Scheduling Advance Reservation Requests for WDM Networks with Static Traffic Demands

T. Daniel Wallace*, Abdallah Shami*, and Chadi Assi+

*Department of Electrical and Computer Engineering,
The University of Western Ontario, London, Ontario, Canada

+ Faculty of Engineering and Computer Science,
Concordia University, Montreal, Quebec, Canada

Abstract

Telecommunication and grid computing applications demand high bandwidth data channels that offer guarantees with respect to service availability. Such applications include: remote surgery, remote experimentation, video on-demand (VOD), teleconferencing, and bulk transfers. Furthermore, by forecasting traffic patterns internet service providers (ISP) attempt to optimize network resources in order to lower operational costs during peak periods of bandwidth consumption. Advance reservation for wavelength division multiplexed (WDM) networks can address some of these issues by reserving high volume communication channels (i.e., lightpaths) beforehand. In this paper we develop a mathematical model to solve the problem of scheduling lightpaths in advance. The optimal solution is presented as a mixed integer linear program (MILP) with the assumption that all traffic is static and the network is centrally controlled. Furthermore we have developed two novel meta-heuristics based on: 1) a greedy implementation (local search), and 2) simulated annealing (SA). The meta-heuristics have shown to produce good approximate solutions in a reasonable amount of time.

1 Introduction

Researchers in the fields of computer science and engineering envision grid computing as the future of real-time parallel and distributed computing. Supporting multiple regions of a country with additional reserves of electricity, a power grid is an ideal analogy for a grid network. A grid network can provide supplementary services that include, but are not limited to, data repositories, processing power, and network bandwidth. Historically the advancements made in grid computing technology have been driven primarily by the e-Science community. The definition of e-Science is generally described as a globally distributed collaboration for scientific computing enabled by the Internet [1]. Grid networks for e-Science have been established in numerous locations in North America and Europe with each being dependent on long distance backhaul communication links for information and resource sharing. With a goal of producing over 15 Petabytes of data per year to be accessed and analyzed globally, the Large Hadron Collider (LHC) Computing Grid Project, aims to build and maintain a data storage and analysis infrastructure for the worlds largest physics laboratory, the European Organization for Nuclear Research (CERN) [2]. A working example of a grid infrastructure, in this case linking three US based test-bed projects, is the Biomedical Informatics Research Network (BIRN). BIRN allows national collaborations in biomedical engineering [3], allowing the transfer of information and test-results between the engineers involved in their respective projects. Another functional grid community is the George E. Brown Network for Earthquake Engineering and Simulation (NEES) [9] program. Composed of a diverse group of individuals and organizations, NEES has a mandate to study the effects of impact and aftermath of seismic events on common societal structures such as buildings, bridges, roads, etc.

Due to the enormous volumes of data generated by e-Science applications (i.e., Terabytes and even Petabytes), efficient transportation between end-users over a shared network becomes a complex problem. For example, today a Terabyte file would require the use of a postal courier service for the timely delivery of its electronic data. This of course is due to a lack of quality of service (QoS) in the public Internet's best effort service model. Turning the attention towards a more business-oriented example, in 2002 the movie industry created the first specifications for digital cinema. It is expected that these specifications will inspire theatres to also go "all digital" (i.e., digital projectors). Unfortunately the transportation and distribution infrastructure of how a digital movie gets from the film studio to a local theatre is still under debate [21].

In this paper, we address the problem of static lightpath establishment (SLE) for connection requests that are made in advance. With SLE, all connection requests are known ahead of time and the RWA problem is performed offline [22]. Furthermore, the original SLE problem assumes that all connection requests are reserved and established on the same date and last for an undetermined length of time. However, since a WDM network can be considered a semi-dynamic system when future traffic demands are predictable (i.e., lightpath reservations provide a lifespan) the original SLE problem will not be able to efficiently optimize network resources. Therefore, we extend the original SLE problem, so that requests specify not only a source and destination but also a start time and duration. From now on we will refer to this new problem as the advance reservation static lightpath establishment (ARSLE) problem. It is our aim that the ARSLE problem will address network connectivity issues for grid applications that demand a high bandwidth connection between grid resources which use advance and/or co-allocation reservation services. For this research, we assume there are two independent levels of scheduling. While the first level schedules resources at the network layer (e.g., lightpaths), the second level schedules the end devices required by grid applications (e.g. supercomputers). We further assume that both the application layer and the network layer have independent scheduling mechanisms and therefore scheduling on either layer is not performed concurrently. Since machine scheduling is a topic that is beyond the scope of this paper, our intent is to solely focus on scheduling at the network layer. However, this does not mean that solving the duel network/machine scheduling problem is infeasible, simply that the solution requires additional knowledge regarding constraints such as device capacity and application availability. These issues will of course add more overhead to any solution as well as significantly increase the overall complexity of such a problem.

The rest of this paper is organized as follows. Section 2 presents an overview of the literature related to this work. Section 3 formally describes the ARSLE problem and provides a new MILP to solve the problem optimally. Section 4 provides a simple example problem for the purpose of explaining the MILP. Section 5 describes the meta-heuristics and algorithms used to find a good approximate solution to the ARSLE problem. Finally, Section 6 evaluates the performance of the meta-heuristics and the optimal solution, and Section 7 summarizes the work.

