

The analysis, experiments, and results presented in this paper are aimed at complementing the ongoing wave of research in creating service oriented communication stacks for Lambda-Grids [4]. This paper explores the semantics of advance reservation requests and incorporates them in a user model. The metrics presented and the bounding con-

ditions derived help to assess and to quantify the impact of usage patterns on lightpath provisioning. We have outlined a framework for future research in the advance reservation paradigm. Toward this end, we have introduced FONTS, an effective tool for generating advance reservation request traffic.

Traces generated by FONTS not only help to overcome the paucity of information on lightpath utilization in the real world, but also help in the comprehensive validation of resource provisioning strategies by simulating complex scenarios created by the play of multiple variables in the reservation system. FONTS recognizes the switching nodes in an optical network but is independent of their interconnection. Hence it is scalable and can also be used to model network traffic in existing networks. It has been used in the DWDM-RAM project to ascertain network blocking probability results [13]. It is also being used by the authors in validating the effects of novel scheduling algorithms for lightpath provisioning [9]. In this paper we use FONTS to study the effect of traffic distribution on the rejection probability, while the effect on utilization is currently being studied. In keeping with the philosophy of exposing services through web portals, a web interface to FONTS has been developed. Traces can be configured, generated, evaluated, and downloaded via this interface [11].

Acknowledgment

The authors would like to thank Tan Trieu for developing the web interface to FONTS.

References

- [1] J. MacLaren, editor, "Advance Reservations: State of the Art", *Grid Resource Allocation Agreement Protocol, GGF*, (2.0), June 2003. Available: <http://www.fz-juelich.de/zam/RD/coop/ggf/graap/graap-wg.html>.
- [2] G. Clapp, editor, "Grid Working Draft on Network Services", *Grid High Performance Networking Research Group, GGF*, (0), Feb. 2004. Available: <http://forge.gridforum.org/projects/ghpn-wg/>.
- [3] T. Ferrari, editor, "Grid Working Draft on Network Services Usecases", *Grid High Performance Networking Research Group, GGF*, (1.1), June 2004. Available: <http://forge.gridforum.org/projects/ghpn-wg/>.
- [4] D. Simeonidou and R. Nejabati, editors, "Optical Network Infrastructure for Grid (Working Draft)", *Grid High Performance Networking Research Group, GGF*, (2), 2004. Available: <http://forge.gridforum.org/projects/ghpn-rg/>.
- [5] W. Allcock et al., "Data management and transfer in high-performance computational grid environments", *Parallel Comput.*, 28(5):749–771, 2002.
- [6] R. Boutaba et al., "Grid-Controlled Lightpaths for High Performance Grid Applications", *Journal of Grid Computing*, 1(4):387–394, 2003.
- [7] E. G. Coffman, Jr., P. Jelenkovic, and B. Poonen, "Reservation Probabilities", *Advances in Performance Analysis*, 1998.
- [8] D. Ferrari, A. Gupta, and G. Ventre. "Distributed Advance Reservation of Real-Time Connections". In *Network and Operating System Support for Digital Audio and Video*, pages 16–27, 1995.
- [9] S. Figueira et al., "Advance Reservation of Lightpaths in Optical-Network Based Grids", *ICST/IEEE Gridnets*, 2004.
- [10] S. Figueira et al., "DWDM-RAM: Enabling Grid Services with Dynamic Optical Networks", *CCGRID'04, 4th IEEE/ACM International Symposium on Cluster Computing and the Grid*, 2004.
- [11] FONTS - Flexible Optical Network Traffic Simulator. Available: <http://students.engr.scu.edu/~snaiksat/fonts/>.
- [12] L. Gommans et al., "Authorization of a QoS path based on generic AAA", *Future Gener. Comput. Syst.*, 19(6):1009–1016, 2003.
- [13] T. Lavian et al., "A Platform for Large-Scale Grid Data Service on Dynamic High-Performance Networks", *ICST/IEEE Gridnets*, 2004.
- [14] J. Mambretti et al., "The Photonic TeraStream: enabling next generation applications through intelligent optical networking at iGRID2002", *Future Generation Computer Systems*, 19(6):897–908, Aug. 2003.
- [15] National LambdaRail Architecture, 2004. Available: <http://www.nlr.net/architecture.html>.
- [16] E. Radius. "SURFnet6: Hybrid Optical and Packet Switching Infrastructure", May 2004. Available: http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn14/20040510_ER_SURFnet6.pdf.
- [17] O. Schelen and S. Pink, "Sharing Resources through Advance Reservation Agents", *Proceedings of the IFIP International Workshop on Quality of Service*, May 1997.
- [18] N. Taesombut and A. A. Chien, "Distributed Virtual Computers (DVC): Simplifying the Development of High Performance Grid Applications", *CCGRID'04, 4th IEEE/ACM International Symposium on Cluster Computing and the Grid*, 2004.
- [19] S. R. Thorpe, D. Stevenson, and G. K. Edwards, "Using Just-in-Time to Enable Optical Networking for Grids", *ICST/IEEE Gridnets*, 2004.
- [20] S. Vazhkudai and J. Schopf, "Using Regression Techniques to Predict Large Data Transfers", *International Journal of High Performance Computing Applications - Special Issue on Grid Computing: Infrastructure and Applications*, January 2003.
- [21] L. C. Wolf et al., "Issues of Reserving Resources in Advance", *Proceedings of Network and Operating System Support for Digital Audio and Video*, pages 27–37, 1995.
- [22] H. Zhang, J. P. Jue, and B. Mukherjee, "A Review of Routing and Wavelength Assignment Approaches for Wavelength-Routed Optical WDM Networks", *SPIE Optical Networks Magazine*, pages 47–60, 2000.
- [23] J. Zheng and H. T. Mouftah, "Routing and Wavelength Assignment for Advance Reservation in Wavelength-Routed WDM Optical Networks", *IEEE International Conference on Communications, ICC 2002*, 5:2722–2726, 2002.