2 Literature Review

Allocating the resources of an optical network in advance is actually a special case of the static lightpath establishment (SLE) problem [15]. SLE is a dual problem: 1) a network design problem in which the objective is to minimize the number of lightpaths that pass through a particular link at any given time (i.e., minimize congestion); and 2) a network operation problem of maximizing the number of established connections for a fixed number of wavelengths. An integer linear program (ILP) for the original SLE problem has been formulated and shown to be NP-complete in [22] and [15]. For the design problem of minimizing the number of wavelengths needed to support a set of lightpaths, Wang et al. proposed a tabu search (TS) algorithm in [20], while Banerjee and Sharan use an evolutionary algorithm in [5] to generate an approximate solution.

Since SLE will allocate network resources to start at time zero and last for an indefinite period of time, connection requests only need to specify a source and destination. However, to apply advance reservation techniques, researchers have had to fix both a start time and duration to each connection request. This type of advance reservation problem for SLE is presented by Kuri et al. in [14] and [13] by presenting the idea of scheduled lightpath demands (SLDs). In [14] the authors formulate a mathematical model to find the optimal solution that minimizes wavelength usage for SLD in a WDM network. Furthermore, they also propose a simulated annealing (SA) and tabu search (TS) algorithm to find good approximate solutions for minimizing congested links. The authors of [16] also focus on SLD by proposing two unique algorithms, namely the independent sets algorithm (ISA) and the time window algorithm (TWA). These two algorithms are based on circular arc graph theory, where prior to routing the SLDs are divided into subsets of the time disjoint demands with the ISA, while TWA divides the SLDs into groups of time-overlapping demands.

There are some grid applications, e.g., bulk transfers, that may know in advance when and where data needs to be transmitted, but are flexible regarding the exact start time. These types of advance reservation network problems are interesting because they can be treated as standard scheduling problems, where the goal may be to minimize either the average completion time or the tardiness of a requested bulk transfer. Chen and Lee in [7] discuss flexible start times for advance reservation scheduling by attempting to optimize network resources through the release and reconfiguration of candidate reservations in a video on-demand (VOD) system. In a similar fashion, the authors of [8] propose an algorithm for scheduling advance reservations with laxities in grid networks. In this case the laxity of task or job is defined as "the difference between its deadline and the time at which it would finish executing on that resource if it starts executing at its start time" [8]. Smith et al. also study advance reservation scheduling in [17] by analyzing the performance of different job types, e.g., jobs that can and cannot be restarted.

3 Problem Formulation

In this section, we describe the ARSLE problem in a formal manner by explaining its objective while introducing the notation used to solve the optimal solution. The ARSLE problem is a scheduling problem which attempts to minimize the average tardiness (i.e., delay) of all advance connection requests. Scheduling is performed on an optical network by efficiently allocating connection requests that specify a source, destination, requested service time (RST), and duration. In this case, the start time of each request is a variable that is greater or equal to its RST. To solve the optimal solution for the ARSLE problem, we use mathematical modeling to formulate a mixed integer linear program (MILP). It should be noted that the techniques used to formulate the optimal solution have been influenced by the mathematical models developed in [19] and [15], respectively. Here, it should be noted that in [19], the authors formulate a MILP to minimize the weighted completion time a set of jobs will require when scheduled to machines within a manufacturing system, while the ILP in [15] maximizes the number of established connections that can be handled in an optical network. By utilizing both formulations together, the former provides the means to prevent connection requests from overlapping in time, while the latter addresses routing and wavelength assignment constraints.

3.1 Notation

We now define the notation used to formulate the MILP proposed to minimize the average tardiness of a connection request. It is assumed that the following parameters are given:

- J is the set of jobs (i.e., the set of connection requests).
- L is the set of links.
- W is the set of wavelengths.
- r_j is the requested service time (RST) of job j.

- d_j is the duration of job j.
- h is a very large number.
- $a_{j,p} = \begin{cases} 1, & \text{if path } p \text{ is an alternate path for job } j \\ 0, & \text{otherwise.} \end{cases}$
- $b_{l,p} = \begin{cases} 1, & \text{if link } l \text{ is on path } p \\ 0, & \text{otherwise.} \end{cases}$

3.2 MILP Formulation

The mixed integer linear program is now defined as follows:

$$\textbf{Minimze:} \quad \frac{\sum\limits_{j \in J}(S_j - r_j)}{|J|} \tag{1}$$

Subject to the following constraints:

$$S_j \geq r_j, \quad \forall j \in J \tag{2}$$

$$\sum_{p \in P}(d_j \cdot b_{l,p} \cdot X_{j,p,w} + h \cdot b_{l,p} \cdot X_{j,p,w} + h \cdot b_{l,p} \cdot X_{k,p,w}) \\ + S_j - S_k + h \cdot Y_{j,k} \leq 3h, \quad \forall j \in J, \forall k \in J, j \neq k, \forall l \in L, \forall w \in W \tag{3}$$

$$Y_{j,k} + Y_{k,j} = 1, \quad \forall j \in J, \forall k \in J, j \neq k \tag{4}$$

$$\sum_{p \in P, w \in W} X_{j,p,w} \cdot a_{j,p} = 1, \quad \forall j \in J \tag{5}$$

$$X_{j,p,w} \in \{0, 1\}, \quad \forall j \in J, \forall p \in P, \forall w \in W \tag{6}$$

$$Y_{j,k} \in \{0, 1\}, \quad \forall j \in J, \forall k \in J, j \neq k \tag{7}$$

$$r_j \geq 0, \quad \forall j \in J \tag{8}$$

$$S_j \geq 0, \quad \forall j \in J \tag{9}$$

The objective function (1) minimizes the average tardiness of a connection request. Constraint (2) ensures that a connection request does not start before its release time. The set of constraints (3) and (4) ensure that at most one request can be scheduled on a certain link and wavelength at a time. Constraint (5) says one request can only be scheduled on one wavelength and on one path and ensures that a request can only be scheduled on one of its alternate paths. While expressions (6) and (7) are binary constraints, Equations (8) and (9) specify non-negative values.

4 Example Problem

In this section, we present a simple example for the purpose of explaining and understanding the solution provided by the MILP presented in Section 2. The following example was solved using ILOG CPLEX version 9.1 by converting the MILP into an AMPL (A Modeling Language for Mathematical Programming) [10] script. First consider the small bidirectional network in Fig. 1 and the set of connection requests in Table 1. The table provides both the source and destination of each connection request as well as its RST and duration. Furthermore, it should be assumed there are 2 wavelengths per link.

Table 1: Connection Requests for Example Problem

Job_j	Source_j	Destination_j	RST_j	Duration_j
J_1	B	A	0	3
J_2	A	B	0	6
J_3	C	A	4	4
J_4	B	C	3	3
J_5	C	A	2	5
J_6	A	B	5	2
J_7	C	A	3	7
J_8	B	A	0	6

The objective of the ARSLE problem is to minimize the average tardiness

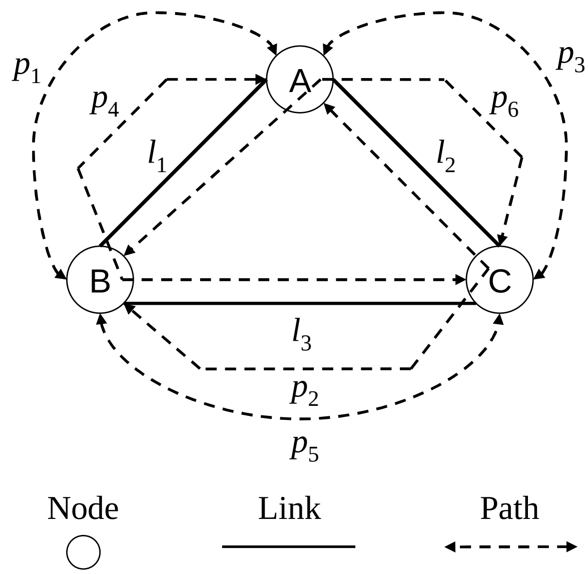

Figure 1: Network Topology for Example Problem

Table 2: Route and Wavelength Assignments Solved by MILP

Job_j	Path_j	Wavelength_j	Start Time_j	Delay_j
J_1	2(B-C-A)	1	0	0
J_2	1(A-B)	1	0	0
J_3	4(C-B-A)	1	6	2
J_4	5(B-C)	2	3	0
J_5	3(C-A)	2	2	0
J_6	1(A-B)	2	6	1
J_7	3(C-A)	1	3	0
J_8	1(B-A)	2	0	0

of a connection request. Using the topology from Fig. 1 with the set of input parameters in Table 1, the MILP returns an objective of 0.375, which means that on average each request will experience a delay in service of 0.375 time units. While Table 2 lists the lightpath assignments with start times and delay in tabular form, Fig. 2 illustrates the schedule of each link over time.

5 Approximate Solutions Using Meta-heuristics

In this section, we propose using meta-heuristics to solve the advance reservation static lightpath establishment problem. Meta-heuristics can be de-

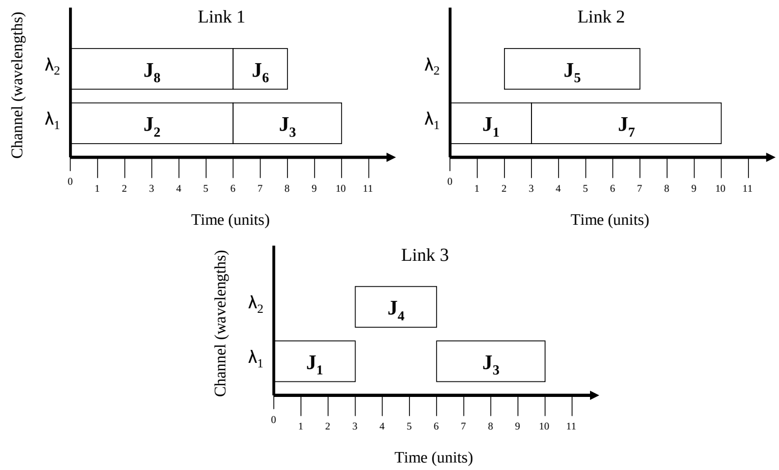

Figure 2: Visualized Link Schedule Solved by the MILP

scribed in general terms as a set of strategies that guide subordinate heuristics in such a way to produce good quality solutions [6]. Typically, meta-heuristics are applied to combinatorial optimization problems that are considered exponentially complex. Usually the optimal solution of an exponentially complex problem cannot be found in a reasonable amount of time when the data set is very large. However, it is common practice to use meta-heuristics that find good approximate solutions in a realistic amount of time, albeit there is no guarantee on optimality. In this work, we have chosen to study the impact simulated annealing (SA) will have on the ARSLE problem. Although there are a number of alternate meta-heuristics that may also be suitable for this problem, we have elected to examine SA since there is extensive literature written about its method [18, 4]. This paper should not reflect an exhaustive study on meta-heuristics with respect to the ARSLE problem.

5.1 Greedy Advance RWA Algorithm

The greedy advance routing and wavelength assignment algorithm (GARWA) is a local search on the ARSLE problem, i.e., at each stage of the algorithm the local optimum is found. The GARWA algorithm considers a set of requests in any arbitrary order, where each request specifies a source, destination, RST, and duration. When solving the ARSLE problem, the algorithm simply performs routing and wavelength assignment for k-shortest paths with w wavelengths per link in a sequential manner (i.e., it schedules each request one after another without regard for requests that exist further down the list). In this way, the algorithm will find a set of earliest start times corresponding to every path and wavelength combination for each request. The algorithm will then assign the request to the path and wavelength combination with the minimum start time in that set. It should be noted that this algorithm follows the wavelength continuity constraint, which specifies that a lightpath must use the same wavelength along all links of its route. This method of search will favour requests that are shorter in length (i.e., number of hops) from source to destination because on average fewer resources will be used when compared to requests that traverse many hops. Therefore, one could intuitively assume that this method would also favour requests which have shorter durations over those with longer ones. We present the pseudocode for the greedy ARWA algorithm is Appendix A.1.

5.2 Simulated Annealing

Simulated annealing was first suggested as a meta-heuristic algorithm for solving combinatorial optimization problems by Kirkpatrick et al. in [12]. The algorithm is based on a natural process involving the cooling of liquid solids in states of higher energy to crystalline lattices of minimal energy. In this way, during periods where the temperature is very high, the particles of a solid will drastically rearrange themselves and wander randomly through states of higher energy but, during cooler periods the particles will only make moderate changes. SA is advantageous over many iterative algorithms because it avoids being trapped within local minima by probabilistically moving to worse configurations by means of an acceptance rule which is governed by the Metropolis criterion. The Metropolis criterion specifies a nonzero probability of switching to a configuration with a higher cost and is defined in pseudocode as follows: if $\Delta C_{i,j} \leq 0$ then accept configuration change, else if $\exp(-\Delta C_{i,j}/c) > \text{rand}(0, 1]$

then accept configuration change. In the definition of the Metropolis criterion, C is the cost of a configuration and c is a control parameter, which are analogous to energy and temperature, respectively [18]. Appendix A.2 provides the pseudocode for the general simulated annealing algorithm.

To adapt SA to a new combinatorial optimization problem, it is necessary to develop the following three mechanisms: 1) a cost or objective function, 2) a perturbation function, and 3) a cooling schedule.

1) The cost function of ARSLE has been discussed in Section 3. To reiterate: the cost function is simply the objective of the problem. Therefore the cost function that minimizes the average tardiness of a request is given by Equation (1).

2) The perturbation or generation procedure is a device for making transitions from one state to another. The perturbation is more formally defined such that configuration i can generate configuration j by randomly exchanging an element in configuration i with one from the neighbourhood of i [18]. The perturbation procedure for our problem is as follows: First, choose at random job j from the set J. Next, perform the GARWA algorithm for only job j. Finally, update the start time, path, and wavelength of job j to the next earliest start time such that the new path/wavelength combination for this job is not the same as its predecessor.

3) The final mechanism is the cooling schedule, which is broken up into four parts, namely: an initial temperature, a final temperature, the length of the Markov chain, and a decrementing rule. While the initial and final temperatures are trivial conceptually, the length of the Markov chain and the decrementing rule are a bit more involved. The length of the Markov chain is simply the number of configurations or solutions the problem will explore at any given temperature, the reader is referred to [18, 4], for an in depth analysis of Markov chains. Furthermore, the decrement rule is simply how the temperature decreases. In many instances the decrement rule is a negative exponential process. For our purposes, we will evaluate two cooling schedules which have previously appeared in the literature. The first cooling schedule is based on what is referred to as a simple implementation as described in [18], while the second, proposed by Huang et al. in [11] is considered more elaborate. We have made only one modification to the cooling schedules. In both cases, the schedules use the following decrement rule

$$c_{k+1} = \alpha^k \cdot c_k, \quad k = 0, 1, 2 \ldots \tag{10}$$

Finally, it should be noted that to effectively implement the simple annealing schedule, it is necessary to describe the size of the problem. We have defined the size of the ARSLE problem by

$$N = |W| \cdot |J| \cdot p, \tag{11}$$

where W is the set of wavelengths, J is the set of requests or jobs, and p is the maximum number of alternate paths allowed per request.

6 Numerical Results

The evaluation section is broken into two parts. In the first part it is our intention to compare the performance of the proposed meta-heuristics with that of the MILP (i.e., the optimal solution) on a small network. It is necessary to evaluate the MILP using a network topology with only a few nodes and a small number of connection requests since the time to compute the optimal solution will drastically increase with size of the problem (e.g., nodes, requests, wavelengths, etc.). The second part of the evaluation analyzes the meta-heuristics on a larger scale by contrasting the performance results of each scheduling model using the NSFNET topology with additional wavelengths and more connection requests. Both SA heuristics, as well as the GARWA algorithm, were implemented in C++. The MILP was written in AMPL syntax and solved using ILOG CPLEX version 9.1. All tests were performed on a desktop PC with a Pentium 4 3GHz processor and 2GB RAM running Windows XP.

6.1 Small Problem Size

The network topology used in these simulations is a small hypothetical campus-to-campus optical network depicted in Fig. 3. It is assumed that each link is bidirectional and equipped with 2 wavelengths. For each request, our traffic model generates source and destination pairs that are uniformly distributed among all nodes. Source and destination pairs will also have 2 alternate paths to choose from. To the best of our knowledge there is no distribution curve which can effectively represent advance reservation traffic, therefore the requested service time of a request is uniformly distributed between 0 and some maximum window size. In these simulations the window size is assumed to be a length of 60 minutes. Furthermore, each request is accompanied by a holding time that follows an exponential process where the average holding time $1/\mu = 30$ minutes. Finally, the number of consecutive Markov chains with no change in cost (simple annealing schedule) is calculated as 1% of N (i.e., the problem size). Using each of the previously discussed solution methods, the average tar-

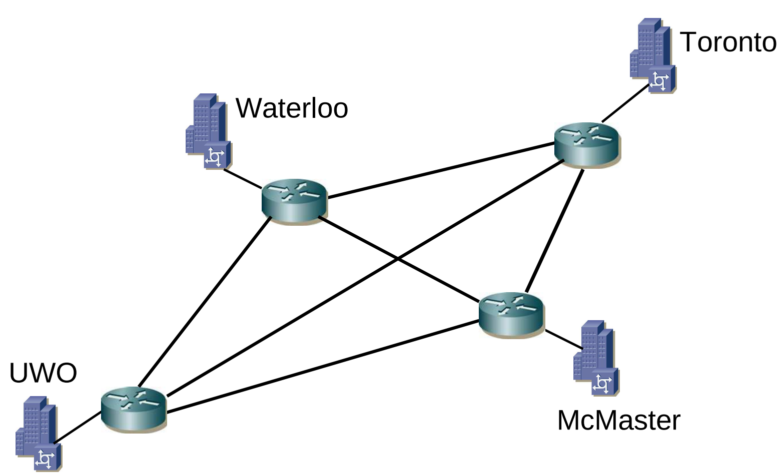

Figure 3: Hypothetical Campus-to-Campus Optical Network

diness with respect to the number of requests is shown in Fig. 4. The graph shows that the average tardiness experienced by a request will increase with the number of requests. Of course this would be anticipated since more connection requests would increase the service delay. It can be seen in the graph that both SA schedules consistently outperform the GARWA algorithm but fail to reach the optimal solution. This is expected since SA gives no guarantee on optimality, but should be able to find good approximate solutions that are better than local searches such as greedy algorithms. We refer the reader to

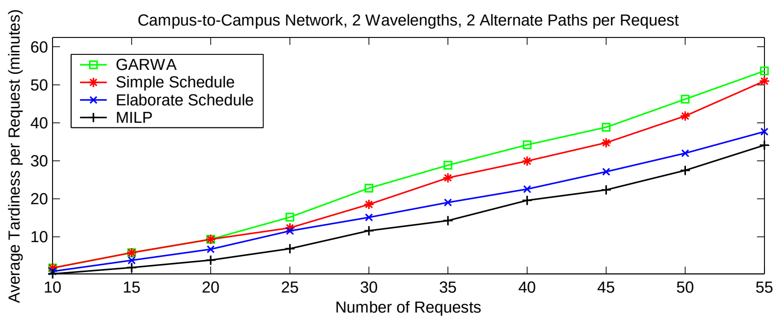

Figure 4: Average Tardiness for Small Problem

Table 3 and Table 4 for a more detailed statistical comparison of each method when the number of requests is 30 and 50, respectively. From the tables it is clear there is a significant increase in computation (solve) time when solving for the optimal solution. Therefore we can conclude that the time required to produce an approximate solution is negligible compared to solving the optimal solution. It should be noted that the results when using the GARWA heuristic have been omitted from the table because the solve times are in the order of microseconds. Finally, Fig. 5 (a) and (b) show the probability density function in terms of tardiness experienced per request for each of the proposed solution methods.

6.2 Large Problem Size

The network topology used in these simulations is the 16 node nationwide network NSFNET (Fig. 6). It should be assumed that each of the 25

Table 3: Solution Comparison with 30 Requests

Method	Avg. Solve Time (sec)	Average Tardiness (min)			
		Min	Max	Std	Mean
GARWA	N/A	15.01	33.12	6.58	22.79
Simple	0.022	10.01	27.66	6.76	18.52
Elaborate	0.106	8.37	24.06	4.79	15.14
MILP	3668.7	4.55	23.69	5.68	11.59

Table 4: Solution Comparison with 50 Requests

Method	Avg. Solve Time (sec)	Average Tardiness (min)			
		Min	Max	Std	Mean
GARWA	N/A	33.24	58.41	8.03	46.21
Simple	0.045	32.07	51.29	6.61	41.81
Elaborate	0.322	18.11	42.12	7.07	31.98
MILP	8804.7	15.41	33.32	5.50	27.65

links in NSFNET is bidirectional and is equipped with 8 wavelengths. In these simulations the maximum window size is increased to a length of 180 minutes. Furthermore, the maximum number of alternate paths per source and destination pair is increased to 5. Again, the source and destination of each request is generated randomly and uniformly distributed among all nodes. As well as each request is accompanied by a holding time that follows an exponential process where the average holding time $1/\mu = 30$ minutes. Finally, we also calculate the number of consecutive Markov chains with no change in cost (simple annealing schedule) as 1% of N (i.e., the problem size). Since it was shown in the previous section that solving for the optimal solution can be very expensive in terms of time, it will not be desirable to attempt to solve a problem with many nodes and many requests. Therefore it is our intent in this section to compare the performance of only the proposed meta-heuristics. Fig. 7 (a) shows that average tardiness with respect to the number of requests for each meta-heuristic. Again, it is evident that the elaborate annealing schedule produces superior results. However, as we can see in Fig. 7 (b), by choosing to employ the elaborate schedule we must incur a significant cost in time. This cost can be attributed to the nature of the final temperature. The final temperature for

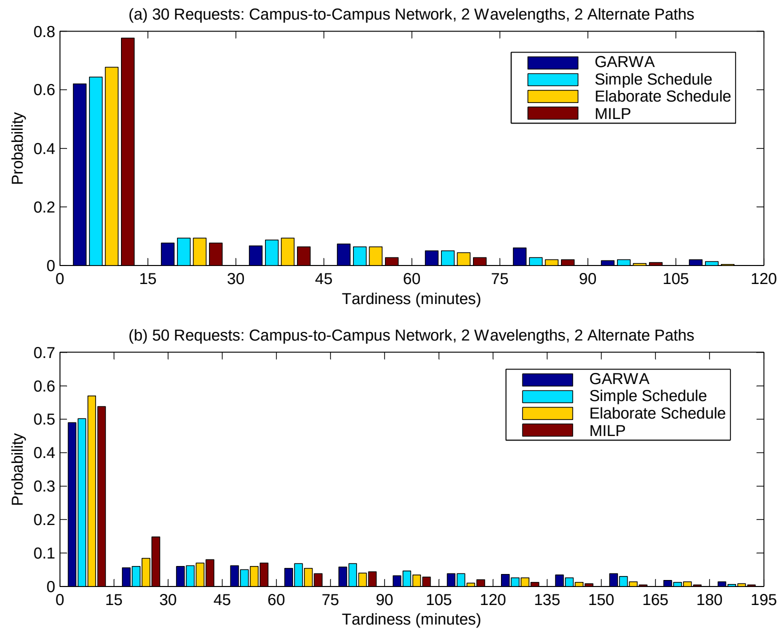

Figure 5: Probability Density Function for Small Problem

the simple schedule is determined by failing to improve on the current solution over some fixed number of Markov chains. In contrast, the elaborate schedule uses a variable number of Markov chains while it attempts to place the solution into steady state equilibrium. Finally, the probability density function in terms of delay (tardiness) experienced by a request is illustrated in Fig. 8 (a) and (b) when the number of requests is 300 and 500, respectively.

7 Conclusion

A new combinatorial optimization problem referred to as the advance reservation static lightpath establishment (ARSLE) problem has been presented. ARSLE is a scheduling problem which minimizes the average tardiness of requests. We have presented a new mixed integer linear program (MILP) to solve the ARSLE problem optimally. However, since the time to compute the optimal solution increases exponentially, we have proposed heuristics based on SA to solve good approximate solutions. Furthermore, we proposed two SA sched-

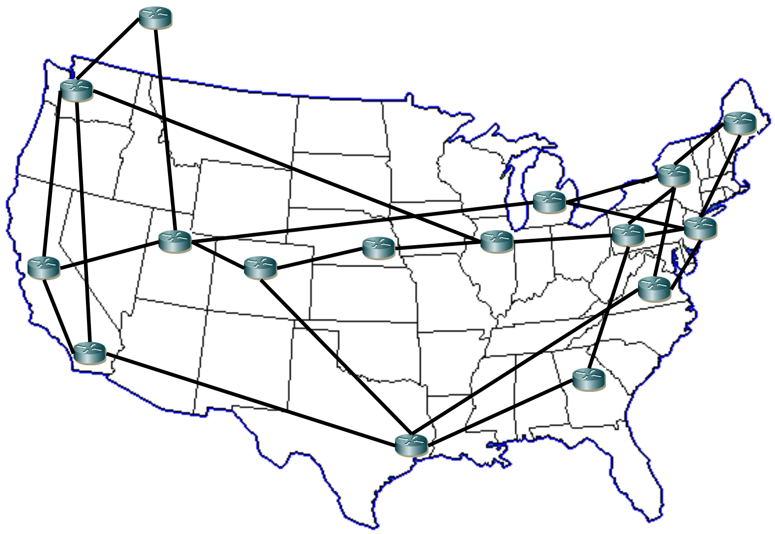

Figure 6: Nationwide Backbone Network Topology NSFNET

ules based on what the literature views as: 1) a simple schedule, and 2) an elaborate schedule. Through experimental results, it was shown that although the elaborate schedule generated solutions of better quality, the simple schedule converged on an approximate solution faster.

This paper primarily focused on the meta-heuristic referred to as simulated annealing (SA). However, there is actually a wide range of different meta-heuristics with a variety of solution strategies. Some include but are not limited to: tabu search (TS), genetic algorithms (GAs), and evolutionary algorithms (EAs). Therefore, we would recommend that future works in this area consider some of these alternative solution methods for the ARSLE problem as one may prove to be better than another.

In this work, the problem specifies that all advance connection requests are scheduled at one time. However, this problem may not address the entire spectrum of applications. For instance, in many cases, wide area networks may want to only reconfigure connections (i.e., lightpaths) on a per request bases instead of all at the same time. Therefore, we would like to propose that future work take a slightly different angle on this problem by only minimizing the tardiness of a single or new request. In this new problem, requests may now

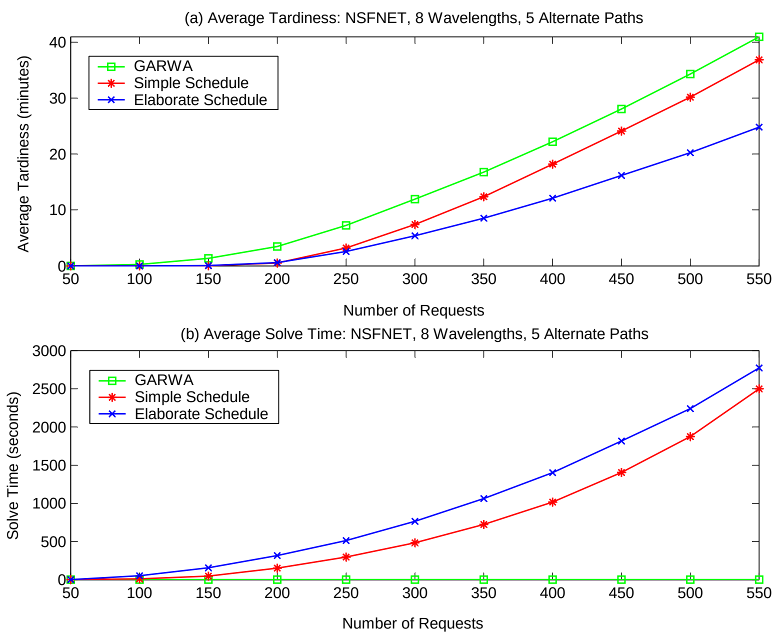

Figure 7: Average Tardiness and Solve Time for Large Problem

arrive dynamically, and the current state of the network is updated such that all reservations retain their reserved start time, but may change paths and/or wavelengths such that the tardiness of the new advance connection request is minimized.

Obviously the ARSLE problem only addresses lightpaths, i.e., optical connections that utilize the entire bandwidth of a wavelength channel across the network. Since the motivation for this work was to address applications with high bandwidth demands, it was intuitive to assume the use of optical networks for the communication technology. However, it could be expected that there will also be a demand for advance connection requests that would only require a fraction of the communication channel. It is therefore our intent to expand this research to encompass the MPLS framework. This of course would mandate the addition of a bandwidth parameter to each advance connection request so that both quality of service and efficiency is maintained while scheduling traffic engineered tunnels.

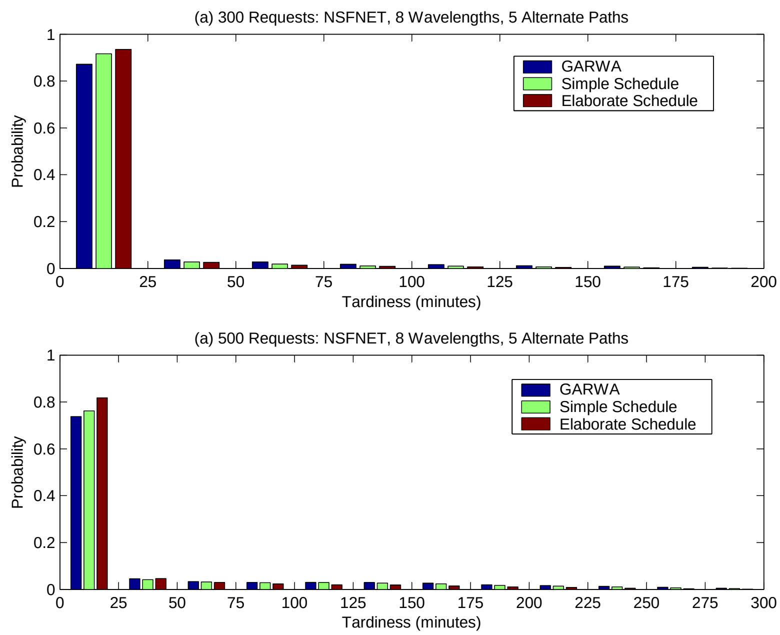

Figure 8: Probability Density Function for Large Problem

A Appendix

A.1 Greedy ARWA Pseudocode

In the pseudocode to follow, it should be assumed that the functions ISINTERSECTING and ISOVERLAPPING will return true if two separate paths use a common link and if two requests overlap in time, respectively.

for $r = 1$ to $RequestCount$ **do**
  for $p = 1$ to $r.AlternatePathCount$ **do**
    for $w = 1$ to $WavelengthCount$ **do**
      $r.StartTimes[p][w] = r.ReleaseTime$;
      $StartTimeUpdate =$ true;
      while $StartTimeUpdate =$ true **do**
        $TempStartTime = r.StartTimes[p][w]$;
        for all i such that i is a reserved request **do**
          if ISINTERSECTING$(r.AltPathSet[p], i.PathId)$
          and $r.Id \neq i.Id$ **and** $w = r.Wavelength$
          and ISOVERLAPPING$(r, i, r.StartTimes[p][w])$ **then**
            $r.StartTimes[p][w] = i.StartTime + i.Duration$;
          end if
        end for
        if $TempStartTime = r.StartTimes[p][w]$ **then**

$StartTimeUpdate = \text{false}$;
else
$TempStartTime = r.StartTimes[p][w]$;
end if
end while
end for
end for
$BestStartTime = r.StartTimes[1][1]$;
$BestPathIndex = 1$;
$BestWavelength = 1$;
for $p = 1$ to $r.AlternatePathCount$ **do**
for $w = 1$ to $WavelengthCount$ **do**
if $StartTimes[p][w] < BestStartTime$ **then**
$BestStartTime = StartTimes[p][w]$;
$BestPathIndex = p$;
$BestWavelength = w$;
end if
end for
end for
$r.Wavelength = BestWavelength$;
$r.PathId = r.AltPathSet[BestPathIndex]$;
$r.Reserved$;
end for

A.2 General Simulated Annealing Pseudocode

For the readers benefit, the generalized simulated annealing algorithm, compiled from [18, 4], is now provided in pseudocode.

$k = 0$;
$S_c =$ initial solution;
$S_b = S_c$;
$c_k =$ initial temperature;
while stopping criteria is not true **do**
$S_n = \text{PERTURB}(S_c)$;
if $C(S_n) - C(S_c) \leq 0$ **then**
accept neighbour solution;
else if $\exp(-[C(S_n) - C(S_c)]/c) < \text{rand}[0, 1)$ **then**
accept neighbour solution;
end if
if $C(S_n) < C(S_b)$ **then**
$S_b = S_n$;
end if
$c_{k+1} = \text{DECREMENTRULE}(c_k)$;
$k = k + 1$;
end while

References

[1] http://www.nesc.ac.uk/nesc/define.html, accessed May 2006.

[2] http://info.web.cern.ch/Press/PressReleases/Releases2005/PR06.05E.html, accessed May 2006.

[3] http://www.nih.gov/news/pr/dec2004/ncrr-08.htm, accessed May 2006.

[4] E. Aarts and J. Korst. *Simulated annealing and Boltzmann machines: a stochastic approach to combinatorial optimization and neural computing.* John Wiley & Sons, Inc. New York, NY, USA, 1989.

[5] N. Banerjee and S. Sharan. A evolutionary algorithm for solving the single objective static routing and wavelength assignment problem in WDM networks. *Intelligent Sensing and Information Processing, 2004. Proceedings of International Conference on*, pages 13–18, 2004.

[6] C. Blum and A. Roli. Metaheuristics in combinatorial optimization: Overview and conceptual comparison. *ACM Computing Surveys (CSUR)*, 35(3):268–308, 2003.

[7] YT Chen and KH Lee. A flexible service model for advance reservation. *Computer Networks*, 37(3-4):251–262, 2001.

[8] U. Farooq, S. Majumdar, and E.W. Parsons. Impact of Laxity on Scheduling with Advance Reservations in Grids. *MASCOTS*, 5:319–324.

[9] I. Foster and C. Kesselman. *The grid: blueprint for a new computing infrastructure.* Morgan Kaufmann Publishers Inc. San Francisco, CA, USA, 1998.

[10] R. Fourer, D.M. Gay, and B.W. Kernighan. *AMPL: A Modeling Language for Mathematical Programming.* Thomson/Brooks/Cole, 2003.

[11] MD Huang, F. Romeo, and A. Sangiovanni-Vincentelli. An efficient general cooling schedule for simulated annealing. *IEEE International Conference on Computer Aided Design*, pages 381–384, 1986.

[12] S. Kirkpatrick, CD Gelati Jr, and MP Vecchi. Optimization by Simulated Annealing. *Biology and Computation: A Physicist's Choice*, 1994.

[13] J. Kuri, N. Puech, and M. Gagnaire. Diverse routing of scheduled lightpath demands in an optical transport network. *Design of Reliable Communication Networks, 2003.(DRCN 2003). Proceedings. Fourth International Workshop on*, pages 69–76, 2003.

[14] J. Kuri, N. Puech, M. Gagnaire, E. Dotaro, and R. Douville. Routing and wavelength assignment of scheduled lightpath demands. *Selected Areas in Communications, IEEE Journal on*, 21(8):1231–1240, 2003.

[15] R. Ramaswami and K.N. Sivarajan. Routing and wavelength assignment in all-optical networks. *IEEE/ACM Transactions on Networking (TON)*, 3(5):489–500, 1995.

[16] CV Saradhi and M. Gurusamy. Graph theoretic approaches for routing and wavelength assignment of scheduled lightpath demands in WDM optical networks. *Broadband Networks, 2005 2nd International Conference on*, pages 368–376, 2005.

[17] W. Smith, I. Foster, and V. Taylor. Scheduling with advanced reservations. *Parallel and Distributed Processing Symposium, 2000. IPDPS 2000. Proceedings. 14th International*, pages 127–132, 2000.

[18] PJM Van Laarhoven and E.H.L. Aarts. *Simulated Annealing: Theory and Applications*. Springer, 1987.

[19] C. Wang, H. Ghenniwa, and W. Shen. Heuristic scheduling algorithm for flexible manufacturing systems with partially overlapping machine capabil-

ities. *Mechatronics and Automation, 2005 IEEE International Conference*, 3.

[20] Y. Wang, T.H. Cheng, and M.H. Lim. A Tabu search algorithm for static routing and wavelength assignment problem. *Communications Letters, IEEE*, 9(9):841–843, 2005.

[21] R. Wintner. Bits on the Big Screen Computer bits are replacing film in the biggest change in cinema since Technicolor. *IEEE SPECTRUM*, 43(12):36, 2006.

[22] H. Zang, J.P. Jue, and B. Mukherjee. A review of routing and wavelength assignment approaches for wavelength-routed optical WDM networks. *Optical Networks Magazine*, 1(1):47–60, 2000.